

Universality of Computational Lower Bounds for Submatrix Detection

Matthew Brennan

Massachusetts Institute of Technology. Department of EECS.

BRENNANM@MIT.EDU

Guy Bresler

Massachusetts Institute of Technology. Department of EECS.

GUY@MIT.EDU

Wasim Huleihel

Tel-Aviv University. Department of EE.

WASIM8@GMAIL.COM

Editors: Alina Beygelzimer and Daniel Hsu

Abstract

In the general submatrix detection problem, the task is to detect the presence of a small $k \times k$ submatrix with entries sampled from a distribution \mathcal{P} in an $n \times n$ matrix of samples from \mathcal{Q} . This formulation includes a number of well-studied problems, such as biclustering when \mathcal{P} and \mathcal{Q} are Gaussians and the planted dense subgraph formulation of community detection when the submatrix is a principal minor and \mathcal{P} and \mathcal{Q} are Bernoulli random variables. These problems all seem to exhibit a universal phenomenon: there is a statistical-computational gap depending on \mathcal{P} and \mathcal{Q} between the minimum k at which this task can be solved and the minimum k at which it can be solved in polynomial time.

Our main result is to tightly characterize this computational barrier as a tradeoff between k and the KL divergences between \mathcal{P} and \mathcal{Q} through average-case reductions from the planted clique conjecture. These computational lower bounds hold given mild assumptions on \mathcal{P} and \mathcal{Q} arising naturally from classical binary hypothesis testing. In particular, our results recover and generalize the planted clique lower bounds for Gaussian biclustering in [Ma and Wu \(2015\)](#); [Brennan et al. \(2018\)](#) and for the sparse and general regimes of planted dense subgraph in [Hajek et al. \(2015\)](#); [Brennan et al. \(2018\)](#). This yields the first universality principle for computational lower bounds obtained through average-case reductions.

To reduce from planted clique to submatrix detection for a specific pair \mathcal{P} and \mathcal{Q} , we introduce two techniques for average-case reductions: (1) multivariate rejection kernels which perform an algorithmic change of measure and lift to a larger submatrix while obtaining an optimal tradeoff in KL divergence, and (2) a technique for embedding adjacency matrices of graphs as principal minors in larger matrices that handles distributional issues arising from their diagonal entries and the matching row and column supports of the $k \times k$ submatrix. We suspect that these techniques have applications in average-case reductions to other problems and are likely of independent interest. We also characterize the statistical barrier in our general formulation of submatrix detection.

Keywords: statistical-computational gaps, average-case reductions, planted clique conjecture, universality, submatrix detection, community detection

1. Introduction

In the general submatrix detection problem, the task is to detect the presence of a small $k \times k$ submatrix with entries sampled from a distribution \mathcal{P} in an $n \times n$ matrix of samples from \mathcal{Q} . This problem arises in many natural contexts for specific pairs of distributions $(\mathcal{P}, \mathcal{Q})$. When \mathcal{P} and

\mathcal{Q} are Gaussians, this yields the well-studied problem of biclustering arising from applications in analyzing microarray data [Shabalin et al. \(2009\)](#). A large body of work has studied the information-theoretic lower bounds, algorithms and limitations of restricted classes of algorithms for biclustering [Butucea and Ingster \(2013\)](#); [Montanari et al. \(2015\)](#); [Shabalin et al. \(2009\)](#); [Kolar et al. \(2011\)](#); [Balakrishnan et al. \(2011\)](#); [Chen and Xu \(2016\)](#); [Cai et al. \(2017\)](#). When the $k \times k$ submatrix is a principal minor and \mathcal{P} and \mathcal{Q} are Bernoulli random variables, general submatrix detection becomes the planted dense subgraph formulation of community detection. This problem has also been studied extensively from algorithmic and information-theoretic viewpoints [Arias-Castro et al. \(2014\)](#); [Butucea and Ingster \(2013\)](#); [Verzelen et al. \(2015\)](#); [Chen and Xu \(2016\)](#); [Montanari \(2015\)](#); [Candogan and Chandrasekaran \(2018\)](#); [Hajek et al. \(2016a\)](#).

The best known algorithms for both the Gaussian and Bernoulli problems seem to exhibit a peculiar phenomenon: there appears to be a statistical-computational gap between the minimum k at which this task can be solved and the minimum k at which it can be solved in polynomial time. Tight statistical-computational gaps for both biclustering and several parameter regimes of planted dense subgraph were recently established through average-case reductions from the planted clique conjecture [Ma and Wu \(2015\)](#); [Hajek et al. \(2015\)](#); [Brennan et al. \(2018\)](#). Furthermore, the regimes in which these problems are information-theoretically impossible, statistically possible but computational hard and admit polynomial time algorithms appear to have a common structure. This raises the following natural question:

Question 1.1 *Are the statistical-computational gaps for general submatrix detection a universal phenomenon regardless of the specific pair of distributions $(\mathcal{P}, \mathcal{Q})$?*

We answer this question for a wide class of pairs of distributions $(\mathcal{P}, \mathcal{Q})$. Our main result is to tightly characterize this computational barrier as a tradeoff between k and the KL divergences between \mathcal{P} and \mathcal{Q} through average-case reductions from the planted clique conjecture. These computational lower bounds hold given mild assumptions on \mathcal{P} and \mathcal{Q} arising naturally from classical binary hypothesis testing. Our results recover and widely generalize the planted clique lower bounds for Gaussian biclustering in [Ma and Wu \(2015\)](#); [Brennan et al. \(2018\)](#) and for the sparse and general regimes of planted dense subgraph in [Hajek et al. \(2015\)](#); [Brennan et al. \(2018\)](#). This yields the first universality principle for computational lower bounds obtained through average-case reductions. We also characterize the statistical barrier in our general formulation of submatrix detection.

Average-case reductions are notoriously brittle in the sense that most natural maps designed for worst-case problems fail to faithfully map a natural distribution to a natural distribution over the target problem. A universality result obtained through an average-case reduction necessarily overcomes this barrier in a strong way as it would have to simultaneously map to an entire set of natural distributions over the target problem. A main contribution of the paper is to introduce techniques handling subtle technical obstacles that arise when devising such a reduction.

Our results are close in flavour to several previous works showing universal phenomena in the context of submatrix problems. In [Montanari \(2015\)](#), approximate message passing algorithms were shown to recover the support of the planted submatrix under regularity conditions on $(\mathcal{P}, \mathcal{Q})$. [Hajek et al. \(2016b\)](#) analyzed semidefinite programming algorithms also under regularity conditions on $(\mathcal{P}, \mathcal{Q})$. In [Hajek et al. \(2017\)](#), the information-theoretic thresholds for submatrix localization – the recovery variant of our detection problem – were shown under very mild assumptions on $(\mathcal{P}, \mathcal{Q})$, characterizing the statistical limit of the problem over a large universality class. In [Lesieur et al. \(2015\)](#); [Deshpande et al. \(2015\)](#); [Krzakala et al. \(2016\)](#), information-theoretic thresholds for

universal formulations of rank-one matrix estimation were established by analyzing approximate message passing and Guerra’s interpolation argument. Our work is the first to analyze both the information-theoretic and computational lower bounds for submatrix detection universally and the first work we are aware of producing a universality class of computational lower bounds for any problem. The conditions on $(\mathcal{P}, \mathcal{Q})$ for our computational lower bounds are mild and similar to those in [Hajek et al. \(2017\)](#).

1.1. Contributions to Techniques for Average-Case Reductions

One of our main contributions is to introduce two new techniques for average-case reductions that are of independent interest. This work is part of a growing body of literature establishing statistical-computational gaps in high-dimensional inference problems based on average-case reductions. Previous reductions include lower bounds for testing k -wise independence [Alon et al. \(2007\)](#), RIP certification [Wang et al. \(2016a\)](#); [Koiran and Zouzias \(2014\)](#), matrix completion [Chen \(2015\)](#) and sparse PCA [Berthet and Rigollet \(2013b,a\)](#); [Wang et al. \(2016b\)](#); [Gao et al. \(2017\)](#); [Brennan and Bresler \(2019\)](#). A number of techniques were introduced in [Brennan et al. \(2018\)](#) to provide the first web of average-case reductions to problems including planted independent set, planted dense subgraph, sparse spiked Wigner, sparse PCA, the subgraph stochastic block model and biclustering. More detailed surveys of this area can be found in the introduction section of [Brennan et al. \(2018\)](#) and in [Wu and Xu \(2018\)](#). In this work, we introduce the following two techniques to map from a generalization of planted clique to general submatrix detection:

- **Multivariate Rejection Kernels:** These are randomized maps that perform an algorithmic change of measure and lift to a larger submatrix while obtaining an optimal tradeoff in KL divergence that matches the target lower bounds for submatrix detection.
- **Planting Diagonals by Embedding as a Minor:** This is a technique for embedding adjacency matrices of graphs as principal minors in larger matrices. It handles distributional issues arising from missing diagonal entries in adjacency matrices and the matching row and column supports of the $k \times k$ submatrix.

The latter technique solves a central obstacle in the reductions of [Ma and Wu \(2015\)](#), [Hajek et al. \(2015\)](#), and [Brennan et al. \(2018\)](#) to biclustering and planted dense subgraph. In [Hajek et al. \(2015\)](#), it is noted that the main issue leading to the complicated analysis of their reduction arises from the missing diagonal entries in the adjacency matrix of planted clique, which on lifting get mapped to “holes” in the community. In [Brennan et al. \(2018\)](#), this same obstacle was overcome through DISTRIBUTIONAL-LIFTING, an involved technique first performing a change of measure and then iteratively lifting the resulting problem. However, neither of these reductions generalize to the universal setting with arbitrary $(\mathcal{P}, \mathcal{Q})$. Our second technique resolves this obstacle cleanly and is crucial to yielding lower bounds for arbitrary $(\mathcal{P}, \mathcal{Q})$.

1.2. Outline of the Paper

The paper is structured as follows. In Section 2, we formally define the general submatrix problem and motivate some assumptions on $(\mathcal{P}, \mathcal{Q})$ from classical binary hypothesis testing. In Section 3, we give general statements of our universality results for computational and statistical barriers. We also specialize these results to universality classes over which we can obtain complete characterizations

of the computational phase diagram for submatrix detection. In Section B, we provide preliminaries on average-case reductions in total variation. In Section C, we introduce and analyze multivariate rejection kernels. In Section D, we give our general average-case reduction TO-SUBMATRIX. In Section E, we deduce computational lower bounds from this reduction and analyze simple test statistics showing achievability of these lower bounds. In Section F, we establish the statistical limits of submatrix detection. In Section G, we discuss the strength of the assumptions giving rise to our main three universality classes UC-A, UC-B and UC-C and the distributions that they contain. In Section H, we discuss open problems remaining after this work.

2. Submatrix Problems and Conditions for Universality

2.1. General Submatrix Detection

The primary focus in this work are detection problems, wherein an algorithm is given a set of observations and tasked with distinguishing between two hypotheses:

- a *uniform* hypothesis H_0 , under which observations are generated from the natural noise distribution for the problem; and
- a *planted* hypothesis H_1 , under which observations are generated from the same noise distribution but modified by planting a latent sparse structure.

In the problems we consider, H_0 and H_1 are typically both simple hypothesis consisting of a single distribution. As discussed in Brennan et al. (2018) and Hajek et al. (2015), lower bounds for simple vs. simple hypothesis testing formulations are stronger and technically more difficult than for formulations involving composite hypotheses. For a given detection problem, the goal is to design an algorithm $\mathcal{A}(X) \in \{0, 1\}$ that classifies an input X with low asymptotic Type I+II error

$$\limsup_{n \rightarrow \infty} \{\mathbb{P}_{H_0}[\mathcal{A}(X) = 1] + \mathbb{P}_{H_1}[\mathcal{A}(X) = 0]\}$$

where n is the parameter indicating the size of X . If the asymptotic Type I+II error of \mathcal{A} is zero, then we say \mathcal{A} solves the detection problem. We now define the universal formulation of submatrix detection that will be our main object of study. Throughout this paper, $(\mathcal{P}, \mathcal{Q})$ will either denote a fixed pair of distributions over a measurable space (X, \mathcal{B}) or, when there is a natural problem parameter n , implicitly denote a pair of sequences of distributions $\mathcal{P} = (\mathcal{P}_n)$ and $\mathcal{Q} = (\mathcal{Q}_n)$.

Definition 1 (General Symmetric Index Set Submatrix Detection) *Given a pair of distributions $(\mathcal{P}, \mathcal{Q})$ over a measurable space (X, \mathcal{B}) , let $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$ denote the hypothesis testing problem with observation $M \in X^{n \times n}$ and hypotheses*

$$H_0 : M \sim \mathcal{Q}^{\otimes n \times n} \quad \text{and} \quad H_1 : M \sim \mathcal{M}(n, k, \mathcal{P}, \mathcal{Q})$$

where $\mathcal{M}(n, k, \mathcal{P}, \mathcal{Q})$ is the distribution of matrices M with entries $M_{ij} \sim \mathcal{P}$ if $i, j \in S$ and $M_{ij} \sim \mathcal{Q}$ otherwise that are conditionally independent given S , which is chosen uniformly at random over all k -subsets of $[n]$.

Similarly, asymmetric index set submatrix detection $\text{ASD}(n, k, \mathcal{P}, \mathcal{Q})$ is formulated with $M_{ij} \sim \mathcal{P}$ for all $(i, j) \in S \times T$ where both S and T are chosen independently and uniformly at random

over all k -subsets of $[n]$. Note that both information-theoretic and computational lower bounds for SSD are stronger than for ASD, since an instance of ASD can be obtained from SSD by randomly permuting its column indices. Similarly, algorithms for ASD are stronger since permuting the columns of SSD and applying an ASD blackbox yields an algorithm for SSD. In this work, we characterize the statistical and computational barriers in these and related problems through average-case reductions from the planted clique conjecture. For this to be possible, it is necessary to impose assumptions on $(\mathcal{P}, \mathcal{Q})$ so that submatrix detection is well-posed. The next section is devoted to identifying several reasonable and natural assumptions on $(\mathcal{P}, \mathcal{Q})$.

2.2. Natural Assumptions from Classical Binary Hypothesis Testing

Our objective is to examine the statistical-computational tradeoffs that arise in submatrix detection as a high-dimensional problem with hidden structure. More precisely, we aim to capture the tradeoff between the dimension k of the hidden submatrix and how distinguishable the two distributions \mathcal{P} and \mathcal{Q} are. For this to be possible, the problem of testing between \mathcal{P} and \mathcal{Q} needs to be well-posed. Consider the classical binary hypothesis testing formulation of this task with i.i.d. observations X_1, X_2, \dots, X_m where

$$H_0 : X_1, X_2, \dots, X_m \sim_{\text{i.i.d.}} \mathcal{Q} \quad \text{and} \quad H_1 : X_1, X_2, \dots, X_m \sim_{\text{i.i.d.}} \mathcal{P}$$

Note that given the latent submatrix indices S , the problem of distinguishing between H_0 and H_1 in SSD reduces exactly to this classical binary hypothesis testing task with $m = k^2$ samples.

In order to capture the tradeoffs that arise because of hidden structure in high dimensions, SSD ought to be easy to solve given S . By the Neyman-Pearson Lemma, the optimal test is a log-likelihood ratio (LLR) test that outputs H_1 if

$$\sum_{i=1}^m L(X_i) \geq m\tau \quad \text{where} \quad L(x) = \log \frac{d\mathcal{P}}{d\mathcal{Q}}(x)$$

for some threshold $m\tau$ and where $L : X \rightarrow \mathbb{R}$ is the LLR or logarithm of the Radon-Nikodym derivative between \mathcal{P} and \mathcal{Q} . We assume that L can be computed efficiently so this test is computationally feasible. Note that if \mathcal{Q} and \mathcal{P} are not close to mutually absolutely continuous distributions in total variation, then there would be a non-negligible probability of seeing samples from one not in the support of the other. We assume they are exactly mutually absolutely continuous for simplicity and so L is well-defined. We also assume that the expectations of the LLR with respect to \mathcal{P} and \mathcal{Q} are finite, or in other words that $d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})$ and $d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})$ are finite.

The error of this test resolves to the tails of the distribution of the LLR under each of \mathcal{P} and \mathcal{Q} at the threshold $m\tau$. Standard Chernoff bounds on these tails yield that if $\tau \in [-d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}), d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})]$,

$$\begin{aligned} \mathbb{P}_{H_0} \left[\sum_{i=1}^m L(X_i) \geq m\tau \right] &\leq \exp(-m \cdot E_{\mathcal{Q}}(\tau)) \\ \mathbb{P}_{H_1} \left[\sum_{i=1}^m L(X_i) < m\tau \right] &\leq \exp(-m \cdot E_{\mathcal{P}}(\tau)) \end{aligned}$$

Here, the Chernoff exponents $E_{\mathcal{P}}, E_{\mathcal{Q}} : \mathbb{R} \rightarrow [-\infty, \infty)$ are the Legendre transforms of the log-moment generating functions

$$E_{\mathcal{Q}}(\tau) = \sup_{\lambda \in \mathbb{R}} \lambda\tau - \psi_{\mathcal{Q}}(\lambda) \quad \text{and} \quad E_{\mathcal{P}}(\tau) = \sup_{\lambda \in \mathbb{R}} \lambda\tau - \psi_{\mathcal{P}}(\lambda)$$

where $\psi_{\mathcal{Q}}(\tau) = \log \mathbb{E}_{\mathcal{Q}}[\exp(\lambda L)]$ and $\psi_{\mathcal{P}}(\tau) = \log \mathbb{E}_{\mathcal{P}}[\exp(\lambda L)]$. Observe that $\psi_{\mathcal{P}}(\lambda) = \psi_{\mathcal{Q}}(\lambda + 1)$ and thus $E_{\mathcal{Q}}(\tau) + \tau = E_{\mathcal{P}}(\tau)$. Note that $\psi'_{\mathcal{Q}}(0) = -d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})$. It is well known that $\lambda\tau - \psi_{\mathcal{Q}}(\lambda)$ is concave and has derivative at zero given by $\tau + d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})$. This implies that the maximizer λ^* to the concave optimization $E_{\mathcal{Q}}(\tau)$ can be taken to be nonnegative if $\tau \geq -d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})$. The same is true for $E_{\mathcal{P}}(\tau)$ if $\tau \leq d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})$, justifying the Chernoff bounds above. We also remark that $E_{\mathcal{Q}}$ and $E_{\mathcal{P}}$ are nonnegative convex functions and are minimized at $E_{\mathcal{Q}}(-d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})) = E_{\mathcal{P}}(d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})) = 0$.

Note that statements of the form $E_{\mathcal{Q}}(\tau) \geq \beta$ or $E_{\mathcal{P}}(\tau) \geq \beta$ correspond to large deviation principles (LDP) for the LLR under \mathcal{Q} and \mathcal{P} . Our main contribution is show that the tradeoff between k and the KL divergence between \mathcal{P} and \mathcal{Q} dictates the computational barrier for submatrix detection as long as the LLR has LDPs under \mathcal{Q} and \mathcal{P} . We devise an average-case reduction to show this given the planted clique conjecture. Our results are described in more detail in the next section. The natural assumptions arising from the discussion above are summarized in the following definition of computable pairs $(\mathcal{P}, \mathcal{Q})$, that we will adopt throughout the rest of the paper.

Definition 2 (Computable Pair of Distributions) *Define a pair of sequences of distributions $(\mathcal{P}, \mathcal{Q})$ over a measurable space (X, \mathcal{B}) where $\mathcal{P} = (\mathcal{P}_n)$ and $\mathcal{Q} = (\mathcal{Q}_n)$ to be computable if:*

1. *there is an oracle producing a sample from \mathcal{Q}_n in $\text{poly}(n)$ time;*
2. *\mathcal{P}_n and \mathcal{Q}_n are mutually absolutely continuous and the likelihood ratio satisfies*

$$\mathbb{E}_{x \sim \mathcal{Q}_n} \left[\frac{d\mathcal{P}_n}{d\mathcal{Q}_n}(x) \right] = \mathbb{E}_{x \sim \mathcal{P}_n} \left[\left(\frac{d\mathcal{P}_n}{d\mathcal{Q}_n}(x) \right)^{-1} \right] = 1$$

where $\frac{d\mathcal{P}_n}{d\mathcal{Q}_n}$ is the Radon-Nikodym derivative.

3. *the KL divergences $d_{\text{KL}}(\mathcal{P}_n \parallel \mathcal{Q}_n)$ and $d_{\text{KL}}(\mathcal{Q}_n \parallel \mathcal{P}_n)$ are both finite; and*
4. *there is an oracle computing $\frac{d\mathcal{P}_n}{d\mathcal{Q}_n}(x)$ in $\text{poly}(n)$ time for each $x \in X$.*

We assume that algorithms solving our submatrix detection problems have access to these oracles and to the values $d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})$ and $d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})$. The oracles appearing in this definition can be viewed as part of the computational model that we adopt. In particular, when these oracles can be implemented in computational models such as **BPP**, so can our reductions. We remark that the assumptions and discussion in this section are similar to the setup in [Hajek et al. \(2017\)](#), which showed universality of information-theoretic lower bounds for submatrix recovery. See Sections 2.1 and 3 in [Hajek et al. \(2017\)](#) for further discussion of related assumptions on \mathcal{P} and \mathcal{Q} .

3. Summary of Results

Our main result is an average-case reduction from planted clique showing a computational lower bound for submatrix detection in terms of KL divergence when the LLR has LDPs under \mathcal{Q} and \mathcal{P} . We now briefly define the planted clique and planted dense subgraph problems as well as the planted clique and planted dense subgraph conjectures.

The planted dense subgraph problem $\text{PDS}(n, k, p, q)$ with edge densities $0 < q < p \leq 1$ is the hypothesis testing problem between

$$H_0 : G \sim \mathcal{G}(n, q) \quad \text{and} \quad H_1 : G \sim \mathcal{G}(n, k, p, q)$$

where $\mathcal{G}(n, q)$ denotes an Erdős-Rényi random graph with edge probability q . Here, $\mathcal{G}(n, k, p, q)$ denotes the random graph formed by sampling $\mathcal{G}(n, q)$ and replacing the induced graph on a subset S of size k chosen uniformly at random with a sample from $\mathcal{G}(k, p)$. The planted clique problem $\text{PC}(n, k, p)$ is then $\text{PDS}(n, k, 1, p)$. There are many polynomial-time algorithms in the literature for finding the planted clique in $\mathcal{G}(n, k, p)$, including approximate message passing, semidefinite programming, nuclear norm minimization and several combinatorial approaches [Feige and Krauthgamer \(2000\)](#); [McSherry \(2001\)](#); [Feige and Ron \(2010\)](#); [Ames and Vavasis \(2011\)](#); [Dekel et al. \(2014\)](#); [Deshpande and Montanari \(2015\)](#); [Chen and Xu \(2016\)](#). All of these algorithms require that $k = \Omega(\sqrt{n})$ if p is constant, despite the fact that the planted clique can be found by exhaustive search as soon as k is larger than $2 \log_{1/p} n$. This leads to the following conjecture.

Conjecture 3 (PC Conjecture) *Fix some constant $p \in (0, 1)$. Suppose that $\{\mathcal{A}_n\}$ is a sequence of randomized polynomial time algorithms $\mathcal{A}_n : \mathcal{G}_n \rightarrow \{0, 1\}$ and k_n is a sequence of positive integers satisfying that $\limsup_{n \rightarrow \infty} \log_n k_n < \frac{1}{2}$. Then if G is an instance of $\text{PC}(n, k, p)$, it holds that*

$$\liminf_{n \rightarrow \infty} (\mathbb{P}_{H_0} [\mathcal{A}_n(G) = 1] + \mathbb{P}_{H_1} [\mathcal{A}_n(G) = 0]) \geq 1.$$

The PC Conjecture can be seen, through a simple reduction erasing random edges, to imply a similar barrier at $k = o(\sqrt{n})$ for $\text{PDS}(n, k, p, q)$ if $0 < q < p \leq 1$ are constants. We refer to this as the PDS Conjecture. We can now state our main computational lower bounds. Judging the quality of these bounds is easy in cases where they can be achieved by efficient algorithms. We describe a set of universality classes of $(\mathcal{P}, \mathcal{Q})$ for which this is true in [Section 3.1](#).

Theorem 4 (Main Computational Lower Bounds) *Let $p \in (0, 1)$ be a fixed constant and $(\mathcal{P}, \mathcal{Q})$ be a computable pair over (X, \mathcal{B}) such that either:*

- $k = \Omega(\sqrt{n})$ and $\frac{k^4}{n^2} \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \rightarrow 0$ and the LLR between $(\mathcal{P}, \mathcal{Q})$ satisfies the LDP

$$E_{\mathcal{P}}(m) \geq \omega(m \log n)$$

for some positive m with $d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \leq m = o(n^2/k^4)$

- $k = o(\sqrt{n})$ and $d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) < \log p^{-1}$ and the LLR between $(\mathcal{P}, \mathcal{Q})$ satisfies the LDP

$$E_{\mathcal{P}}(\log p^{-1}) \geq 2 \log n + \omega(1)$$

Then assuming the PC conjecture at density p , there is no randomized polynomial time algorithm solving $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$ with asymptotic Type I+II error less than one.

This is the simplest theorem statement of our lower bounds. A more general computational lower bound starting from the PDS Conjecture and including the heteroskedastic formulation of submatrix detection with different pairs $(\mathcal{P}_{ij}, \mathcal{Q}_{ij})$ at each entry is stated in [Section E](#). We also give an alternative version requiring weaker bounds on $E_{\mathcal{P}}$ and showing a slightly weaker formulation of the same computational lower bounds. In addition to computational lower bounds, we show information-theoretic lower bounds and give inefficient and polynomial time tests providing upper bounds at the barriers in submatrix detection. In [Sections E.2](#) and [F.2](#), we give general statements of these results given lower bounds on $E_{\mathcal{P}}$ and $E_{\mathcal{Q}}$ similar to the conditions in the theorem above.




Figure 1: Computational and statistical barriers in $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$ where $(\mathcal{P}, \mathcal{Q})$ is in UC-A, UC-B and UC-C with $k = \tilde{\Theta}(n^\beta)$ and $d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}) = \tilde{\Theta}(n^{-\alpha})$.

3.1. Universality Class with a Complete Phase Diagram

We now outline several assumptions on a computable pair $(\mathcal{P}, \mathcal{Q})$ that allow our results to completely characterize the computational phase diagram for $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$, by providing lower bounds on $E_{\mathcal{P}}$ and $E_{\mathcal{Q}}$. The first class we consider is a universality class that allows our average-case reduction to show lower bounds for submatrix detection.

Definition 5 (Universality Class UC-A) Define $(\mathcal{P}, \mathcal{Q})$ to be in the universality class UC-A if $(\mathcal{P}, \mathcal{Q})$ is computable and for any fixed $\epsilon \in (0, 1)$, it holds that

$$E_{\mathcal{P}}(n^\epsilon \cdot d_{\text{KL}}(\mathcal{P}_n \| \mathcal{Q}_n)) = \Omega(n^\epsilon \cdot d_{\text{KL}}(\mathcal{P}_n \| \mathcal{Q}_n) \cdot \log n)$$

The next universality class UC-B that we consider ensures the simple test statistics introduced in Sections E.2 and F.2 show achievability at the computational and statistical barriers. This class is introduced as Assumption 2 in Hajek et al. (2017) and is weaker than sub-Gaussianity of the LLR.

Definition 6 (Universality Class UC-B) Define $(\mathcal{P}, \mathcal{Q})$ to be in the universality class UC-B if $(\mathcal{P}, \mathcal{Q})$ is computable and there is a constant $C \geq 1$ such that

$$\begin{aligned} \psi_{\mathcal{P}}(\lambda) - d_{\text{KL}}(\mathcal{P} \| \mathcal{Q}) \cdot \lambda &\leq C \cdot d_{\text{KL}}(\mathcal{P} \| \mathcal{Q}) \cdot \lambda^2 \quad \text{for all } \lambda \in [-1, 0] \\ \psi_{\mathcal{Q}}(\lambda) + d_{\text{KL}}(\mathcal{Q} \| \mathcal{P}) \cdot \lambda &\leq C \cdot d_{\text{KL}}(\mathcal{Q} \| \mathcal{P}) \cdot \lambda^2 \quad \text{for all } \lambda \in [-1, 1] \end{aligned}$$

Our last universality class ensures that the information-theoretic lower bound that we show in Section F.1 matches the upper bound from Section F.2. Let $d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}) = d_{\text{KL}}(\mathcal{P} \| \mathcal{Q}) + d_{\text{KL}}(\mathcal{Q} \| \mathcal{P})$ denote symmetric KL divergence.

Definition 7 (Universality Class UC-C) Define $(\mathcal{P}, \mathcal{Q})$ to be in the universality class UC-C if $(\mathcal{P}, \mathcal{Q})$ there is a constant $C' > 0$ such that $\chi^2(\mathcal{P} \| \mathcal{Q}) \leq C' \cdot d_{\text{SKL}}(\mathcal{P}, \mathcal{Q})$.

In Sections E and F, we specialize our theorems on information-theoretic and computational upper and lower bounds to these three universality classes. This yields the following characterization of the computational phase diagram for $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$ when $(\mathcal{P}, \mathcal{Q})$ is in UC-A, UC-B and UC-C. These regimes are depicted in Figure 1. Here, \ll hides factors that are sub-polynomial n .

Theorem 8 (Submatrix Detection Phase Diagram) *The regions of the computational phase diagram in $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$ are:*

- (Statistically Impossible) if $(\mathcal{P}, \mathcal{Q})$ is in UC-C then SSD is impossible if

$$d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}) \ll \frac{1}{k} \wedge \frac{n^2}{k^4}$$

- (PC-Hard) if $(\mathcal{P}, \mathcal{Q})$ is in UC-A then SSD is PC-hard but possible if

$$\frac{1}{k} \wedge \frac{n^2}{k^4} \ll d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}) \ll \frac{n^2}{k^4} \wedge 1$$

- (Polynomial Time Algorithms) if $(\mathcal{P}, \mathcal{Q})$ is in UC-B then SSD can be solved in $\text{poly}(n)$ time if

$$\frac{n^2}{k^4} \wedge 1 \ll d_{\text{SKL}}(\mathcal{P}, \mathcal{Q})$$

In Section G, we discuss the three universality classes UC-A, UC-B and UC-C and sub-classes of distributions that they contain. For example, these three classes contain the following pairs $(\mathcal{P}, \mathcal{Q})$:

- Pairs $(\mathcal{P}, \mathcal{Q})$ with sub-Gaussian LLR i.e. with $L = \log \frac{d\mathcal{P}}{d\mathcal{Q}}(x)$ for $x \sim \mathcal{Q}$ is sub-Gaussian.
- Pairs $(\mathcal{P}, \mathcal{Q})$ with bounded LLR i.e. with $L = \log \frac{d\mathcal{P}}{d\mathcal{Q}}(x)$ for $x \sim \mathcal{Q}$ is bounded almost surely.
- A wide variety of other pairs $(\mathcal{P}, \mathcal{Q})$ from a common exponential family. The computations in Appendix B of Hajek et al. (2017) provide a simple method for determining if pairs $(\mathcal{P}, \mathcal{Q})$ from an exponential family are in UC-B. This same method can be applied to check membership in UC-A and UC-C.

The class UC-B is discussed at length in Sections 2.1, 3 and Appendix B of Hajek et al. (2017). In Section G, we observe that these properties imply that the following three important computable pairs are in all three of the classes UC-A, UC-B and UC-C:

- (\mathcal{D}_{BC}) $\mathcal{P} = \mathcal{N}(\mu, 1)$ and $\mathcal{Q} = \mathcal{N}(0, 1)$ where $\mu = n^{-\alpha}$ for some $\alpha > 0$, in which case SSD corresponds to Gaussian biclustering;
- (\mathcal{D}_{SP}) $\mathcal{P} = \text{Bern}(p)$ and $\mathcal{Q} = \text{Bern}(q)$ where $p = cq = cn^{-\alpha}$ for some constant $c > 1$ and $\alpha > 0$, in which case the above diagonal entries of SSD are the adjacency matrix of an instance of sparse planted dense subgraph; and
- (\mathcal{D}_{GP}) $\mathcal{P} = \text{Bern}(p)$ and $\mathcal{Q} = \text{Bern}(q)$ where $p = q + \Theta(n^{-\gamma})$ and $q = n^{-\alpha}$ for some constants $\gamma > \alpha > 0$, in which case the above diagonal entries of SSD are the adjacency matrix of an instance of general planted dense subgraph.

The fact that these three examples are in our universality classes implies that our average-case reduction recovers and generalizes the planted clique lower bounds for Gaussian biclustering in Ma and Wu (2015); Brennan et al. (2018) and for the sparse and general regimes of planted dense subgraph in Hajek et al. (2015); Brennan et al. (2018). For the two graph problems, this is achieved by constructing a graph from the above diagonal terms of the matrix output by the reduction.

Algorithm MRK(B)

Parameters: Input $B \in \{0, 1\}$, parameter n , number of iterations N , dimension $\ell = \text{poly}(n)$, Bernoulli probabilities $0 < q < p \leq 1$ and ℓ pairs of computable sequences of distributions $(\mathcal{P}^i, \mathcal{Q}^i)$ for $i \in \{1, 2, \dots, \ell\}$ over the measurable space (X, \mathcal{B})

1. Initialize $z = (z_1, z_2, \dots, z_\ell)$ arbitrarily in the support of $\mathcal{Q}_n^1 \otimes \mathcal{Q}_n^2 \otimes \dots \otimes \mathcal{Q}_n^\ell$
2. Until z is set or N iterations have elapsed:
 - (1) Form $z' = (z'_1, z'_2, \dots, z'_\ell)$ where $z'_i \sim \mathcal{Q}_n^i$ is sampled independently for each $i \in \{1, 2, \dots, \ell\}$ and compute the log-likelihood ratio

$$L_n(z') = \sum_{i=1}^{\ell} \log \frac{d\mathcal{P}_n^i}{d\mathcal{Q}_n^i}(z'_i)$$

- (2) Proceed to the next iteration if it does not hold that

$$\log \left(\frac{1-p}{1-q} \right) \leq L_n(z') \leq \log \left(\frac{p}{q} \right)$$

- (3) If $B = 0$, then set $z \leftarrow z'$ with probability $1 - \frac{q}{p} \cdot \exp(L_n(z'))$
- (4) If $B = 1$, then set $z \leftarrow z'$ with probability $\frac{q}{p} \cdot \exp(L_n(z')) - \frac{q(1-p)}{p(1-q)}$

3. Output z

Figure 2: Multivariate rejection kernel algorithm with Bernoulli input $B \in \{0, 1\}$.

4. Overview of Average-Case Reduction Techniques

In this section, we give a preview of multivariate rejection kernels as well as our reduction from planted clique to general submatrix detection. Some of the material in this section is repeated in the Appendix for clarity. We provide formal theorem statements with guarantees for our reductions in Sections C and D.

4.1. Multivariate Rejection Kernels

The first average-case technique we introduce is multivariate rejection kernels, which effectively perform an algorithmic change of measure simultaneously with a lift. These are the maps $\text{MRK}(B)$, shown in Figure 2, that send a binary input $B \in \{0, 1\}$ to a higher dimensional space X^ℓ simultaneously satisfying two Markov transition properties:

1. if $B \sim \text{Bern}(p)$, then $\text{MRK}(B)$ is close to $\mathcal{P}_n^1 \otimes \mathcal{P}_n^2 \otimes \dots \otimes \mathcal{P}_n^\ell$ in total variation; and
2. if $B \sim \text{Bern}(q)$, then $\text{MRK}(B)$ is close to $\mathcal{Q}_n^1 \otimes \mathcal{Q}_n^2 \otimes \dots \otimes \mathcal{Q}_n^\ell$ in total variation.

where $0 < q < p \leq 1$ are fixed and $(\mathcal{P}^i, \mathcal{Q}^i)$ are pairs of computable sequences of distributions over a common measurable space (X, \mathcal{B}) for $i \in \{1, 2, \dots, \ell\}$. The maps MRK will be a key part of our reduction to SSD, lifting planted dense subgraph and planted clique instances to instances of submatrix detection with independent entries and the correct marginals.

These maps MRK use rejection sampling to transform the distribution of the input from either one of two input Bernoulli distributions to two ℓ -fold product distributions, without knowing which of the two distributions was the input. Multivariate rejection kernels are closely related to previous average-case reduction techniques in the literature. Univariate Markov transitions taking $\text{Bern}(1)$ and $\text{Bern}(q)$ to either a fixed pair of Bernoulli distributions or a fixed pair of normal distributions were introduced in [Ma and Wu \(2015\)](#) and [Gao et al. \(2017\)](#). More closely related to our map are the rejection kernel framework introduced in [Brennan et al. \(2018\)](#) and the reduction from planted clique to planted dense subgraph in [Hajek et al. \(2015\)](#). The rejection kernels in [Brennan et al. \(2018\)](#) used a similar rejection sampling scheme as we use here to map binary inputs to distributions on \mathbb{R} . In [Brennan et al. \(2018\)](#), it is assumed that both of the target distributions f_X and g_X have explicit density or mass functions and have sampling oracles. In contrast, our map MRK only requires the Radon-Nikodym derivatives $\frac{d\mathcal{P}_n^i}{d\mathcal{Q}_n^i}$ exist and that there is a sampling oracle for the target noise distributions \mathcal{Q}_n^i rather than for \mathcal{P}_n^i as well, keeping the assumptions on $(\mathcal{P}^i, \mathcal{Q}^i)$ relatively minimal.

All of the univariate maps in [Ma and Wu \(2015\)](#), [Gao et al. \(2017\)](#) and [Hajek et al. \(2015\)](#) rely on sampling explicit mass or density functions that are linear combinations of the target distributions. This requires time at least the size of the support of the target distribution or of a sufficiently fine net of its support. In the multivariate case mapping to distributions on X^ℓ , the size of this support grows exponentially in ℓ and this sampling scheme is not feasible in polynomial time. The rejection sampling scheme used here circumvents this issue and is crucial in our reduction to general submatrix detection.

4.2. Average-Case Reduction to Submatrix Detection

Figure 3 shows the pseudocode for our main average-case reduction TO-SUBMATRIX from planted dense subgraph to general submatrix detection. As stated, our reduction produces an instance of the more general problem of heteroskedastic submatrix detection, where each entry is allowed to have a different computable pair $(\mathcal{P}_{ij}, \mathcal{Q}_{ij})$. It also begins with an instance of planted dense subgraph, providing a more general class of reductions than just from planted clique. The additional submatrix problems this implies hardness for are discussed in Section E.

As shown, this reduction begins by cloning the adjacency matrix of a planted dense subgraph instance to produce two independent samples of the above diagonal portions of their adjacency matrix. These are then embedded as two halves of a principal minor in a larger matrix in Step 2. This random embedding hides the previously missing diagonal terms in total variation. Analyzing this step is one of our technical contributions. The resulting Bernoulli submatrix problem is then lifted using MRK maps to an instance of heteroskedastic submatrix detection.

This reduction succeeds in mapping to the target submatrix problem as long as the pairs of planted and noise distributions $(\mathcal{P}_{ij}, \mathcal{Q}_{ij})$ are computable and satisfy an LDP of their LLR. Our reduction implies the computational lower bounds for Gaussian biclustering and planted dense subgraph in [Ma and Wu \(2015\)](#), [Hajek et al. \(2015\)](#) and [Brennan et al. \(2018\)](#). The step of embedding as a principal minor in Step 2 circumvents the arguments in [Hajek et al. \(2015\)](#) and [Brennan et al. \(2018\)](#) showing that the total variation error introduced by missing diagonal entries is small. It also

Algorithm TO-SUBMATRIX

Inputs: Graph $G \in \mathcal{G}_n$, parameters $0 < q < p \leq 1$, intermediate dimension N , expansion factor ℓ , number of iterations N_{it} , subgraph size k , pairs of computable sequences of distributions $(\mathcal{P}_{ij}, \mathcal{Q}_{ij})$ for $1 \leq i, j \leq N\ell$ over the measurable space (X, \mathcal{B})

1. Apply GRAPH-CLONE to G with edge probabilities $P = p$ and $Q = 1 - \sqrt{(1-p)(1-q)} + \mathbf{1}_{\{p=1\}}(\sqrt{q} - 1)$ and $t = 2$ clones to obtain (G_1, G_2)
2. Sample $s_1 \sim \text{Bin}(n, p)$, $s_2 \sim \text{Bin}(N, Q)$ and a set $S \subseteq [N]$ with $|S| = n$ uniformly at random. Sample $T_1 \subseteq S$ and $T_2 \subseteq [N] \setminus S$ with $|T_1| = s_1$ and $|T_2| = \max\{s_2 - s_1, 0\}$ uniformly at random. Now form the matrix $M^1 \in \{0, 1\}^{N \times N}$ where

$$M_{ij}^1 = \begin{cases} \mathbf{1}_{\{\pi(i), \pi(j)\} \in E(G_1)} & \text{if } i < j \text{ and } i, j \in S \\ \mathbf{1}_{\{\pi(i), \pi(j)\} \in E(G_2)} & \text{if } i > j \text{ and } i, j \in S \\ \mathbf{1}_{\{i \in T_1\}} & \text{if } i = j \text{ and } i, j \in S \\ \mathbf{1}_{\{i \in T_2\}} & \text{if } i = j \text{ and } i, j \notin S \\ \sim_{\text{i.i.d.}} \text{Bern}(Q) & \text{if } i \neq j \text{ and } i \notin S \text{ or } j \notin S \end{cases}$$

where $\pi : S \rightarrow [n]$ is a bijection chosen uniformly at random

3. Let τ be a random permutation of $[N\ell]$ and form the matrix $M^2 \in X^{N\ell \times N\ell}$ by setting

$$(M_{ij}^2 : s\ell + 1 \leq \tau^{-1}(i) \leq (s+1)\ell \wedge t\ell + 1 \leq \tau^{-1}(j) \leq (t+1)\ell) = \text{MRK}_{st} \left(M_{(s+1)(t+1)}^1 \right)$$

for each $0 \leq s, t < N$, where MRK_{st} is the multivariate rejection kernel with N_{it} iterations sending $\text{Bern}(p)$ and $\text{Bern}(Q)$ to the random set of ℓ^2 pairs $(\mathcal{P}_{ij}^n, \mathcal{Q}_{ij}^n)$ satisfying that $s\ell + 1 \leq \tau^{-1}(i) \leq (s+1)\ell$ and $t\ell + 1 \leq \tau^{-1}(j) \leq (t+1)\ell$

4. Output the matrix M^2

Figure 3: Reduction TO-SUBMATRIX for showing universal computational lower bounds for submatrix problems based on the hardness of planted clique or planted dense subgraph. The pseudocode for the subroutine GRAPH-CLONE is shown in Figure 6 in Section D.

simplifies the reductions PC-LIFTING and DISTRIBUTIONAL-LIFTING in Brennan et al. (2018). Our reduction extends the Gaussian biclustering lower bounds in Ma and Wu (2015) and Brennan et al. (2018) to the cases of symmetric indices.

Acknowledgements

We thank Philippe Rigollet, Yury Polyanskiy, Ankur Moitra, Elchanan Mossel, Jonathan Weed, Frederic Koehler, Enric Boix Adserà, Austin J. Stromme, Vishesh Jain and Yash Deshpande for inspiring discussions on related topics. This work was supported in part by the grant ONR N00014-17-1-2147.

References

- Louigi Addario-Berry, Nicolas Broutin, Luc Devroye, and Gábor Lugosi. On combinatorial testing problems. *The Annals of Statistics*, 38(5):3063–3092, 2010.
- Noga Alon, Alexandr Andoni, Tali Kaufman, Kevin Matulef, Ronitt Rubinfeld, and Ning Xie. Testing k -wise and almost k -wise independence. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 496–505. ACM, 2007.
- Brendan PW Ames and Stephen A Vavasis. Nuclear norm minimization for the planted clique and biclique problems. *Mathematical programming*, 129(1):69–89, 2011.
- Ery Arias-Castro, Nicolas Verzelen, et al. Community detection in dense random networks. *The Annals of Statistics*, 42(3):940–969, 2014.
- Sivaraman Balakrishnan, Mladen Kolar, Alessandro Rinaldo, Aarti Singh, and Larry Wasserman. Statistical and computational tradeoffs in biclustering. In *NIPS 2011 workshop on computational trade-offs in statistical learning*, volume 4, 2011.
- Quentin Berthet and Philippe Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In *COLT*, pages 1046–1066, 2013a.
- Quentin Berthet and Philippe Rigollet. Optimal detection of sparse principal components in high dimension. *The Annals of Statistics*, 41(4):1780–1815, 2013b.
- Aditya Bhaskara, Moses Charikar, Eden Chlamtac, Uriel Feige, and Aravindan Vijayaraghavan. Detecting high log-densities: an $o(n^{1/4})$ approximation for densest k -subgraph. *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 201–210, 2010.
- Matthew Brennan and Guy Bresler. Optimal average-case reductions to sparse pca: From weak assumptions to strong hardness. *arXiv preprint arXiv:1902.07380*, 2019.
- Matthew Brennan, Guy Bresler, and Wasim Huleihel. Reducibility and computational lower bounds for problems with planted sparse structure. In *COLT*, pages 48–166, 2018.
- Cristina Butucea and Yuri I Ingster. Detection of a sparse submatrix of a high-dimensional noisy matrix. *Bernoulli*, 19(5B):2652–2688, 2013.
- T Tony Cai, Tengyuan Liang, Alexander Rakhlin, et al. Computational and statistical boundaries for submatrix localization in a large noisy matrix. *The Annals of Statistics*, 45(4):1403–1430, 2017.
- Utkan Onur Candogan and Venkat Chandrasekaran. Finding planted subgraphs with few eigenvalues using the schur–horn relaxation. *SIAM Journal on Optimization*, 28(1):735–759, 2018.
- Yudong Chen. Incoherence-optimal matrix completion. *IEEE Transactions on Information Theory*, 61(5):2909–2923, 2015.
- Yudong Chen and Jiaming Xu. Statistical-computational tradeoffs in planted problems and submatrix localization with a growing number of clusters and submatrices. *Journal of Machine Learning Research*, 17(27):1–57, 2016.

- Eden Chlamtac, Michael Dinitz, and Robert Krauthgamer. Everywhere-sparse spanners via dense subgraphs. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 758–767. IEEE, 2012.
- Eden Chlamtáč, Michael Dinitz, and Yury Makarychev. Minimizing the union: Tight approximations for small set bipartite vertex expansion. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 881–899. SIAM, 2017.
- E. Chlamtáč and P. Manurangsi. Sherali-adams integrality gaps matching the log-density threshold. *arXiv preprint arXiv:1804.07842*, 2018.
- Yael Dekel, Ori Gurel-Gurevich, and Yuval Peres. Finding hidden cliques in linear time with high probability. *Combinatorics, Probability and Computing*, 23(1):29–49, 2014.
- Yash Deshpande and Andrea Montanari. Finding hidden cliques of size \sqrt{N}/e in nearly linear time. *Foundations of Computational Mathematics*, 15(4):1069–1128, 2015.
- Yash Deshpande, Emmanuel Abbe, and Andrea Montanari. Asymptotic mutual information for the two-groups stochastic block model. *arXiv preprint arXiv:1507.08685*, 2015.
- Uriel Feige and Robert Krauthgamer. Finding and certifying a large hidden clique in a semirandom graph. *Random Structures and Algorithms*, 16(2):195–208, 2000.
- Uriel Feige and Dorit Ron. Finding hidden cliques in linear time. In *21st International Meeting on Probabilistic, Combinatorial, and Asymptotic Methods in the Analysis of Algorithms (AofA'10)*, pages 189–204. Discrete Mathematics and Theoretical Computer Science, 2010.
- Chao Gao, Zongming Ma, and Harrison H Zhou. Sparse cca: Adaptive estimation and computational barriers. *The Annals of Statistics*, 45(5):2074–2101, 2017.
- Bruce Hajek, Yihong Wu, and Jiaming Xu. Achieving exact cluster recovery threshold via semidefinite programming. *IEEE Transactions on Information Theory*, 62(5):2788–2797, 2016a.
- Bruce Hajek, Yihong Wu, and Jiaming Xu. Semidefinite programs for exact recovery of a hidden community. In *Conference on Learning Theory*, pages 1051–1095, 2016b.
- Bruce Hajek, Yihong Wu, and Jiaming Xu. Information limits for recovering a hidden community. *IEEE Transactions on Information Theory*, 63(8):4729–4745, 2017.
- Bruce E Hajek, Yihong Wu, and Jiaming Xu. Computational lower bounds for community detection on random graphs. In *COLT*, pages 899–928, 2015.
- Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. *The Annals of Statistics*, pages 286–295, 1983.
- Pascal Koiran and Anastasios Zouzias. Hidden cliques and the certification of the restricted isometry property. *IEEE Transactions on Information Theory*, 60(8):4999–5006, 2014.
- Mladen Kolar, Sivaraman Balakrishnan, Alessandro Rinaldo, and Aarti Singh. Minimax localization of structural information in large noisy matrices. In *Advances in Neural Information Processing Systems*, pages 909–917, 2011.

- Florent Krzakala, Jiaming Xu, and Lenka Zdeborová. Mutual information in rank-one matrix estimation. In *2016 IEEE Information Theory Workshop (ITW)*, pages 71–75. IEEE, 2016.
- Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. Mmse of probabilistic low-rank matrix estimation: Universality with respect to the output channel. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 680–687. IEEE, 2015.
- Zongming Ma and Yihong Wu. Computational barriers in minimax submatrix detection. *The Annals of Statistics*, 43(3):1089–1116, 2015.
- Frank McSherry. Spectral partitioning of random graphs. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 529–537. IEEE, 2001.
- Andrea Montanari. Finding one community in a sparse graph. *Journal of Statistical Physics*, 161(2):273–299, 2015.
- Andrea Montanari, Daniel Reichman, and Ofer Zeitouni. On the limitation of spectral methods: From the gaussian hidden clique problem to rank-one perturbations of gaussian tensors. In *Advances in Neural Information Processing Systems*, pages 217–225, 2015.
- Andrey A Shabalin, Victor J Weigman, Charles M Perou, Andrew B Nobel, et al. Finding large average submatrices in high dimensional data. *The Annals of Applied Statistics*, 3(3):985–1012, 2009.
- Nicolas Verzelen, Ery Arias-Castro, et al. Community detection in sparse random networks. *The Annals of Applied Probability*, 25(6):3465–3510, 2015.
- Tengyao Wang, Quentin Berthet, and Yaniv Plan. Average-case hardness of rip certification. In *Advances in Neural Information Processing Systems*, pages 3819–3827, 2016a.
- Tengyao Wang, Quentin Berthet, and Richard J Samworth. Statistical and computational trade-offs in estimation of sparse principal components. *The Annals of Statistics*, 44(5):1896–1930, 2016b.
- Yihong Wu and Jiaming Xu. Statistical problems with planted structures: Information-theoretical and computational limits. *arXiv preprint arXiv:1806.00118*, 2018.

Contents

1	Introduction	1
2	Submatrix Problems and Conditions for Universality	4
3	Summary of Results	6
4	Overview of Average-Case Reduction Techniques	10
A	Notation	17
B	Average-Case Reductions in Total Variation	17
C	Multivariate Rejection Kernels	19
D	Average-Case Reduction to Submatrix Detection	26
E	Computational Barriers in Submatrix Detection	36
F	Statistical Limit of Submatrix Detection	44
G	The Universality Classes UC-A, UC-B and UC-C	48
H	Further Questions	51

Appendix A. Notation

In this paper, we adopt the following notation. Let $\mathcal{L}(X)$ denote the distribution law of a random variable X and given two laws \mathcal{L}_1 and \mathcal{L}_2 , let $\mathcal{L}_1 + \mathcal{L}_2$ denote $\mathcal{L}(X+Y)$ where $X \sim \mathcal{L}_1$ and $Y \sim \mathcal{L}_2$ are independent. Given a distribution \mathcal{P} , let $\mathcal{P}^{\otimes n}$ denote the distribution of (X_1, X_2, \dots, X_n) where the X_i are i.i.d. according to \mathcal{P} . Similarly, let $\mathcal{P}^{\otimes m \times n}$ denote the distribution on $\mathbb{R}^{m \times n}$ with i.i.d. entries distributed as \mathcal{P} . Let d_{TV} , d_{KL} and χ^2 denote total variation distance, KL divergence and χ^2 divergence, respectively. Let $[n] = \{1, 2, \dots, n\}$ and let $\mathbf{1}_S$ denote the vector $v \in \mathbb{R}^n$ with $v_i = 1$ if $i \in S$ and $v_i = 0$ if $i \notin S$ where $S \subseteq [n]$.

Appendix B. Average-Case Reductions in Total Variation

B.1. Reductions in Total Variation and the Computational Model

As introduced in [Berthet and Rigollet \(2013a\)](#) and [Ma and Wu \(2015\)](#), we give approximate reductions in total variation to show that lower bounds for one hypothesis testing problem imply lower bounds for another. These reductions yield an exact correspondence between the asymptotic Type I+II errors of the two problems. This is formalized in the following lemma, which is Lemma 3.1 from [Brennan et al. \(2018\)](#) specialized to the case of simple vs. simple hypothesis testing. Its proof is short and follows from the definition of total variation.

Lemma 9 (Lemma 3.1 in [Brennan et al. \(2018\)](#)) *Let \mathcal{P}_D and \mathcal{P}'_D be detection problems with hypotheses H_0, H_1 and H'_0, H'_1 , respectively. Let X be an instance of \mathcal{P}_D and let Y be an instance of \mathcal{P}'_D . Suppose there is a polynomial time computable map \mathcal{A} satisfying*

$$d_{TV} \left(\mathcal{L}_{H_0}(\mathcal{A}(X)), \mathcal{L}_{H'_0}(Y) \right) + d_{TV} \left(\mathcal{L}_{H_1}(\mathcal{A}(X)), \mathcal{L}_{H'_1}(Y) \right) \leq \delta$$

If there is a randomized polynomial time algorithm solving \mathcal{P}'_D with Type I+II error at most ϵ , then there is a randomized polynomial time algorithm solving \mathcal{P}_D with Type I+II error at most $\epsilon + \delta$.

If $\delta = o(1)$, then given a blackbox solver \mathcal{B} for \mathcal{P}'_D , the algorithm that applies \mathcal{A} and then \mathcal{B} solves \mathcal{P}_D and requires only a single query to the blackbox. An algorithm that runs in randomized polynomial time refers to one that has access to $\text{poly}(n)$ independent random bits and must run in $\text{poly}(n)$ time where n is the size of the instance of the problem. For clarity of exposition, in our reductions we assume that explicit expressions can be exactly computed and that we can sample a biased random bit $\text{Bern}(p)$ in polynomial time. We also assume that the oracles described in Definition 2 can be computed in $\text{poly}(n)$ time.

B.2. Properties of Total Variation

Throughout the proof of our main theorem, we will use the following well-known facts and inequalities concerning total variation distance.

Fact B.1 *The distance d_{TV} satisfies the following properties:*

1. (Triangle Inequality) *Given three distributions P, Q and R on a measurable space $(\mathcal{X}, \mathcal{B})$, it follows that*

$$d_{TV}(P, Q) \leq d_{TV}(P, R) + d_{TV}(Q, R)$$

2. (Data Processing) Let P and Q be distributions on a measurable space $(\mathcal{X}, \mathcal{B})$ and let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a Markov transition kernel. If $A \sim P$ and $B \sim Q$ then

$$d_{TV}(\mathcal{L}(f(A)), \mathcal{L}(f(B))) \leq d_{TV}(P, Q)$$

3. (Tensorization) Let P_1, P_2, \dots, P_n and Q_1, Q_2, \dots, Q_n be distributions on a measurable space $(\mathcal{X}, \mathcal{B})$. Then

$$d_{TV}\left(\prod_{i=1}^n P_i, \prod_{i=1}^n Q_i\right) \leq \sum_{i=1}^n d_{TV}(P_i, Q_i)$$

4. (Conditioning on an Event) For any distribution P on a measurable space $(\mathcal{X}, \mathcal{B})$ and event $A \in \mathcal{B}$, it holds that

$$d_{TV}(P(\cdot|A), P) = 1 - P(A)$$

5. (Conditioning on a Random Variable) For any two pairs of random variables (X, Y) and (X', Y') each taking values in a measurable space $(\mathcal{X}, \mathcal{B})$, it holds that

$$d_{TV}(\mathcal{L}(X), \mathcal{L}(X')) \leq d_{TV}(\mathcal{L}(Y), \mathcal{L}(Y')) + \mathbb{E}_{y \sim Y} [d_{TV}(\mathcal{L}(X|Y=y), \mathcal{L}(X'|Y'=y))]$$

where we define $d_{TV}(\mathcal{L}(X|Y=y), \mathcal{L}(X'|Y'=y)) = 1$ for all $y \notin \text{supp}(Y')$.

Given an algorithm \mathcal{A} and distribution \mathcal{P} on inputs, let $\mathcal{A}(\mathcal{P})$ denote the distribution of $\mathcal{A}(X)$ induced by $X \sim \mathcal{P}$. If \mathcal{A} has k steps, let \mathcal{A}_i denote the i th step of \mathcal{A} and \mathcal{A}_{i-j} denote the procedure formed by steps i through j . Each time this notation is used, we clarify the intended initial and final variables when \mathcal{A}_i and \mathcal{A}_{i-j} are viewed as Markov kernels. The next lemma encapsulates the structure of all of our analyses of average-case reductions.

Lemma 10 *Let \mathcal{A} be an algorithm that can be written as $\mathcal{A} = \mathcal{A}_m \circ \mathcal{A}_{m-1} \circ \dots \circ \mathcal{A}_1$ for a sequence of steps $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_m$. Suppose that the probability distributions $\mathcal{P}_0, \mathcal{P}_1, \dots, \mathcal{P}_m$ are such that $d_{TV}(\mathcal{A}_i(\mathcal{P}_{i-1}), \mathcal{P}_i) \leq \epsilon_i$ for each $1 \leq i \leq m$. Then it follows that*

$$d_{TV}(\mathcal{A}(\mathcal{P}_0), \mathcal{P}_m) \leq \sum_{i=1}^m \epsilon_i$$

Proof This follows from a simple induction on m . Note that the case when $m = 1$ follows by definition. Now observe that by the data-processing and triangle inequalities in Fact B.1, we have that if $\mathcal{B} = \mathcal{A}_{m-1} \circ \mathcal{A}_{m-2} \circ \dots \circ \mathcal{A}_1$ then

$$\begin{aligned} d_{TV}(\mathcal{A}(\mathcal{P}_0), \mathcal{P}_m) &\leq d_{TV}(\mathcal{A}_m \circ \mathcal{B}(\mathcal{P}_0), \mathcal{A}_m(\mathcal{P}_{m-1})) + d_{TV}(\mathcal{A}_m(\mathcal{P}_{m-1}), \mathcal{P}_m) \\ &\leq d_{TV}(\mathcal{B}(\mathcal{P}_0), \mathcal{P}_{m-1}) + \epsilon_m \\ &\leq \sum_{i=1}^m \epsilon_i \end{aligned}$$

where the last inequality follows from the induction hypothesis applied with $m - 1$ to \mathcal{B} . This completes the induction and proves the lemma. \blacksquare

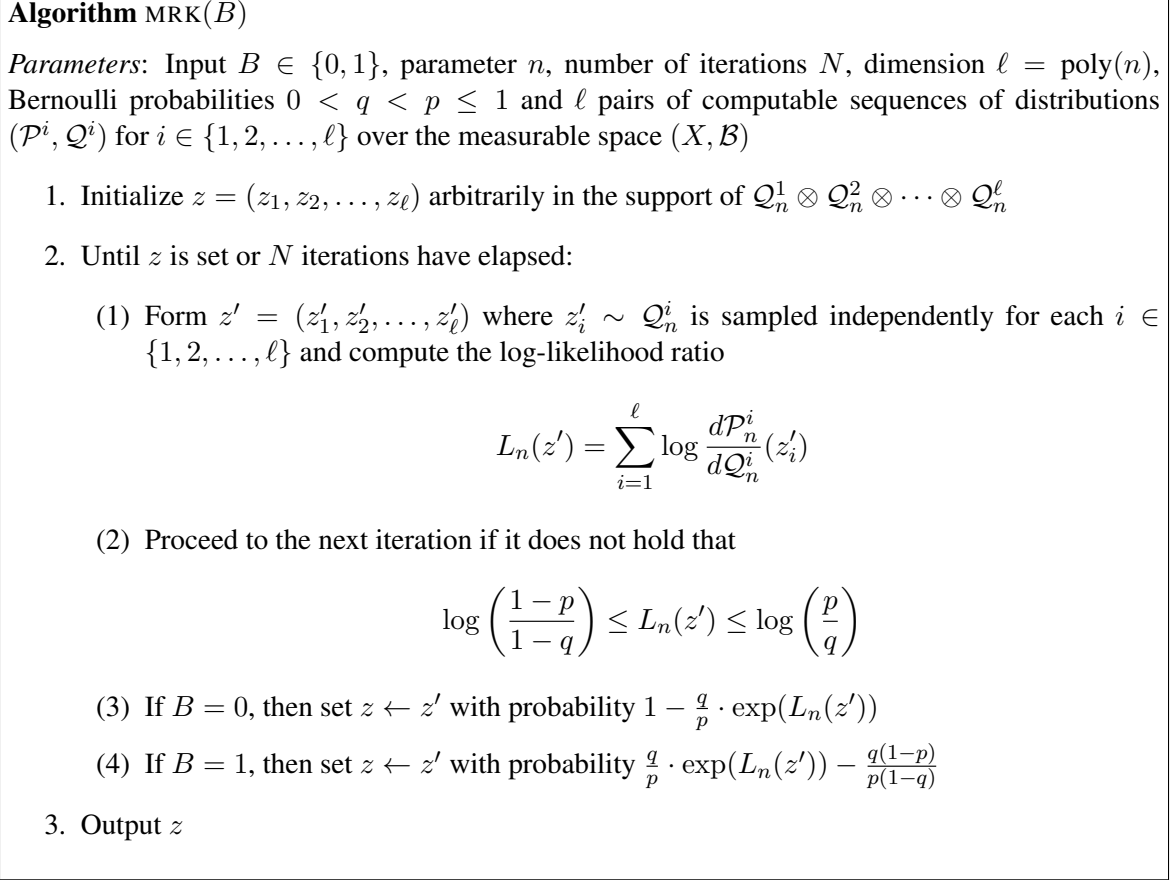


Figure 4: Multivariate rejection kernel algorithm with Bernoulli input $B \in \{0, 1\}$.

Appendix C. Multivariate Rejection Kernels

C.1. General MRK Algorithm and Analysis

In this section, we introduce multivariate rejection kernels, which effectively perform an algorithmic change of measure simultaneously with a lift. More precisely, the map $\text{MRK}(B)$ sends a binary input $B \in \{0, 1\}$ to a higher dimensional space X^ℓ simultaneously satisfying two Markov transition properties:

1. if $B \sim \text{Bern}(p)$, then $\text{MRK}(B)$ is close to $\mathcal{P}_n^1 \otimes \mathcal{P}_n^2 \otimes \dots \otimes \mathcal{P}_n^\ell$ in total variation; and
2. if $B \sim \text{Bern}(q)$, then $\text{MRK}(B)$ is close to $\mathcal{Q}_n^1 \otimes \mathcal{Q}_n^2 \otimes \dots \otimes \mathcal{Q}_n^\ell$ in total variation.

where $0 < q < p \leq 1$ are fixed and $(\mathcal{P}^i, \mathcal{Q}^i)$ are pairs of computable sequences of distributions over a common measurable space (X, \mathcal{B}) for $i \in \{1, 2, \dots, \ell\}$. The maps MRK will be a key part of our reduction to SSD, lifting planted dense subgraph and planted clique instances to instances of submatrix detection with independent entries and the correct marginals.

These maps MRK use rejection sampling to transform the distribution of the input from either one of two input Bernoulli distributions to two ℓ -fold product distributions, without knowing which of the two distributions was the input. Multivariate rejection kernels are closely related to previous average-case reduction techniques in the literature. Univariate Markov transitions taking δ_1 and $\text{Bern}(q)$ to either a fixed pair of Bernoulli distributions or a fixed pair of normal distributions were introduced in [Ma and Wu \(2015\)](#) and [Gao et al. \(2017\)](#). More closely related to our map are the rejection kernel framework introduced in [Brennan et al. \(2018\)](#) and the reduction from planted clique to planted dense subgraph in [Hajek et al. \(2015\)](#). The rejection kernels in [Brennan et al. \(2018\)](#) used a similar rejection sampling scheme as we use here to map binary inputs to distributions on \mathbb{R} . In [Brennan et al. \(2018\)](#), it is assumed that both of the target distributions f_X and g_X have explicit density or mass functions and have sampling oracles. In contrast, our map MRK only requires the Radon-Nikodym derivatives $\frac{d\mathcal{P}_n^i}{d\mathcal{Q}_n^i}$ exist and that there is a sampling oracle for the target noise distributions \mathcal{Q}_n^i rather than for \mathcal{P}_n^i as well, keeping the assumptions on $(\mathcal{P}^i, \mathcal{Q}^i)$ relatively minimal.

The reduction in [Hajek et al. \(2015\)](#) relies on a multivariate Markov transition mapping $\text{Bern}(1)$ and $\text{Bern}(q)$ to distributions of the form $\text{Bern}(cQ)^{\otimes \ell}$ and $\text{Bern}(Q)^{\otimes \ell}$ for some constant $c > 1$. This is achieved by first performing a univariate map from $\text{Bern}(1)$ and $\text{Bern}(q)$ to $\text{Bin}(\ell, cQ)$ and $\text{Bin}(\ell, Q)$, respectively, and then observing that these counts are sufficient statistics for the target product distributions. As discussed in the survey [Wu and Xu \(2018\)](#), this approach extends naturally to other target distributions that have a common sufficient statistic with an explicit mass or density function. For example, the sum of entries is such a sufficient statistic for the family of distribution $\mathcal{N}(\mu, 1)^{\otimes \ell}$ for $\mu \in \mathbb{R}$. On generalizing to any computable pairs $(\mathcal{P}^i, \mathcal{Q}^i)$, such a sufficient statistic common to $\mathcal{P}_n^1 \otimes \mathcal{P}_n^2 \otimes \cdots \otimes \mathcal{P}_n^\ell$ and $\mathcal{Q}_n^1 \otimes \mathcal{Q}_n^2 \otimes \cdots \otimes \mathcal{Q}_n^\ell$ may not exist. We remark that all of the univariate maps in [Ma and Wu \(2015\)](#), [Gao et al. \(2017\)](#) and [Hajek et al. \(2015\)](#) rely on sampling explicit mass or density functions that are linear combinations of the target distributions. This requires time at least the size of the support of the target distribution or of a sufficiently fine net of its support. In the multivariate case mapping to distributions on X^ℓ , the size of this support grows exponentially in ℓ and this sampling scheme is not feasible in polynomial time. The rejection sampling scheme used here circumvents this issue. We now prove the main total variation guarantees of $\text{MRK}(B)$, generalizing the argument in Lemma 5.1 in [Brennan et al. \(2018\)](#).

Lemma 11 (Multivariate Rejection Kernels) *Suppose that n is a parameter, $0 < q < p \leq 1$ and N is a positive integer. Let $(\mathcal{P}^i, \mathcal{Q}^i)$ for $i \in \{1, 2, \dots, \ell\}$ be pairs of computable sequences of distributions over the measurable space (X, \mathcal{B}) and let*

$$S = \left\{ x \in X^\ell : \log \left(\frac{1-p}{1-q} \right) \leq L_n(x) \leq \log \left(\frac{p}{q} \right) \right\}$$

where $L_n(x) = \sum_{i=1}^{\ell} \log \frac{d\mathcal{P}_n^i}{d\mathcal{Q}_n^i}(x_i)$ for each $x = (x_1, x_2, \dots, x_\ell) \in X^\ell$. Then there is a map $\text{MRK} : \{0, 1\} \rightarrow X^\ell$ that can be computed in $\text{poly}(n, \ell, N)$ time satisfying that

$$\begin{aligned} d_{TV} \left(\text{MRK}(\text{Bern}(p)), \mathcal{P}_n^1 \otimes \mathcal{P}_n^2 \otimes \cdots \otimes \mathcal{P}_n^\ell \right) &\leq \Delta \\ d_{TV} \left(\text{MRK}(\text{Bern}(q)), \mathcal{Q}_n^1 \otimes \mathcal{Q}_n^2 \otimes \cdots \otimes \mathcal{Q}_n^\ell \right) &\leq \Delta \end{aligned}$$

where if $\mathcal{P}_n^* = \mathcal{P}_n^1 \otimes \mathcal{P}_n^2 \otimes \dots \otimes \mathcal{P}_n^\ell$ and $\mathcal{Q}_n^* = \mathcal{Q}_n^1 \otimes \mathcal{Q}_n^2 \otimes \dots \otimes \mathcal{Q}_n^\ell$, then Δ is given by

$$\Delta = \frac{\mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \notin S] + \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \notin S]}{p - q} + \max \left\{ \left(\mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \notin S] + \frac{q}{p} \right)^N, \left(\frac{q}{p} \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \notin S] + \frac{p - 2pq + q^2}{p - pq} \right)^N \right\}$$

Proof Let MRK be the $\text{poly}(n, \ell, N)$ time algorithm shown in Figure 4. For the sake of analysis, consider continuing to iterate Step 2 even after z is set for the first time for a total of N iterations. Let A_i^0 and A_i^1 be the events that z is set in the i th iteration of Step 2 when $B = 0$ and $B = 1$, respectively. For $B = 0$, let $B_i^0 = (A_1^0)^C \cap (A_2^0)^C \cap \dots \cap (A_{i-1}^0)^C \cap A_i^0$ be the event that z is set for the first time in the i th iteration of Step 2. Let $C^0 = A_1^0 \cup A_2^0 \cup \dots \cup A_N^0$ be the event that z is set in some iteration of Step 2. Define B_i^1 and C^1 analogously. Let z_0 be the initialization of z in Step 1.

Now let $Z_0 \sim \mathcal{D}_0 = \mathcal{L}(\text{MRK}(0))$ and $Z_1 \sim \mathcal{D}_1 = \mathcal{L}(\text{MRK}(1))$. We have that $\mathcal{L}(Z_0|B_i^0) = \mathcal{L}(Z_0|A_i^0)$ and $\mathcal{L}(Z_1|B_i^1) = \mathcal{L}(Z_1|A_i^1)$ since A_i^t is independent of $A_1^t, A_2^t, \dots, A_{i-1}^t$ for each $t \in \{0, 1\}$ and the sample z' chosen in the i th iteration of Step 2. Observe that independence between Steps 2.1, 2.3 and 2.4 ensures that

$$\begin{aligned} \mathbb{P}[A_i^0] &= \mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\left(1 - \frac{q}{p} \cdot \exp(L_n(x)) \right) \cdot \mathbf{1}_S(x) \right] \\ &= \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \in S] - \frac{q}{p} \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \in S] \\ \mathbb{P}[A_i^1] &= \mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\left(\frac{q}{p} \cdot \exp(L_n(x)) - \frac{q(1-p)}{p(1-q)} \right) \cdot \mathbf{1}_S(x) \right] \\ &= \frac{q}{p} \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \in S] - \frac{q(1-p)}{p(1-q)} \cdot \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \in S] \end{aligned}$$

since $\frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x) = \exp(L_n(x))$. By the independence of the A_i^0 , we have that

$$\begin{aligned} 1 - \mathbb{P}[C^0] &= \prod_{i=1}^N (1 - \mathbb{P}[A_i^0]) = \left(1 - \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \in S] + \frac{q}{p} \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \in S] \right)^N \\ &\leq \left(\mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \notin S] + \frac{q}{p} \right)^N \end{aligned}$$

Similarly, we have that

$$\begin{aligned} 1 - \mathbb{P}[C^1] &= \prod_{i=1}^N (1 - \mathbb{P}[A_i^1]) = \left(1 - \frac{q}{p} \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \in S] + \frac{q(1-p)}{p(1-q)} \cdot \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \in S] \right)^N \\ &\leq \left(\frac{q}{p} \cdot (1 - \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \in S]) + 1 - \frac{q}{p} + \frac{q(1-p)}{p(1-q)} \right)^N \\ &\leq \left(\frac{q}{p} \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \notin S] + \frac{p - 2pq + q^2}{p - pq} \right)^N \end{aligned}$$

We now have that $\mathcal{L}(Z_0|A_i^0)$ and $\mathcal{L}(Z_1|A_i^1)$ are each absolutely continuous with respect to \mathcal{Q}_n^* with Radon-Nikodym derivatives

$$\begin{aligned} \frac{d\mathcal{L}(Z_0|B_i^0)}{d\mathcal{Q}_n^*}(x) &= \frac{d\mathcal{L}(Z_0|A_i^0)}{d\mathcal{Q}_n^*}(x) = \mathbb{P}[A_i^0]^{-1} \cdot \left(1 - \frac{q}{p} \cdot \exp(L_n(x))\right) \cdot \mathbf{1}_S(x) \\ &= \frac{p - q \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x)}{p \cdot \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \in S] - q \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \in S]} \cdot \mathbf{1}_S(x) \\ \frac{d\mathcal{L}(Z_1|B_i^1)}{d\mathcal{Q}_n^*}(x) &= \frac{d\mathcal{L}(Z_1|A_i^1)}{d\mathcal{Q}_n^*}(x) = \mathbb{P}[A_i^1]^{-1} \cdot \left(\frac{q}{p} \cdot \exp(L_n(z')) - \frac{q(1-p)}{p(1-q)}\right) \cdot \mathbf{1}_S(x) \\ &= \frac{(1-q) \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x) - (1-p)}{(1-q) \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \in S] - (1-p) \cdot \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \in S]} \cdot \mathbf{1}_S(x) \end{aligned}$$

Now observe that since the conditional laws $\mathcal{L}(Z_0|B_i^0)$ are all identical, we have that

$$\frac{d\mathcal{D}_0}{d\mathcal{Q}_n^*}(x) = \mathbb{P}[C^0] \cdot \frac{d\mathcal{L}(Z_0|B_1^0)}{d\mathcal{Q}_n^*}(x) + (1 - \mathbb{P}[C^0]) \cdot \mathbf{1}_{z_0}(x)$$

Therefore it follows that

$$\begin{aligned} d_{\text{TV}}(\mathcal{D}_0, \mathcal{L}(Z_0|B_1^0)) &= \frac{1}{2} \cdot \mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\left| \frac{d\mathcal{D}_0}{d\mathcal{Q}_n^*}(x) - \frac{d\mathcal{L}(Z_0|B_1^0)}{d\mathcal{Q}_n^*}(x) \right| \right] \\ &= \frac{1}{2} (1 - \mathcal{P}[C^0]) \cdot \mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\left| \mathbf{1}_{z_0}(x) - \frac{d\mathcal{L}(Z_0|B_1^0)}{d\mathcal{Q}_n^*}(x) \right| \right] \\ &\leq \frac{1}{2} (1 - \mathcal{P}[C^0]) \cdot \mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\mathbf{1}_{z_0}(x) + \frac{d\mathcal{L}(Z_0|B_1^0)}{d\mathcal{Q}_n^*}(x) \right] \\ &= 1 - \mathcal{P}[C^0] \end{aligned}$$

by the triangle inequality. A symmetric argument implies that $d_{\text{TV}}(\mathcal{D}_1, \mathcal{L}(Z_1|B_1^1)) \leq 1 - \mathcal{P}[C^1]$.

Now observe that since $\frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x) \leq \frac{p}{q}$ for $x \in S$, we have that

$$\begin{aligned} &\mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\left| \frac{d\mathcal{L}(Z_0|B_1^0)}{d\mathcal{Q}_n^*}(x) - \frac{p - q \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x)}{p - q} \right| \right] \\ &= \left| \frac{1}{p \cdot \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \in S] - q \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \in S]} - \frac{1}{p - q} \right| \cdot \mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\left(p - q \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x) \right) \cdot \mathbf{1}_S(x) \right] \\ &\quad + \mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\left| \frac{p - q \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x)}{p - q} \right| \cdot \mathbf{1}_{S^c}(x) \right] \\ &\leq \left| 1 - \frac{p \cdot \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \in S] - q \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \in S]}{p - q} \right| + \mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\frac{p + q \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x)}{p - q} \cdot \mathbf{1}_{S^c}(x) \right] \\ &= \left| \frac{p \cdot \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \notin S] - q \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \notin S]}{p - q} \right| + \frac{p \cdot \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \notin S] + q \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \notin S]}{p - q} \\ &\leq \frac{2 \cdot \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \notin S] + 2 \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \notin S]}{p - q} \end{aligned}$$

By a symmetric computation, we have that

$$\mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\left| \frac{d\mathcal{L}(Z_1|B_1^1)}{d\mathcal{Q}_n^*}(x) - \frac{(1-q) \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x) - (1-p)}{p-q} \right| \right] \leq \frac{2 \cdot \mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \notin S] + 2 \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*}[x \notin S]}{p-q}$$

Now observe that

$$\begin{aligned} \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x) &= p \cdot \frac{(1-q) \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x) - (1-p)}{p-q} + (1-p) \cdot \frac{p-q \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x)}{p-q} \\ 1 &= q \cdot \frac{(1-q) \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x) - (1-p)}{p-q} + (1-q) \cdot \frac{p-q \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x)}{p-q} \end{aligned}$$

Now observe that

$$\begin{aligned} d_{\text{TV}}(\text{MRK}(\text{Bern}(p)), \mathcal{P}_n^*) &= d_{\text{TV}}(p \cdot \mathcal{L}(Z_1|B_1^1) + (1-p) \cdot \mathcal{L}(Z_0|B_1^0), \mathcal{P}_n^*) \\ &\quad + d_{\text{TV}}(p \cdot \mathcal{L}(Z_1|B_1^1) + (1-p) \cdot \mathcal{L}(Z_0|B_1^0), \text{MRK}(\text{Bern}(p))) \\ &\leq \frac{p}{2} \cdot \mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\left| \frac{d\mathcal{L}(Z_1|B_1^1)}{d\mathcal{Q}_n^*}(x) - \frac{p-q \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x)}{p-q} \right| \right] \\ &\quad + \frac{(1-p)}{2} \cdot \mathbb{E}_{x \sim \mathcal{Q}_n^*} \left[\left| \frac{d\mathcal{L}(Z_0|B_1^0)}{d\mathcal{Q}_n^*}(x) - \frac{p-q \cdot \frac{d\mathcal{P}_n^*}{d\mathcal{Q}_n^*}(x)}{p-q} \right| \right] \\ &\quad + p \cdot d_{\text{TV}}(\mathcal{D}_1, \mathcal{L}(Z_1|B_1^1)) + (1-p) \cdot d_{\text{TV}}(\mathcal{D}_0, \mathcal{L}(Z_0|B_1^0)) \\ &\leq \Delta \end{aligned}$$

A symmetric argument shows $d_{\text{TV}}(\text{MRK}(\text{Bern}(q)), \mathcal{Q}_n^*) \leq \Delta$, completing the proof of the lemma. \blacksquare

C.2. Homogeneous MRK and Log-Likelihood Ratio LDPs

We now show that when all of the pairs $(\mathcal{P}^i, \mathcal{Q}^i) = (\mathcal{P}, \mathcal{Q})$ are the same and their LLR satisfies LDPs with respect to each of \mathcal{Q} and \mathcal{P} , then the total variation error Δ is small in the lemma above. This will be the form of the lemma that we will apply to show computational lower bounds for SSD.

Lemma 12 *Let $(\mathcal{P}, \mathcal{Q})$ be a computable pair over (X, \mathcal{B}) . Define n, ℓ, p and q as in Lemma 11. Let $\tau_+, \tau_- \geq \ell^{-1} \cdot \log(4(p-q)^{-1})$ and suppose that $(\mathcal{P}, \mathcal{Q})$ satisfies*

$$\log\left(\frac{1-q}{1-p}\right) < -\ell \cdot d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}) \leq \ell \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) < \log\left(\frac{p}{q}\right)$$

and the LLR between $(\mathcal{P}, \mathcal{Q})$ satisfies the LDPs

$$E_{\mathcal{P}}\left(\ell^{-1} \cdot \log\left(\frac{p}{q}\right)\right) \geq \tau_+ \quad \text{and} \quad E_{\mathcal{Q}}\left(\ell^{-1} \cdot \log\left(\frac{1-p}{1-q}\right)\right) \geq \tau_-$$

where the second inequality is only necessary if $p \neq 1$. Then there is a map $\text{MRK} : \{0, 1\} \rightarrow X^\ell$ that can be computed in $\text{poly}(n, \ell, pq^{-1}(p-q)^{-1}, \min(\tau_+, \tau_-))$ time satisfying that

$$\begin{aligned} d_{TV} \left(\text{MRK}(\text{Bern}(p)), \mathcal{P}_n^{\otimes \ell} \right) &\leq \frac{3(e^{-\ell\tau_+} + e^{-\ell\tau_-})}{p-q} \\ d_{TV} \left(\text{MRK}(\text{Bern}(q)), \mathcal{Q}_n^{\otimes \ell} \right) &\leq \frac{3(e^{-\ell\tau_+} + e^{-\ell\tau_-})}{p-q} \end{aligned}$$

Proof Let $c_+ = \log\left(\frac{p}{q}\right) > 0$ and $c_- = \log\left(\frac{1-p}{1-q}\right) < 0$. Observe that

$$\begin{aligned} E_{\mathcal{Q}}(c_+ \ell^{-1}) &= E_{\mathcal{P}}(c_+ \ell^{-1}) + c_+ \ell^{-1} > \tau_+ \\ E_{\mathcal{P}}(c_- \ell^{-1}) &= E_{\mathcal{Q}}(c_- \ell^{-1}) - c_- \ell^{-1} > \tau_- \end{aligned}$$

where the second set of inequalities holds if $p \neq 1$. Now consider applying Lemma 11 with number of iterations

$$N = \left\lceil \frac{\ell \cdot \min(\tau_+, \tau_-)}{\log\left(1 - \frac{q(p-q)}{2p}\right)^{-1}} \right\rceil = O\left(\frac{p}{q(p-q)} \cdot \ell \cdot \min(\tau_+, \tau_-)\right)$$

since $\log(1-x)^{-1} \geq x$ for $x = \frac{q(p-q)}{2p}$. Thus it suffices to bound Δ given that the LLR between $(\mathcal{P}, \mathcal{Q})$ satisfies the LDPs above. For now consider the case where $p \neq 1$. Let $L_n(x) = \sum_{i=1}^{\ell} \log \frac{d\mathcal{P}_n}{d\mathcal{Q}_n}(x_i)$ for each $x \in X^\ell$. Now consider λ_1 with objective value approaching the supremum in the optimization $E_{\mathcal{P}}(c_+ \ell^{-1}) = \sup_{\lambda \in \mathbb{R}} c_+ \ell^{-1} \lambda - \psi_{\mathcal{P}}(\lambda)$. Since the derivative of the objective at $\lambda = 0$ is $c_+ \ell^{-1} - d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) > 0$, we can take $\lambda_1 \geq 0$. Analogously, we can take $\lambda_2 \leq 0$ approaching the supremum in the optimization $E_{\mathcal{P}}(c_- \ell^{-1})$. Applying a Chernoff bound now yields that

$$\begin{aligned} \mathbb{P}_{x \sim \mathcal{P}_n^{\otimes \ell}} [x \notin S] &= \mathbb{P}_{x \sim \mathcal{P}_n^{\otimes \ell}} \left[\exp(\lambda_1 \cdot L_n(x)) > \left(\frac{p}{q}\right)^{\lambda_1} \right] + \mathbb{P}_{x \sim \mathcal{P}_n^{\otimes \ell}} \left[\exp(\lambda_2 \cdot L_n(x)) > \left(\frac{1-p}{1-q}\right)^{\lambda_2} \right] \\ &\leq \left(\frac{p}{q}\right)^{-\lambda_1} \cdot \mathbb{E}_{X \sim \mathcal{P}_n^{\otimes \ell}} [\exp(\lambda_1 \cdot L_n(x))] + \left(\frac{1-p}{1-q}\right)^{-\lambda_2} \cdot \mathbb{E}_{X \sim \mathcal{P}_n^{\otimes \ell}} [\exp(-\lambda_2 \cdot L_n(x))] \\ &= \exp(-c_+ \lambda_1 + \ell \cdot \psi_{\mathcal{P}}(\lambda_1)) + \exp(-c_- \lambda_2 + \ell \cdot \psi_{\mathcal{P}}(\lambda_2)) \\ &\rightarrow \exp(-\ell \cdot E_{\mathcal{P}}(c_+ \ell^{-1})) + \exp(-\ell \cdot E_{\mathcal{P}}(c_- \ell^{-1})) \\ &\leq e^{-\ell\tau_+} + e^{-\ell\tau_-} \end{aligned}$$

where the limit holds on taking λ_1 and λ_2 to approach the two suprema. Note that if $p = 1$, then the second probability above trivially satisfies that

$$\mathbb{P}_{x \sim \mathcal{P}_n^{\otimes \ell}} \left[L_n(x) < \frac{1-p}{1-q} \right] = 0 < e^{-\ell\tau_-}$$

and the bound on $\mathbb{P}_{x \sim \mathcal{P}_n^{\otimes \ell}} [x \notin S]$ above still holds. By a symmetric argument, we also have that

$$\mathbb{P}_{x \sim \mathcal{Q}_n^{\otimes \ell}} [x \notin S] \leq e^{-\ell\tau_+} + e^{-\ell\tau_-}$$

Now observe that since $e^{-\ell\tau_+}, e^{-\ell\tau_-} \leq \frac{1}{4}(p-q)$, we have that

$$\frac{q}{p} \cdot \mathbb{P}_{x \sim \mathcal{P}_n^{\otimes \ell}} [x \notin S] + \frac{p-2pq+q^2}{p-pq} \leq \frac{q}{p} \cdot \left(\frac{1}{2}(p-q) \right) + 1 - \frac{q(p-q)}{p(1-q)} \leq 1 - \frac{q(p-q)}{2p(1-q)} \leq 1 - \frac{q(p-q)}{2p}$$

Similarly, we have

$$\mathbb{P}_{x \sim \mathcal{Q}_n^{\otimes \ell}} [x \notin S] + \frac{q}{p} \leq \frac{1}{2}(p-q) + 1 - \frac{p-q}{p} \leq 1 - \frac{q(p-q)}{2p}$$

Therefore it follows that

$$\begin{aligned} & \max \left\{ \left(\mathbb{P}_{x \sim \mathcal{Q}_n^*} [x \notin S] + \frac{q}{p} \right)^N, \left(\frac{q}{p} \cdot \mathbb{P}_{x \sim \mathcal{P}_n^*} [x \notin S] + \frac{p-2pq+q^2}{p-pq} \right)^N \right\} \\ & \leq \left(1 - \frac{q(p-q)}{2p} \right)^N \leq \max \left(e^{-\ell\tau_+}, e^{-\ell\tau_-} \right) \end{aligned}$$

and hence that $\Delta \leq \frac{3(e^{-\ell\tau_+} + e^{-\ell\tau_-})}{p-q}$, completing the proof of the lemma. \blacksquare

C.3. Entrywise Reductions Fail to Show Tight Computational Lower Bounds

A conceptually simpler idea for an average-case reduction would avoid multivariate rejection kernels and instead begin by mapping to a submatrix problem with a specific pair $(\mathcal{P}^*, \mathcal{Q}^*)$ and then apply a univariate entry-wise map from $(\mathcal{P}^*, \mathcal{Q}^*)$ to $(\mathcal{P}, \mathcal{Q})$. For example, if $(\mathcal{P}^*, \mathcal{Q}^*)$ were Bernoulli random variables, this would correspond to mapping to a submatrix variant of planted dense subgraph and then to $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$. We remark that there is a simple barrier to making this approach produce tight computational lower bounds.

As discussed in the summary of our results in Section 3, the relevant signal in the computational phase diagram for SSD is $d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})$. For this approach to lead to a reduction, there would have to be a univariate map ϕ such that $\phi(\mathcal{P}^*) \sim \mathcal{P}$ and $\phi(\mathcal{Q}^*) \sim \mathcal{Q}$ as long as $d_{\text{KL}}(\mathcal{P}^* \parallel \mathcal{Q}^*) \asymp d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})$. However, looking at the family of pairs of Bernoulli random variables shows that this is not the case. If $(\mathcal{P}^*, \mathcal{Q}^*) = (\text{Bern}(n^{-\alpha}), \text{Bern}(2n^{-\alpha}))$ then we have that

$$d_{\text{KL}}(\mathcal{P}^* \parallel \mathcal{Q}^*) = -n^{-\alpha} \log 2 - (1-n^{-\alpha}) \log \left(\frac{1-2n^{-\alpha}}{1-n^{-\alpha}} \right) = \Theta(n^{-\alpha}) \quad \text{and} \quad d_{\text{TV}}(\mathcal{Q}^* \parallel \mathcal{P}^*) = n^{-\alpha}$$

Furthermore, let $(\mathcal{P}, \mathcal{Q}) = (\text{Bern}(1/2), \text{Bern}(1/2 + n^{-\alpha/2}))$ and note that

$$d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) = -\frac{1}{2} \log(1 + 2n^{-\alpha/2}) - \frac{1}{2} \log(1 - 2n^{-\alpha/2}) = \Theta(n^{-\alpha}) \quad \text{and} \quad d_{\text{TV}}(\mathcal{Q} \parallel \mathcal{P}) = n^{-\alpha/2}$$

By the data-processing applied to d_{TV} , it is impossible for such a map ϕ to exist even though these two pairs have KL divergences on the same order.

Algorithm TO-SUBMATRIX

Inputs: Graph $G \in \mathcal{G}_n$, parameters $0 < q < p \leq 1$, intermediate dimension N , expansion factor ℓ , number of iterations N_{it} , subgraph size k , pairs of computable sequences of distributions $(\mathcal{P}_{ij}, \mathcal{Q}_{ij})$ for $1 \leq i, j \leq N\ell$ over the measurable space (X, \mathcal{B})

1. Apply GRAPH-CLONE to G with edge probabilities $P = p$ and $Q = 1 - \sqrt{(1-p)(1-q)} + \mathbf{1}_{\{p=1\}}(\sqrt{q} - 1)$ and $t = 2$ clones to obtain (G_1, G_2)
2. Sample $s_1 \sim \text{Bin}(n, p)$, $s_2 \sim \text{Bin}(N, Q)$ and a set $S \subseteq [N]$ with $|S| = n$ uniformly at random. Sample $T_1 \subseteq S$ and $T_2 \subseteq [N] \setminus S$ with $|T_1| = s_1$ and $|T_2| = \max\{s_2 - s_1, 0\}$ uniformly at random. Now form the matrix $M^1 \in \{0, 1\}^{N \times N}$ where

$$M_{ij}^1 = \begin{cases} \mathbf{1}_{\{\pi(i), \pi(j)\} \in E(G_1)} & \text{if } i < j \text{ and } i, j \in S \\ \mathbf{1}_{\{\pi(i), \pi(j)\} \in E(G_2)} & \text{if } i > j \text{ and } i, j \in S \\ \mathbf{1}_{\{i \in T_1\}} & \text{if } i = j \text{ and } i, j \in S \\ \mathbf{1}_{\{i \in T_2\}} & \text{if } i = j \text{ and } i, j \notin S \\ \sim_{\text{i.i.d.}} \text{Bern}(Q) & \text{if } i \neq j \text{ and } i \notin S \text{ or } j \notin S \end{cases}$$

where $\pi : S \rightarrow [n]$ is a bijection chosen uniformly at random

3. Let τ be a random permutation of $[N\ell]$ and form the matrix $M^2 \in X^{N\ell \times N\ell}$ by setting

$$(M_{ij}^2 : s\ell + 1 \leq \tau^{-1}(i) \leq (s+1)\ell \wedge t\ell + 1 \leq \tau^{-1}(j) \leq (t+1)\ell) = \text{MRK}_{st} \left(M_{(s+1)(t+1)}^1 \right)$$

for each $0 \leq s, t < N$, where MRK_{st} is the multivariate rejection kernel with N_{it} iterations sending $\text{Bern}(p)$ and $\text{Bern}(Q)$ to the random set of ℓ^2 pairs $(\mathcal{P}_{ij}^n, \mathcal{Q}_{ij}^n)$ satisfying that $s\ell + 1 \leq \tau^{-1}(i) \leq (s+1)\ell$ and $t\ell + 1 \leq \tau^{-1}(j) \leq (t+1)\ell$

4. Output the matrix M^2

Figure 5: Reduction TO-SUBMATRIX for showing universal computational lower bounds for submatrix problems based on the hardness of planted clique or planted dense subgraph.

Appendix D. Average-Case Reduction to Submatrix Detection

In this section, we give a polynomial time reduction TO-SUBMATRIX to SSD detection with pairs of planted and noise distributions that are computable and satisfy an LDP of their LLR. This reduction is shown in Figure 5 and its total variation guarantees are stated in Theorem 13 below. The reduction begins by cloning the adjacency matrix of a planted dense subgraph instance to produce two independent samples of the above diagonal portions of their adjacency matrix. These are then embedded as two halves of a principal minor in a larger matrix in Step 2. This random embedding hides the previously missing diagonal terms in total variation. Analyzing this step is one of our

technical contributions. The resulting Bernoulli submatrix problem is then lifted using MRK maps to an instance of submatrix detection.

Our reduction implies the computational lower bounds for Gaussian biclustering and planted dense subgraph in [Ma and Wu \(2015\)](#), [Hajek et al. \(2015\)](#) and [Brennan et al. \(2018\)](#). The step of embedding as a principal minor in Step 2 circumvents the arguments in [Hajek et al. \(2015\)](#) and [Brennan et al. \(2018\)](#) showing that the total variation error introduced by missing diagonal entries is small. It also simplifies the reductions PC-LIFTING and DISTRIBUTIONAL-LIFTING in [Brennan et al. \(2018\)](#). Our reduction extends the Gaussian biclustering lower bounds in [Ma and Wu \(2015\)](#) and [Brennan et al. \(2018\)](#) to the cases of symmetric indices.

In Section [D.1](#), we describe and analyze the subroutine GRAPH-CLONE in Step 1 which transforms one instance of planted dense subgraph into many independent instances, preserving both H_0 and H_1 , at a minor loss in the parameters p and q . In Section [D.2](#), we prove several lemmas to analyze Step 2 of TO-SUBMATRIX and in Section [D.3](#), we complete the proof of Theorem [13](#). In the next section, we apply Theorem [13](#) to deduce computational lower bounds for various forms of submatrix detection. Note that the reduction TO-SUBMATRIX handles a more general setup where the pairs $(\mathcal{P}_{ij}, \mathcal{Q}_{ij})$ are permitted to differ from entry to entry. The additional submatrix problems this implies hardness for are discussed in Section [E](#). TO-SUBMATRIX also begins with an instance of planted dense subgraph, providing a more general class of reductions than just from planted clique.

Let $\mathcal{M}_d((\mathcal{Q}_{ij})_{1 \leq i, j \leq d})$ denote the distribution on $d \times d$ matrices M with independent entries satisfying that $M_{ij} \sim \mathcal{Q}_{ij}$ for each $1 \leq i, j \leq d$. Let $\mathcal{M}_d((\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq d}, k)$ denote the distribution on $d \times d$ matrices M generated as follows: choose a k -subset $S \subseteq [d]$ uniformly at random and then independently sample $M_{ij} \sim \mathcal{P}_{ij}$ for each $i, j \in S$ and $M_{ij} \sim \mathcal{Q}_{ij}$ otherwise.

Theorem 13 (Reduction to Submatrix Detection) *Let $n, k \leq n, N, N_{\text{it}}, \ell$ and $0 < q < p \leq 1$ be parameters such that $q = n^{-O(1)}$,*

$$N \geq \left(\frac{p}{Q} + \epsilon \right) n, \quad k \leq \frac{Q\epsilon n}{2} \quad \text{and} \quad \frac{k^2}{N} \leq \min \left\{ \frac{Q}{1-Q}, \frac{1-Q}{Q} \right\}$$

where $\epsilon > 0$ and $Q = 1 - \sqrt{(1-p)(1-q)} + \mathbf{1}_{\{p=1\}}(\sqrt{q} - 1)$. Let $(\mathcal{P}_{ij}, \mathcal{Q}_{ij})$ for $1 \leq i, j \leq N\ell$ be pairs of computable sequences of distributions indexed by n over the measurable space (X, \mathcal{B}) . For each pair of ℓ -subsets $U, V \subseteq [N\ell]$, let $\Delta_{U,V}$ be the total variation upper bound as defined in Lemma [11](#) for the multivariate rejection kernel $\text{MRK}_{U,V}$ with N_{it} iterations sending $\text{Bern}(p)$ and $\text{Bern}(Q)$ to the ℓ^2 pairs $(\mathcal{P}_{ij}^n, \mathcal{Q}_{ij}^n)$ where $i \in U$ and $j \in V$. Let $\Delta = \max_{U,V} \Delta_{U,V}$ be the maximum such upper bound. The algorithm $\mathcal{A} = \text{TO-SUBMATRIX}$ in Figure [5](#) runs in $\text{poly}(N, \ell)$ time and satisfies that

$$\begin{aligned} d_{\text{TV}}(\mathcal{A}(\mathcal{G}(n, k, p, q)), \mathcal{M}_{N\ell}((\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq N\ell}, k\ell)) &\leq N^2 \cdot \Delta + 4 \cdot \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right) \\ &\quad + \sqrt{\frac{k^2(1-Q)}{2QN}} + \sqrt{\frac{k^2 Q}{2N(1-Q)}} \\ d_{\text{TV}}(\mathcal{A}(\mathcal{G}(n, q)), \mathcal{M}_{N\ell}((\mathcal{Q}_{ij})_{1 \leq i, j \leq N\ell})) &\leq N^2 \cdot \Delta + 4 \cdot \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right) \end{aligned}$$

Before proceeding to the proofs in this section, we first establish some additional notation. When all of the pairs $(\mathcal{P}_{ij}, \mathcal{Q}_{ij})$ are the same $(\mathcal{P}, \mathcal{Q})$, then denote $\mathcal{M}_d((\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq d}, k)$ by

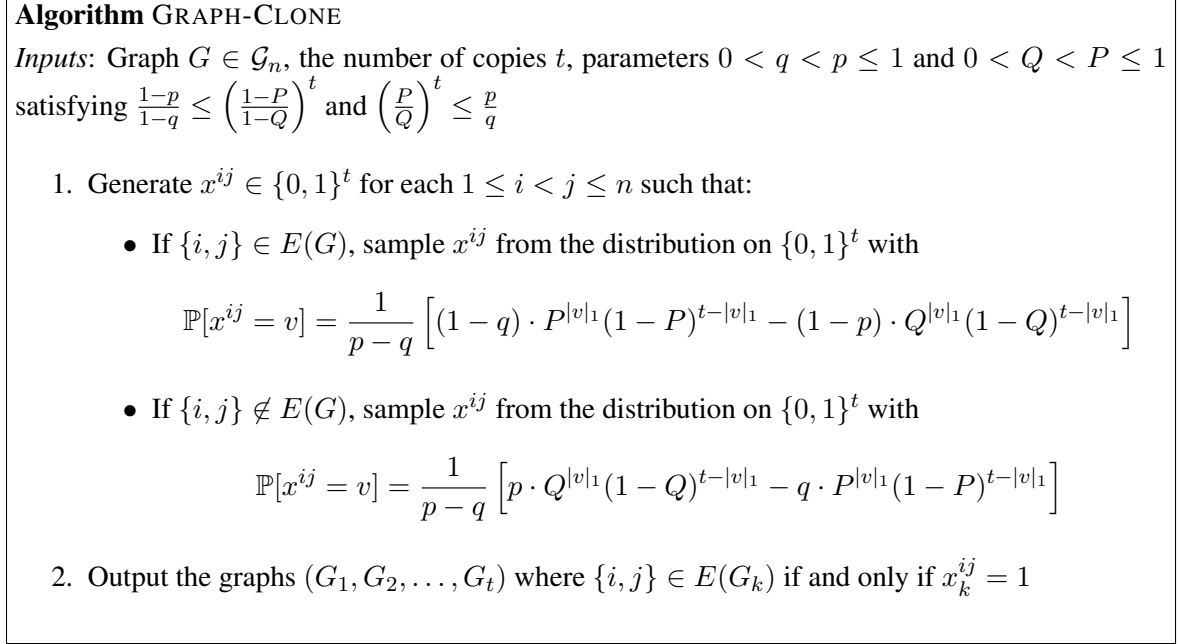


Figure 6: Subroutine GRAPH-CLONE for producing independent samples from planted graph problems.

$\mathcal{M}_d(\mathcal{P}, \mathcal{Q}, k)$. Given a k -subset $S \subseteq [d]$, let $\mathcal{M}_d((\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq d}, S)$ denote $\mathcal{M}_d((\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq d}, k)$ conditioned on the selection of S as the planted index set. Define $\mathcal{M}_d(\mathcal{P}, \mathcal{Q}, S)$ analogously. Similarly, let $\mathcal{V}_N(\mathcal{P}, \mathcal{Q}, S)$ denotes the distribution of vectors $V \in X^N$ with independent entries and $V_i \sim \mathcal{P}$ if $i \in S$ and $V_i \sim \mathcal{Q}$ if $i \notin S$. Let $\mathcal{V}_N(\mathcal{P}, \mathcal{Q}, k)$ denote the mixture of $\mathcal{V}_N(\mathcal{P}, \mathcal{Q}, S)$ induced by choosing S uniformly at random from all k -subsets of $[N]$. Let $\mathcal{G}(n, S, p, q)$ denote the distribution of planted dense subgraph instances from $\mathcal{G}(n, k, p, q)$ conditioned on the subgraph being planted on the vertex set S where $|S| = k$. Given an algorithm \mathcal{A} with k steps, let \mathcal{A}_i denote the i th step of \mathcal{A} and \mathcal{A}_{i-j} denote the procedure formed by steps i through j . Each time this notation is used, we clarify the intended initial and final variables when \mathcal{A}_i and \mathcal{A}_{i-j} are viewed as Markov kernels.

D.1. Graph Cloning

We begin with the subroutine GRAPH-CLONE which produces several independent samples from a planted subgraph problems given a single sample. This procedure is a simple generalization of PDS-CLONING in Section 10 of [Brennan et al. \(2018\)](#) and can be viewed as an exact variant of a simple multivariate rejection kernel to products of Bernoulli random variables.

Lemma 14 (Graph Cloning) *Let $t \in \mathbb{N}$, $0 < q < p \leq 1$ and $0 < Q < P \leq 1$ satisfy that*

$$\frac{1-p}{1-q} \leq \left(\frac{1-P}{1-Q}\right)^t \quad \text{and} \quad \left(\frac{P}{Q}\right)^t \leq \frac{p}{q}$$

Then the algorithm $\mathcal{A} = \text{GRAPH-CLONE}$ runs in $\text{poly}(t, n)$ time and satisfies that for each $S \subseteq [n]$,

$$\mathcal{A}(\mathcal{G}(n, q)) \sim \mathcal{G}(n, Q)^{\otimes t} \quad \text{and} \quad \mathcal{A}(\mathcal{G}(n, S, p, q)) \sim \mathcal{G}(n, S, P, Q)^{\otimes t}$$

Proof Let $R_0, R_1 : \{0, 1\}^t \rightarrow \mathbb{R}$ be given by

$$\begin{aligned} R_0(v) &= \frac{1}{p-q} \left[p \cdot Q^{|v|_1} (1-Q)^{t-|v|_1} - q \cdot P^{|v|_1} (1-P)^{t-|v|_1} \right] \\ R_1(v) &= \frac{1}{p-q} \left[(1-q) \cdot P^{|v|_1} (1-P)^{t-|v|_1} - (1-p) \cdot Q^{|v|_1} (1-Q)^{t-|v|_1} \right] \end{aligned}$$

Now observe that for each $v \in \{0, 1\}^t$, the fact that $P > Q$ implies that

$$\frac{1-p}{1-q} \leq \left(\frac{1-P}{1-Q} \right)^t \leq \frac{P^{|v|_1} (1-P)^{t-|v|_1}}{Q^{|v|_1} (1-Q)^{t-|v|_1}} \leq \left(\frac{P}{Q} \right)^t \leq \frac{p}{q}$$

which implies that $R_0(v) \geq 0$ and $R_1(v) \geq 0$ for each $v \in \{0, 1\}^t$. Furthermore, we have that

$$\sum_{v \in \{0, 1\}^t} R_0(v) = \sum_{v \in \{0, 1\}^t} R_1(v) = 1$$

which implies that R_0 and R_1 are well-defined probability mass functions. Also observe that

$$\begin{aligned} (1-p) \cdot R_0(v) + p \cdot R_1(v) &= P^{|v|_1} (1-P)^{t-|v|_1} \\ (1-q) \cdot R_0(v) + q \cdot R_1(v) &= Q^{|v|_1} (1-Q)^{t-|v|_1} \end{aligned}$$

Therefore it follows that if $\mathbf{1}_{\{i,j\} \in E(G)} \sim \text{Bern}(p)$, then $x^{ij} \sim \text{Bern}(P)^{\otimes t}$ and if $\mathbf{1}_{\{i,j\} \in E(G)} \sim \text{Bern}(q)$, then $x^{ij} \sim \text{Bern}(Q)^{\otimes t}$. Since the edge indicators $\mathbf{1}_{\{i,j\} \in E(G)}$ are independent in each of $\mathcal{G}(n, q)$ and $\mathcal{G}(n, S, p, q)$, this implies that $(G_1, G_2, \dots, G_t) \sim \mathcal{G}(n, Q)^{\otimes t}$ if $G \sim \mathcal{G}(n, q)$ and $(G_1, G_2, \dots, G_t) \sim \mathcal{G}(n, P)^{\otimes t}$ if $G \sim \mathcal{G}(n, k, p, q)$, completing the proof of the lemma. \blacksquare

D.2. Planting Diagonals by Embedding as a Principal Minor

The next two lemmas are a key technical component in the analysis of TO-SUBMATRIX. Specifically, they are crucial to showing the correctness of Step 2 in TO-SUBMATRIX, which plants missing diagonal entries while randomly embedding entries derived from the adjacency matrix of the input instance as a principal minor into a larger matrix to hide the planted entries in total variation. We remark that the applications of Cauchy-Schwarz reducing the proof of the second lemma to bounding χ^2 divergences are unlikely to be tight. However, the resulting bounds are sufficient for our purposes. We also remark that Lemma 16 can be proven by directly bounding sums of differences of binomial coefficients. Instead, our approach yields more elegant computations, as carried out in Lemma 15, and can be generalized to bound sums of random variables beyond binomial distributions.

Lemma 15 *Suppose that \mathcal{P} and \mathcal{Q} are probability distributions on a measurable space (X, \mathcal{B}) where \mathcal{P} is absolutely continuous with respect to \mathcal{Q} . Then for any positive integers k and m with $k^2 \cdot \chi^2(\mathcal{P} \parallel \mathcal{Q}) \leq m$, we have that*

$$\chi^2(\mathcal{V}_m(\mathcal{P}, \mathcal{Q}, k) \parallel \mathcal{Q}^{\otimes m}) \leq \frac{2k^2 \cdot \chi^2(\mathcal{P} \parallel \mathcal{Q})}{m}$$

Proof Let $f : X \rightarrow [0, \infty)$ be the Radon-Nikodym derivative $f = \frac{d\mathcal{P}}{d\mathcal{Q}}$. Note that $\mathcal{V}_m(\mathcal{P}, \mathcal{Q}, k)$ is absolutely continuous with respect to $\mathcal{Q}^{\otimes m}$ with Radon-Nikodym derivative

$$\frac{d\mathcal{V}_m(\mathcal{P}, \mathcal{Q}, k)}{d\mathcal{Q}^{\otimes m}}(x) = \mathbb{E}_{S \sim \mathcal{U}_{k,m}} \left[\prod_{i \in S} f(x_i) \right]$$

for each $x \in X^m$ where $\mathcal{U}_{k,m}$ is the uniform distribution on k -subsets of $[m]$. Now note that by Fubini's theorem

$$\begin{aligned} & \chi^2(\mathcal{V}_m(\mathcal{P}, \mathcal{Q}, k) \parallel \mathcal{Q}^{\otimes m}) + 1 \\ &= \mathbb{E}_{x \sim \mathcal{Q}^{\otimes m}} \left[\left(\frac{d\mathcal{V}_m(\mathcal{P}, \mathcal{Q}, k)}{d\mathcal{Q}^{\otimes m}}(x) \right)^2 \right] \\ &= \mathbb{E}_{x \sim \mathcal{Q}^{\otimes m}} \left[\mathbb{E}_{S \sim \mathcal{U}_{k,m}} \left[\prod_{i \in S} f(x_i) \right] \cdot \mathbb{E}_{T \sim \mathcal{U}_{k,m}} \left[\prod_{i \in S} f(x_i) \right] \right] \\ &= \mathbb{E}_{S, T \sim \mathcal{U}_{k,m}} \left[\mathbb{E}_{x \sim \mathcal{Q}^{\otimes m}} \left[\left(\prod_{i \in S} f(x_i) \right) \left(\prod_{i \in T} f(x_i) \right) \right] \right] \\ &= \mathbb{E}_{S, T \sim \mathcal{U}_{k,m}} \left[\prod_{i \in S \cap T} \mathbb{E}_{x_i \sim \mathcal{Q}} [f(x_i)^2] \prod_{i \in S \setminus T} \mathbb{E}_{x_i \sim \mathcal{Q}} [f(x_i)] \prod_{i \in T \setminus S} \mathbb{E}_{x_i \sim \mathcal{Q}} [f(x_i)] \right] \\ &= \mathbb{E}_{S, T \sim \mathcal{U}_{k,m}} \left[(1 + \chi^2(\mathcal{P} \parallel \mathcal{Q}))^{|S \cap T|} \right] \end{aligned}$$

where the last equality holds since $\mathbb{E}_{x_i \sim \mathcal{Q}} [f(x_i)] = 1$ and $1 + \chi^2(\mathcal{P} \parallel \mathcal{Q}) = \mathbb{E}_{x_i \sim \mathcal{Q}} [f(x_i)^2]$. We now apply an argument in [Addario-Berry et al. \(2010\)](#) to bound this last quantity. Observe that $|S \cap T| \sim \text{Hypergeometric}(m, k, k)$ and is identically distributed to $|[k] \cap S| = \sum_{i=1}^k \mathbf{1}_{\{i \in S\}}$. As shown in Section 3.2 of [Joag-Dev and Proschan \(1983\)](#), the variables $\mathbf{1}_{\{i \in S\}}$ are negatively associated which implies that

$$\begin{aligned} \mathbb{E}_{S, T \sim \mathcal{U}_{k,m}} \left[(1 + \chi^2(\mathcal{P} \parallel \mathcal{Q}))^{|S \cap T|} \right] &= \mathbb{E}_{S \sim \mathcal{U}_{k,m}} \left[\prod_{i=1}^k (1 + \chi^2(\mathcal{P} \parallel \mathcal{Q}))^{\mathbf{1}_{\{i \in S\}}} \right] \\ &\leq \prod_{i=1}^k \mathbb{E}_{S \sim \mathcal{U}_{k,m}} \left[(1 + \chi^2(\mathcal{P} \parallel \mathcal{Q}))^{k \cdot \mathbf{1}_{\{i \in S\}}} \right] \\ &= \left(\frac{k}{m} (1 + \chi^2(\mathcal{P} \parallel \mathcal{Q})) + 1 - \frac{k}{m} \right)^k \\ &= \left(1 + \frac{k}{m} \cdot \chi^2(\mathcal{P} \parallel \mathcal{Q}) \right)^k \\ &\leq \exp \left(\frac{k^2 \cdot \chi^2(\mathcal{P} \parallel \mathcal{Q})}{m} \right) \leq 1 + \frac{2k^2 \cdot \chi^2(\mathcal{P} \parallel \mathcal{Q})}{m} \end{aligned}$$

if $k^2 \cdot \chi^2(\mathcal{P} \parallel \mathcal{Q}) \leq m$. This completes the proof of the lemma. \blacksquare

We now apply this lemma to bound the total variation between support sizes of the matrix diagonals produced in TO-SUBMATRIX and the target matrix distributions.

Lemma 16 (Planting Diagonals) *Suppose that $0 < Q < P \leq 1$ and $N \geq \left(\frac{P}{Q} + \epsilon\right)n$ where $\epsilon > 0$. Let $k \leq n$ satisfy that*

$$k \leq \frac{Q\epsilon n}{2} \quad \text{and} \quad \frac{k^2}{N} \leq \min \left\{ \frac{Q}{1-Q}, \frac{1-Q}{Q} \right\}$$

Let $t_1 \sim \text{Bin}(k, P)$, $t_2 \sim \text{Bin}(n - k, P)$ and $t_3 \sim \text{Bin}(N, Q)$ be independent and set $t_4 = \max\{t_3 - t_1 - t_2, 0\}$. Then it holds that

$$\begin{aligned} d_{\text{TV}}(\mathcal{L}(t_1, t_2 + t_4), \text{Bin}(k, P) \otimes \text{Bin}(N - k, Q)) &\leq 4 \cdot \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right) + \sqrt{\frac{k^2(1-Q)}{2NQ}} \\ &\quad + \sqrt{\frac{k^2 Q}{2N(1-Q)}} \\ d_{\text{TV}}(\mathcal{L}(t_1 + t_2 + t_4), \text{Bin}(N, Q)) &\leq 4 \cdot \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right) \end{aligned}$$

Proof First consider applying Lemma 15 to $\mathcal{P} = \delta_0$ and $\mathcal{Q} = \text{Bern}(Q)$. Note that since $Q \in (0, 1)$, this choice of \mathcal{P} is absolutely continuous with respect to \mathcal{Q} . We have that if K and M are such that $K^2 Q \leq M(1 - Q)$ then

$$\chi^2(\mathcal{V}_M(\delta_0, \text{Bern}(Q), K) \parallel \text{Bern}(Q)^{\otimes M}) \leq \frac{2K^2 Q}{M(1 - Q)}$$

Now by Cauchy-Schwarz and the data-processing property in Fact B.1 on taking the sum of the entries of the vectors, we have that

$$\begin{aligned} d_{\text{TV}}(\text{Bin}(M - K, Q), \text{Bin}(M, Q)) &\leq d_{\text{TV}}(\mathcal{V}_M(\delta_0, \text{Bern}(Q), K), \text{Bern}(Q)^{\otimes M}) \\ &\leq \frac{1}{2} \sqrt{\chi^2(\mathcal{V}_M(\delta_0, \text{Bern}(Q), K) \parallel \text{Bern}(Q)^{\otimes M})} \\ &\leq \sqrt{\frac{K^2 Q}{2M(1 - Q)}} \end{aligned}$$

Now apply Lemma 15 to $\mathcal{P} = \delta_1$ and $\mathcal{Q} = \text{Bern}(Q)$. By the same argument, if $K^2(1 - Q) \leq MQ$, then we have that

$$d_{\text{TV}}(K + \text{Bin}(M - K, Q), \text{Bin}(M, Q)) \leq \sqrt{\frac{K^2(1 - Q)}{2MQ}}$$

Combining these two inequalities and applying the triangle inequality in Fact B.1 yields that

$$\begin{aligned} d_{\text{TV}}(K' + \text{Bin}(M - K, Q), \text{Bin}(M, Q)) &\leq d_{\text{TV}}(K' + \text{Bin}(M - K, Q), \text{Bin}(M - K + K', Q)) \\ &\quad + d_{\text{TV}}(\text{Bin}(M - K + K', Q), \text{Bin}(M, Q)) \\ &\leq \sqrt{\frac{K'^2(1 - Q)}{2(M - K + K')Q}} + \sqrt{\frac{(K - K')^2 Q}{2M(1 - Q)}} \end{aligned}$$

as long as $K' \leq K$ and it holds that

$$K'^2 \leq \frac{(M - K + K')Q}{1 - Q} \quad \text{and} \quad (K - K')^2 \leq \frac{M(1 - Q)}{Q}$$

Note that both of these inequalities are satisfied if $K^2/M \leq \min\{Q^{-1}(1 - Q), Q(1 - Q)^{-1}\}$. Standard Chernoff bounds imply that the binomial distribution satisfies the following concentration inequalities

$$\begin{aligned} \mathbb{P}[\text{Bin}(M, r) > (1 + \delta)rM] &\leq \exp\left(-\frac{\delta^2}{2 + \delta} \cdot rM\right) \\ \mathbb{P}[\text{Bin}(M, r) < (1 - \delta)rM] &\leq \exp\left(-\frac{\delta^2}{2} \cdot rM\right) \end{aligned}$$

for $\delta > 0$. These inequalities can be derived by standard Chernoff bounds. Now observe that if $t_3 \geq QN - \frac{Q\epsilon n}{2} + \frac{m}{2}$ and $t_2 \leq P(n - k) + \frac{Q\epsilon n}{2} - \frac{m}{2}$ hold, then $t_3 \geq m + t_2$ since $QN \geq (P + Q\epsilon)n$. Therefore we have that

$$\begin{aligned} \mathbb{P}[t_4 \neq t_3 - t_1 - t_2 | t_1 = m] &= \mathbb{P}[t_3 < m + t_2] \\ &\leq \mathbb{P}\left[t_3 < QN - \frac{Q\epsilon n}{2} + \frac{m}{2}\right] + \mathbb{P}\left[t_2 > P(n - k) + \frac{Q\epsilon n}{2} - \frac{m}{2}\right] \\ &\leq \mathbb{P}\left[t_3 < QN - \frac{Q\epsilon n}{4}\right] + \mathbb{P}\left[t_2 > P(n - k) + \frac{Q\epsilon n}{4}\right] \\ &\leq \exp\left(-\frac{1}{2} \left(\frac{\epsilon n}{4N}\right)^2 \cdot QN\right) \\ &\quad + \exp\left(-\frac{1}{2 + \frac{Q\epsilon n}{4P(n - k)}} \cdot \left(\frac{Q\epsilon n}{4P(n - k)}\right)^2 \cdot P(n - k)\right) \\ &\leq \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right) + \exp\left(-\frac{1}{8P(n - k) + Q\epsilon n} \cdot \frac{Q^2 \epsilon^2 n^2}{4}\right) \\ &\leq 2 \cdot \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right) \end{aligned}$$

The first inequality above is a union bound, the second inequality holds since $2m \leq 2k \leq Q\epsilon n$ and the last inequality holds since $8P(n - k) + Q\epsilon n \leq 8QN$. Marginalizing this bound over t_1 yields that

$$\mathbb{P}[t_4 \neq t_3 - t_1 - t_2] = \mathbb{E}_{m \sim \mathcal{L}(t_1)} \mathbb{P}[t_4 \neq t_3 - t_1 - t_2 | t_1 = m] \leq 2 \cdot \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right)$$

Now by the conditioning property in Fact B.1, we have

$$\begin{aligned} d_{\text{TV}}(\mathcal{L}(t_1 + t_2 + t_4), \mathcal{L}(t_3 | t_1 + t_2 + t_4 = t_3)) &\leq \mathbb{P}[t_1 + t_2 + t_4 \neq t_3] \\ d_{\text{TV}}(\mathcal{L}(t_3), \mathcal{L}(t_3 | t_1 + t_2 + t_4 = t_3)) &\leq \mathbb{P}[t_1 + t_2 + t_4 \neq t_3] \end{aligned}$$

Since $t_3 \sim \text{Bin}(N, Q)$, the triangle inequality in Fact B.1 implies that

$$d_{\text{TV}}(\mathcal{L}(t_1 + t_2 + t_4), \text{Bin}(N, Q)) \leq 2 \cdot \mathbb{P}[t_1 + t_2 + t_4 \neq t_3] \leq 4 \cdot \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right)$$

which proves the second inequality in the lemma. Now observe that by the conditioning property in Fact B.1, we have that

$$\begin{aligned} d_{\text{TV}}(\mathcal{L}(t_2 + t_4|t_1 = m), \mathcal{L}(t_2 + t_4|t_1 = m, t_1 + t_2 + t_4 = t_3)) &\leq \mathbb{P}[t_4 \neq t_3 - t_1 - t_2|t_1 = m] \\ d_{\text{TV}}(\mathcal{L}(t_3), \mathcal{L}(m + t_2 + t_4|t_1 = m, t_1 + t_2 + t_4 = t_3)) &\leq \mathbb{P}[t_4 \neq t_3 - t_1 - t_2|t_1 = m] \end{aligned}$$

Note that $t_3 \sim \text{Bin}(N, Q)$ is independent of t_1 and thus $\mathcal{L}(t_3) = \mathcal{L}(t_3|t_1 = m)$. Applying the inequality derived above using Lemma 15 with $M = N$, $K = k$ and $K' = m$ yields that

$$\begin{aligned} d_{\text{TV}}(\mathcal{L}(t_3), m + \text{Bin}(N - k, Q)) &\leq \sqrt{\frac{m^2(1-Q)}{2(N-k+m)Q}} + \sqrt{\frac{(k-m)^2Q}{2N(1-Q)}} \\ &\leq \sqrt{\frac{k^2(1-Q)}{2NQ}} + \sqrt{\frac{k^2Q}{2N(1-Q)}} \end{aligned}$$

as long as $k^2/N \leq \min\{Q^{-1}(1-Q), Q(1-Q)^{-1}\}$. Applying the triangle inequality twice now yields that

$$\begin{aligned} &d_{\text{TV}}(\mathcal{L}(t_2 + t_4|t_1 = m), \text{Bin}(N - k, Q)) \\ &= d_{\text{TV}}(\mathcal{L}(m + t_2 + t_4|t_1 = m), m + \text{Bin}(N - k, Q)) \\ &\leq 2 \cdot \mathbb{P}[t_4 \neq t_3 - t_1 - t_2|t_1 = m] + \sqrt{\frac{k^2(1-Q)}{2NQ}} + \sqrt{\frac{k^2Q}{2N(1-Q)}} \\ &\leq 4 \cdot \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right) + \sqrt{\frac{k^2(1-Q)}{2NQ}} + \sqrt{\frac{k^2Q}{2N(1-Q)}} \end{aligned}$$

Now note that by the conditioning on a random variable property in Fact B.1, we have that

$$\begin{aligned} &d_{\text{TV}}(\mathcal{L}(t_1, t_2 + t_4), \text{Bin}(k, P) \otimes \text{Bin}(N - k, Q)) \\ &\leq \mathbb{E}_{m \sim \text{Bin}(k, P)} d_{\text{TV}}(\mathcal{L}(t_2 + t_4|t_1 = m), \text{Bin}(N - k, Q)) \end{aligned}$$

Combining this with the inequality derived above completes the proof of the lemma. \blacksquare

D.3. Proof of Theorem 13

We now combine the lemmas in the previous two sections and Lemma 11 to prove Theorem 13. First consider the case where $G \sim \mathcal{G}(n, R, p, q)$ where $R \subseteq [n]$ satisfies $|R| = k$. In the first part of the proof of this proposition, let $M^1 = \mathcal{A}_{1-2}(G)$ be the matrix M^1 after Steps 1 and 2 in \mathcal{A} . First observe by AM-GM that

$$\sqrt{pq} \leq \frac{p+q}{2} = 1 - \frac{(1-p) + (1-q)}{2} \leq 1 - \sqrt{(1-p)(1-q)}$$

Let $Q = 1 - \sqrt{(1-p)(1-q)} + \mathbf{1}_{\{p=1\}}(\sqrt{q} - 1)$. If $p \neq 1$, then it follows that $P = p > Q$, $\frac{1-p}{1-q} = \left(\frac{1-P}{1-Q}\right)^2$ and the inequality above rearranges to $\left(\frac{P}{Q}\right)^2 \leq \frac{p}{q}$. If $p = 1$, then $Q = \sqrt{q}$,

the inequality $\frac{1-p}{1-q} \leq \left(\frac{1-p}{1-Q}\right)^2$ holds trivially and $\left(\frac{p}{Q}\right)^2 = \frac{p}{q}$. Applying Lemma 14 with $t = 2$ therefore yields that $(G_1, G_2) \sim \mathcal{G}(n, R, p, Q)^{\otimes 2}$. Let $U = \pi^{-1}(R)$ be the subset of $[N]$ that the dense subgraph vertices are mapped to in Step 2, on choosing S and π . Let $R' \subseteq [N]$ be a fixed subset with $|R'| = k$. Observe that the matrix M^1 in Step 2 conditioned on $U = R'$ has independent off-diagonal entries satisfying $M_{ij}^1 \sim \text{Bern}(p)$ if $i, j \in R'$ and $M_{ij}^1 \sim \text{Bern}(Q)$ otherwise, matching the off-diagonal distribution of $\mathcal{M}_N(\text{Bern}(p), \text{Bern}(Q), R')$. Furthermore these entries are independent of the diagonal entries of M^1 . Thus the tensorization property in Fact B.1 implies that

$$\begin{aligned} d_{\text{TV}}(\mathcal{L}(M^1|U = R'), \mathcal{M}_N(\text{Bern}(p), \text{Bern}(Q), R')) \\ = d_{\text{TV}}(\mathcal{L}(\text{diag}(M^1)|U = R'), \mathcal{V}_N(\text{Bern}(p), \text{Bern}(Q), R')) \end{aligned}$$

Fix some subset $S' \subseteq [N]$ with $|S'| = n$. Now observe that conditioned on $S = S'$, the entries M_{ii}^1 with $i \in S'$ are i.i.d. distributed as $\text{Bern}(p)$ since the number of $i \in S'$ with $M_{ii}^1 = 1$ is $s_1 \sim \text{Bin}(n, p)$ and $(M_{ii}^1 : i \in S')$ is exchangeable. Therefore, conditioned on $U = R'$ and $S = S'$, the entries of $\text{diag}(M^1)$ are distributed as $(M_{ii}^1 : i \in S') \sim \text{Bern}(p)^{\otimes n}$ and $(M_{ii}^1 : i \notin S')$ is exchangeable with support of size $|T_2| = \max\{s_2 - s_1, 0\}$ where s_1 is the size of the support of $(M_{ii}^1 : i \in S')$ and $s_2 \sim \text{Bin}(N, Q)$ is sampled independently. Since S is chosen uniformly at random, conditioned on $U = R'$, the elements of $S \setminus R'$ are a uniformly at random chosen subset of $[N] \setminus R'$ of size $n - k$. Thus relaxing the conditioning to only $U = R'$ yields that the entries of $\text{diag}(M^1)$ are distributed as $(M_{ii}^1 : i \in R') \sim \text{Bern}(p)^{\otimes k}$ and $(M_{ii}^1 : i \notin R')$ is exchangeable with support of size $t_2 + |T_2| = t_2 + \max\{s_2 - t_1 - t_2, 0\}$ where t_1 is the size of the support of $(M_{ii}^1 : i \notin R')$ and $t_2 \sim \text{Bin}(n - k, p)$ is sampled independently.

Note that the distributions of $\mathcal{L}(\text{diag}(M^1)|U = R')$ and $\mathcal{V}_N(\text{Bern}(p), \text{Bern}(Q), R')$ restricted to the indices in R' and $[N] \setminus R'$ are each exchangeable. Therefore conditioning on the pair of support sizes within R' and $[N] \setminus R'$ and applying the conditioning property in Fact B.1 yields that

$$\begin{aligned} d_{\text{TV}}(\mathcal{L}(\text{diag}(M^1)|U = R'), \mathcal{V}_N(\text{Bern}(p), \text{Bern}(Q), R')) \\ = d_{\text{TV}}(\mathcal{L}(t_1, t_2 + \max\{s_2 - t_1 - t_2, 0\}), \text{Bin}(k, p) \otimes \text{Bin}(N - k, Q)) \\ \leq 4 \cdot \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right) + \sqrt{\frac{k^2(1-Q)}{2QN}} + \sqrt{\frac{k^2Q}{2N(1-Q)}} \end{aligned}$$

by Lemma 16. Applying the conditioning property in Fact B.1 to conditioning on R and $U = R'$ now yields that

$$\begin{aligned} d_{\text{TV}}(\mathcal{L}(\mathcal{A}_{1-2}(\mathcal{G}(n, k, p, q))), \mathcal{M}_N(\text{Bern}(p), \text{Bern}(Q), k)) \\ \leq \mathbb{E}_{R \sim \mathcal{U}_{k,N}} \mathbb{E}_U [d_{\text{TV}}(\mathcal{L}(\text{diag}(M^1)|U), \mathcal{V}_N(\text{Bern}(p), \text{Bern}(Q), U))] \\ \leq 4 \cdot \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right) + \sqrt{\frac{k^2(1-Q)}{2QN}} + \sqrt{\frac{k^2Q}{2N(1-Q)}} \end{aligned}$$

where $\mathcal{U}_{k,N}$ is the uniform distribution on the k -subsets of $[N]$. Now let \mathcal{A}_3 denote Step 3 of \mathcal{A} with input M^1 and output M^2 . Let $M^1 \sim \mathcal{M}_N(\text{Bern}(p), \text{Bern}(Q), R')$ and $M^2 = \mathcal{A}_3(M^1)$. Consider also conditioning on the permutation $\tau = \tau'$ where τ' is a fixed permutation of $[N\ell]$. Let

$U_s = \tau'(\{s\ell + 1, s\ell + 2, \dots, (s + 1)\ell\})$ for each $0 \leq s < N$ and note that $\text{MRK}_{st} = \text{MRK}_{U_s, U_t}$. Now applying Lemma 11 to the rejection kernels MRK_{st} yields that

$$\begin{aligned} d_{\text{TV}} \left(\text{MRK}_{st}(\text{Bern}(p)), \bigotimes_{(i,j) \in U_s \times U_t} \mathcal{P}_{ij}^n \right) &\leq \Delta \quad \text{and} \\ d_{\text{TV}} \left(\text{MRK}_{st}(\text{Bern}(Q)), \bigotimes_{(i,j) \in U_s \times U_t} \mathcal{Q}_{ij}^n \right) &\leq \Delta \end{aligned}$$

Now let $V = \bigcup_{i \in R'} U_{i-1}$ be the set of indices of $[N\ell]$ that R' is mapped to and let M' be sampled as $M' \sim \mathcal{M}_{N\ell}((\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq N\ell}, V)$. The tensorization property in Fact B.1 now yields that

$$\begin{aligned} &d_{\text{TV}}(\mathcal{L}(M^2 | \tau = \tau'), \mathcal{M}_{N\ell}((\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq N\ell}, V)) \\ &\leq \sum_{i, j=1}^N d_{\text{TV}}(\mathcal{L}(M_{ab}^2 : a \in U_{i-1}, b \in U_{j-1}), \mathcal{L}(M'_{ab} : a \in U_{i-1}, b \in U_{j-1})) \\ &= \sum_{(i,j) \in R'^2} d_{\text{TV}} \left(\text{RK}_{(i-1)(j-1)}(\text{Bern}(p)), \bigotimes_{(a,b) \in U_{i-1} \times U_{j-1}} \mathcal{P}_{ab}^n \right) \\ &\quad + \sum_{(i,j) \notin R'^2} d_{\text{TV}} \left(\text{RK}_{(i-1)(j-1)}(\text{Bern}(Q)), \bigotimes_{(a,b) \in U_{i-1} \times U_{j-1}} \mathcal{Q}_{ab}^n \right) \\ &\leq N^2 \cdot \Delta \end{aligned}$$

Now note that when $\tau = \tau'$ is chosen uniformly at random, the set V is a uniformly at random chosen $k\ell$ -subset of $[N\ell]$. Applying the conditioning property in Fact B.1 to conditioning on R' and $\tau = \tau'$ now yields that

$$d_{\text{TV}}(\mathcal{A}_3(\mathcal{M}_N(\text{Bern}(p), \text{Bern}(Q), k)), \mathcal{M}_{N\ell}((\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq N\ell}, k\ell)) \leq N^2 \cdot \Delta$$

Applying Lemma 10 to the steps \mathcal{A}_{1-2} and \mathcal{A}_3 with the sequence of distributions $\mathcal{P}_0 = \mathcal{G}(n, k, p, q)$, $\mathcal{P}_{1-2} = \mathcal{M}_N(\text{Bern}(p), \text{Bern}(Q), k)$ and $\mathcal{P}_3 = \mathcal{M}_{N\ell}((\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq N\ell}, k\ell)$ yields that

$$\begin{aligned} d_{\text{TV}}(\mathcal{A}(\mathcal{G}(n, k, p, q)), \mathcal{M}_{N\ell}((\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq N\ell}, k\ell)) &\leq N^2 \cdot \Delta + 4 \cdot \exp\left(-\frac{Q\epsilon^2 n^2}{32N}\right) \\ &\quad + \sqrt{\frac{k^2(1-Q)}{2QN}} + \sqrt{\frac{k^2 Q}{2N(1-Q)}} \end{aligned}$$

We now follow an analogous and simpler argument to analyze the case $G \sim \mathcal{G}(n, q)$. Let $M^1 = \mathcal{A}_{1-2}(G)$ and note that $(G_1, G_2) \sim \mathcal{G}(n, Q)^{\otimes 2}$ by Lemma 14. Therefore the entries of M^1 are distributed as $M_{ij}^1 \sim_{\text{i.i.d.}} \text{Bern}(Q)$ for all $i \neq j$ independently of $\text{diag}(M^1)$, which is an exchangeable distribution on $\{0, 1\}^N$ with support size $s_1 + \max\{s_2 - s_1, 0\}$ where $s_1 \sim \text{Bin}(n, p)$ and $s_2 \sim \text{Bin}(N, Q)$. Applying the tensorization and conditioning properties in Fact B.1 as in the

previous case yields that

$$\begin{aligned} d_{\text{TV}} \left(\mathcal{L}(\mathcal{A}_{1-2}(\mathcal{G}(n, q))), \text{Bern}(Q)^{\otimes N \times N} \right) &= d_{\text{TV}} \left(\mathcal{L}(\text{diag}(M^1)), \text{Bern}(Q)^{\otimes N} \right) \\ &= d_{\text{TV}} \left(\mathcal{L}(s_1 + \max\{s_2 - s_1, 0\}), \text{Bin}(N, Q) \right) \\ &\leq 4 \cdot \exp \left(-\frac{Q\epsilon^2 n^2}{32N} \right) \end{aligned}$$

by Lemma 16. Conditioning on $\tau = \tau'$ and applying the tensorization property in Fact B.1 yields

$$\begin{aligned} &d_{\text{TV}} \left(\mathcal{L} \left(\mathcal{A}_3 \left(\text{Bern}(Q)^{\otimes N \times N} \right) \Big|_{\tau = \tau'} \right), \mathcal{M}_{N\ell} \left((\mathcal{Q}_{ij})_{1 \leq i, j \leq N\ell} \right) \right) \\ &\leq \sum_{i, j=1}^N d_{\text{TV}} \left(\text{MRK}_{(i-1)(j-1)} \left(\text{Bern}(Q) \right), \bigotimes_{(a, b) \in U_{i-1} \times U_{j-1}} \mathcal{Q}_{ab}^n \right) \leq N^2 \cdot \Delta \end{aligned}$$

Applying the conditioning property in Fact B.1 to conditioning on $\tau = \tau'$ now yields that

$$d_{\text{TV}} \left(\mathcal{L} \left(\mathcal{A}_3 \left(\text{Bern}(Q)^{\otimes N \times N} \right) \right), \mathcal{M}_{N\ell} \left((\mathcal{Q}_{ij})_{1 \leq i, j \leq N\ell} \right) \right) \leq N^2 \cdot \Delta$$

Applying Lemma 10 to \mathcal{A}_{1-2} and \mathcal{A}_3 with distributions $\mathcal{P}_0 = \mathcal{G}(n, q)$, $\mathcal{P}_{1-2} = \text{Bern}(Q)^{\otimes N \times N}$ and $\mathcal{P}_3 = \mathcal{M}_{N\ell} \left((\mathcal{Q}_{ij})_{1 \leq i, j \leq N\ell} \right)$ yields that

$$d_{\text{TV}} \left(\mathcal{L}(\mathcal{G}(n, q)), \mathcal{M}_{N\ell} \left((\mathcal{Q}_{ij})_{1 \leq i, j \leq N\ell} \right) \right) \leq N^2 \cdot \Delta + 4 \cdot \exp \left(-\frac{Q\epsilon^2 n^2}{32N} \right)$$

which completes the proof of the theorem.

Appendix E. Computational Barriers in Submatrix Detection

E.1. Computational Lower Bounds from Our Average-Case Reduction

The average-case reduction from planted dense subgraph in the previous section implies lower bounds for a more general heteroskedastic version of submatrix detection where the pairs of planted and noise distributions are allowed to vary from entry to entry, that we now formally define in the notation from the previous section.

Definition 17 (Heteroskedastic Symmetric Index Set Submatrix Detection) *Given computable pairs $(\mathcal{P}_{ij}, \mathcal{Q}_{ij})$ for $1 \leq i, j \leq n$ over a common measurable space (X, \mathcal{B}) , define $\text{HSSD}(n, k, (\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq n})$ to have observation $M \in X^{n \times n}$ and hypotheses*

$$H_0 : M \sim \mathcal{M}_n \left((\mathcal{Q}_{ij})_{1 \leq i, j \leq n} \right) \quad \text{and} \quad H_1 : M \sim \mathcal{M}_n \left((\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq n}, k \right)$$

The reduction TO-SUBMATRIX from the previous section yields the following lower bounds for HSSD based on the PDS conjecture. We state the implied lower bounds when $k = \Omega(\sqrt{n})$ and $k = o(\sqrt{n})$ separately in the next two theorems.

Theorem 18 (Heteroskedastic PDS Lower Bounds when $k = o(\sqrt{n})$) Let $k = o(\sqrt{n})$, let $0 < q < p \leq 1$ be fixed constants and let $(\mathcal{P}_{ij}, \mathcal{Q}_{ij})$ be a computable pairs over (X, \mathcal{B}) for each $1 \leq i, j \leq n$ such that

$$\sup_{1 \leq i, j \leq n} \mathbb{P}_{X \sim \mathcal{D}_{ij}} \left[\log \frac{d\mathcal{P}_{ij}}{d\mathcal{Q}_{ij}}(X) \notin \left[\log \left(\frac{1-p}{1-q} \right), \log \left(\frac{p}{q} \right) \right] \right] = o(n^{-2})$$

under both the settings $\mathcal{D}_{ij} = \mathcal{P}_{ij}$ and $\mathcal{D}_{ij} = \mathcal{Q}$ for each $i, j \in [n]$. Then assuming the PDS conjecture at densities $0 < q < p \leq 1$, there is no randomized polynomial time algorithm solving $\text{HSSD}(n, k, (\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq n})$ with asymptotic Type I+II error less than one.

Proof Consider applying Lemma 9 and Theorem 13 with $\ell = 1$ and starting planted dense subgraph instance with subgraph size k , n vertices and densities $0 < q < p \leq 1$. Excluding $n^2 \cdot \Delta$, all of the terms in both of the total variation upper bounds in Theorem 13 are $o(1)$ since $k^2 = o(n)$ and $Q < P$ are constants. It suffices to show that $n^2 \cdot \Delta = o(1)$. By the definition of Δ , we have that $\Delta = \max_{i, j \in [n]} \Delta_{ij}$ since $\ell = 1$. Now consider the definition of Δ_{ij} in Lemma 11. The condition in the theorem statement above implies that $\mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \notin S] = o(n^{-2})$ and $\mathbb{P}_{x \sim \mathcal{P}_n^*}[x \notin S] = o(n^{-2})$ and in particular that they both are at most $\frac{1}{4}(p-q)$. By the same argument as in the proof of Lemma 12, this is sufficient to imply that $\Delta_{ij} = o(n^{-2})$. Since this holds for each $i, j \in [n]$, we have that $\Delta = o(n^{-2})$, which proves the theorem. \blacksquare

Theorem 19 (Heteroskedastic PDS Lower Bounds when $k = \Omega(\sqrt{n})$) Let $k = \Omega(\sqrt{n})$, let $0 < q < p \leq 1$ be fixed constants and let $(\mathcal{P}_{ij}, \mathcal{Q}_{ij})$ be a computable pairs over (X, \mathcal{B}) for each $1 \leq i, j \leq n$. Suppose that there is some $m = \omega(k^2/n)$ with $m = o(n)$ such that

$$\sup_{S, T \subseteq [n]: |S|=|T|=m} \mathbb{P}_{X_{S \times T} \sim \mathcal{D}_{S \times T}} \left[\sum_{i \in S} \sum_{j \in T} \log \frac{d\mathcal{P}_{ij}}{d\mathcal{Q}_{ij}}(X_{ij}) \notin \left[\log \left(\frac{1-p}{1-q} \right), \log \left(\frac{p}{q} \right) \right] \right] = o\left(\frac{m^2}{n^2}\right)$$

under both the settings $\mathcal{D}_{S \times T} = \otimes_{i \in S} \otimes_{j \in T} \mathcal{P}_{ij}$ and $\mathcal{D}_{S \times T} = \otimes_{i \in S} \otimes_{j \in T} \mathcal{Q}_{ij}$ for each S, T . Then assuming the PDS conjecture at densities $0 < q < p \leq 1$, there is no randomized polynomial time algorithm solving $\text{HSSD}(n, k, (\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq n})$ with asymptotic Type I+II error less than one.

Proof Consider applying Lemma 9 and Theorem 13 with $\ell = m$ and starting planted dense subgraph instance with subgraph size $\lfloor k/m \rfloor$, $\lfloor n/m \rfloor$ vertices and densities $0 < q < p \leq 1$. Since $m = \omega(k^2/n)$, it follows that $k/m = o(\sqrt{n/m})$ and thus it suffices to show that the total variation upper bounds in Theorem 13 are $o(1)$. As in the proof of the previous theorem, this reduces to showing that $(n/m)^2 \cdot \Delta = o(1)$. For each pair $S, T \subseteq [n]$ with $|S| = |T| = m$, consider the definition of $\Delta_{S, T}$ in Lemma 11. We have that $\mathbb{P}_{x \sim \mathcal{Q}_n^*}[x \notin S] = o(n^{-2})$ and $\mathbb{P}_{x \sim \mathcal{P}_n^*}[x \notin S] = o(m^2/n^2)$ by the guarantees in the theorem statement. Since $m = o(n)$, these probabilities are at most $\frac{1}{4}(p-q)$ for large n . By the same reasoning as in the previous theorem, we have that $\Delta = o(m^2/n^2)$, proving the theorem. \blacksquare

We remark that we ignored issues of divisibility in the previous theorem statement, reducing to an instance with $m \lfloor n/m \rfloor$ vertices and submatrix size $m \lfloor k/m \rfloor$ instead of exactly n and m . This can be resolved by taking all of n, k, m, ℓ to be powers of two without affecting their sizes by more

than a factor of 2. Constructing a sequence of indices with these parameters is enough to rule out polynomial time algorithms given our forms of the PC and PDS conjectures.

From this point forward, we will restrict our attention to the homoskedastic formulation of submatrix detection that we have so far focused on. However, we first remark that these general heteroskedastic lower bounds can be specialized to imply hardness for submatrix problem with dependences between entries induced by natural column-wise and row-wise mixtures. Let $\mathcal{P}(\theta)$ be a family of distributions indexed by θ such that $(\mathcal{P}(\theta), \mathcal{Q})$ is a computable pair for each $\theta \in \Theta$. Now consider the submatrix problem with $H_0 : M \sim \mathcal{Q}^{\otimes n \times n}$ and H_1 distribution formed as follows:

- Select a subset $S \subseteq [n]$ with $|S| = k$ uniformly at random
- Sample $\theta_i \sim_{\text{i.i.d.}} \mathcal{D}$ for each $i \in [n]$
- Sample $M_{ij} \sim \mathcal{P}(\theta_i)$ for each $i, j \in S$ and $M_{ij} \sim \mathcal{Q}$ otherwise independently

for some distribution \mathcal{D} on Θ . For example, if $\mathcal{P}(\theta) = \mathcal{N}(\theta, 1)$, $\mathcal{Q} = \mathcal{N}(0, 1)$ and \mathcal{D} is normally distributed, this model has row-wise dependences resembling sparse PCA. Now suppose that an algorithm \mathcal{A} solve this problem with asymptotic Type I+II error ϵ , then there must be a deterministic choice of the θ_i such that \mathcal{A} solves the problem with asymptotic Type I+II error ϵ . When θ_i are deterministic, this problem is exactly HSSD($n, k, (\mathcal{P}_{ij}, \mathcal{Q}_{ij})_{1 \leq i, j \leq n}$) with $\mathcal{P}_{ij} = \mathcal{P}(\theta_i)$ and $\mathcal{Q}_{ij} = \mathcal{Q}$.

We now combine the heteroskedastic lower bounds in Theorems 18 and 19 with the MRK upper bound on Δ given in Lemma 12 to yield clean statements of the implied computational lower bounds for SSD based on the PDS and PC conjectures.

Corollary 20 (PDS Lower Bounds for Submatrix Detection) *Let $0 < q < p \leq 1$ be fixed constants and $(\mathcal{P}, \mathcal{Q})$ be a computable pair over (X, \mathcal{B}) such that either:*

- $k = \Omega(\sqrt{n})$ and $\frac{k^4}{n^2} \cdot d_{SKL}(\mathcal{P}, \mathcal{Q}) \rightarrow 0$ and the LLR between $(\mathcal{P}, \mathcal{Q})$ satisfies the LDP

$$E_{\mathcal{P}}(m) = \omega(m \log n) \quad \text{and} \quad E_{\mathcal{Q}}(-m) = \omega(m \log n)$$

for some positive m satisfying $\max\{d_{KL}(\mathcal{Q} \parallel \mathcal{P}), d_{KL}(\mathcal{P} \parallel \mathcal{Q})\} \leq m = o(n^2/k^4)$ and where the second inequality is only necessary if $p \neq 1$

- $k = o(\sqrt{n})$, $d_{KL}(\mathcal{Q} \parallel \mathcal{P}) < \log\left(\frac{1-q}{1-p}\right)$ and $d_{KL}(\mathcal{P} \parallel \mathcal{Q}) < \log\left(\frac{p}{q}\right)$ and the LLR between $(\mathcal{P}, \mathcal{Q})$ satisfies the LDP

$$E_{\mathcal{P}}\left(\log\left(\frac{p}{q}\right)\right) \geq 2 \log n + \omega(1) \quad \text{and} \quad E_{\mathcal{Q}}\left(\log\left(\frac{1-p}{1-q}\right)\right) \geq 2 \log n + \omega(1)$$

where the second inequality is only necessary if $p \neq 1$

Then assuming the PDS conjecture at densities $0 < q < p \leq 1$, there is no randomized polynomial time algorithm solving SSD($n, k, \mathcal{P}, \mathcal{Q}$) with asymptotic Type I+II error less than one.

Proof We first consider the case where $k = \Omega(\sqrt{n})$. Consider the reduction in Theorem 19 with blow-up factor $\ell = \omega(k^2/n)$ where ℓ is chosen so that $\ell^{-2} = \omega(m)$ where m is the positive constant in the statement of the corollary. Applying Lemma 12 to the MRK with blow-up factor ℓ^2 yields

$$\Delta \leq \frac{3 \exp(-\ell^2 \cdot E_{\mathcal{P}}(c_+ \ell^{-2})) + 3 \exp(-\ell^2 \cdot E_{\mathcal{Q}}(c_- \ell^{-2}))}{p - q}$$

where $c_+ = \log\left(\frac{p}{q}\right)$ and $c_- = \log\left(\frac{1-p}{1-q}\right)$. Since $E_{\mathcal{P}}(\lambda)$ is convex and minimized at $d_{\text{KL}}(\mathcal{Q}||\mathcal{P})$, it follows that since $c_+\ell^{-2} \geq m$ we have that

$$\frac{E_{\mathcal{P}}(c_+\ell^{-2})}{c_+\ell^{-2} - d_{\text{KL}}(\mathcal{Q}||\mathcal{P})} \geq \frac{E_{\mathcal{P}}(m)}{m - d_{\text{KL}}(\mathcal{Q}||\mathcal{P})} = \omega(\log n)$$

Therefore $E_{\mathcal{P}}(c_+\ell^{-2}) = \omega(\ell^{-2} \log n)$. A symmetric argument shows that $E_{\mathcal{Q}}(c_-\ell^{-2}) = \omega(\ell^{-2} \log n)$. Now it follows that $\Delta = o(n^{-2})$, which yields the first statement on applying Theorem 19. Now consider the case where $k = o(\sqrt{n})$. Applying Lemma 12 to the MRK with blow-up factor 1 to the reduction in Theorem 18 yields that

$$\Delta \leq \frac{3 \exp(-E_{\mathcal{P}}(c_+)) + 3 \exp(-E_{\mathcal{Q}}(c_-))}{p - q} = o(n^{-2})$$

by the given conditions. Combining this with Theorem 18 completes the proof. \blacksquare

Note that the constraints on $E_{\mathcal{Q}}$ are no longer necessary if $p = 1$. The next corollary states our PC lower bounds and is a restatement of the main theorem on computational lower bounds from Section 3.

Corollary 21 (PC Lower Bounds for Submatrix Detection) *Let $p \in (0, 1)$ be a fixed constant and $(\mathcal{P}, \mathcal{Q})$ be a computable pair over (X, \mathcal{B}) such that either:*

- $k = \Omega(\sqrt{n})$ and $\frac{k^4}{n^2} \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) \rightarrow 0$ and the LLR between $(\mathcal{P}, \mathcal{Q})$ satisfies the LDP

$$E_{\mathcal{P}}(m) \geq \omega(m \log n)$$

for some positive m with $d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) \leq m = o(n^2/k^4)$

- $k = o(\sqrt{n})$ and $d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) < \log p^{-1}$ and the LLR between $(\mathcal{P}, \mathcal{Q})$ satisfies the LDP

$$E_{\mathcal{P}}(\log p^{-1}) \geq 2 \log n + \omega(1)$$

Then assuming the PC conjecture at density p , there is no randomized polynomial time algorithm solving $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$ with asymptotic Type I+II error less than one.

We now give another corollary of Theorems 18 and 19 yielding a slight variation of these lower bounds for submatrix detection based on the PC conjecture. This corollary implies that the lower bounds stated in Section 3.1 hold for all $(\mathcal{P}, \mathcal{Q})$ in the universality class UC-A. Note that unlike the previous corollary which yielded clean lower bounds given the PC conjecture for any fixed p , we deduce the desired lower bounds up to a factor of $n^{\epsilon(p)}$ where $\epsilon(p)$ tends to zero with p . The proof of this corollary is very similar to that of Corollary 20, making crucial use of the convexity of $E_{\mathcal{P}}$ and the definition of UC-A.

Corollary 22 (Computational Lower Bounds for UC-A) *Suppose that $(\mathcal{P}, \mathcal{Q})$ is a computable pair in UC-A. Fix any $\epsilon \in (0, 1)$ and suppose that either:*

- $k = \Omega(\sqrt{n})$ and $d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) = o\left(\frac{n^{2-\epsilon}}{k^4}\right)$ or

- $k = o(\sqrt{n})$ and $d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) = o(n^{-\epsilon})$

Then there is a sufficiently small $p = p(\epsilon) > 0$ such that assuming the PC conjecture at density p , there is no randomized polynomial time algorithm solving $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$ with asymptotic Type I+II error less than one.

Proof We first consider the case where $k = \Omega(\sqrt{n})$. Consider the reduction in Theorem 19 with blow-up factor $\ell = \Theta(k^2/n^{1-\epsilon/4})$. Observe that by the given assumption, we have that $n^{\epsilon/2} \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) = o(\ell^{-2})$. Applying Lemma 12 to the MRK with blow-up factor ℓ^2 yields

$$\Delta \leq \frac{3 \exp(-\ell^2 \cdot E_{\mathcal{P}}(\ell^{-2} \log p^{-1}))}{1-p}$$

Since $E_{\mathcal{P}}(\lambda)$ is convex and minimized at $d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})$, it follows that since $\ell^{-2} \log p^{-1} \geq n^{\epsilon/2} \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})$ we have that

$$\frac{E_{\mathcal{P}}(\ell^{-2} \log p^{-1})}{\ell^{-2} \log p^{-1} - d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})} \geq \frac{E_{\mathcal{P}}(n^{\epsilon/2} \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}))}{n^{\epsilon/2} \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) - d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})} \geq c \cdot \log n$$

for some constant $c > 0$ since $(\mathcal{P}, \mathcal{Q})$ is in UC-A. Using the fact that $d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}) \leq \frac{1}{2} \ell n^{-2} \log p^{-1}$ for sufficiently large n , we have that

$$\Delta \leq \exp\left(\frac{c}{2} \log p^{-1} \cdot \log n\right) = o(n^{-2})$$

if p is taken to be sufficiently small. The first statement follows on applying Theorem 19. Now consider the case where $k = o(\sqrt{n})$. Applying Lemma 12 to the MRK with blow-up factor 1 to the reduction in Theorem 18 yields that

$$\Delta \leq \frac{3 \exp(-E_{\mathcal{P}}(\log p^{-1}))}{1-q}$$

by the given conditions. We now apply a similar convexity step, noting that since $n^{-\epsilon} \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \leq \log p^{-1}$ for sufficiently large n , we have that

$$\frac{E_{\mathcal{P}}(\log p^{-1})}{\log p^{-1} - d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})} \geq \frac{E_{\mathcal{P}}(n^{\epsilon} \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}))}{n^{\epsilon} \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) - d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})} \geq c' \cdot \log n$$

for some $c' > 0$. Now it similarly follows that $\Delta \leq \exp\left(\frac{c'}{2} \log p^{-1} \cdot \log n\right) = o(n^{-2})$ if p is taken to be sufficiently small. Combining this with Theorem 18 completes the proof. \blacksquare

E.2. Polynomial Time Test Statistics for Submatrix Detection

In this section, we show algorithmic upper bounds for submatrix detection using two simple test statistics that can be computed in polynomial time. Given a computable pair of distributions $(\mathcal{P}, \mathcal{Q})$ over the measurable space (X, \mathcal{B}) and a matrix $M \in X^{n \times n}$, define

$$T_{\text{sum}}(M) = \frac{1}{n^2} \sum_{i,j=1}^n \log \frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij})$$

$$T_{\text{max}}(M) = \max_{1 \leq i,j \leq n} \log \frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij})$$

Note that both T_{sum} and T_{max} can be computed in $O(n^2 \cdot \mathcal{T})$ time where the Radon-Nikodym derivative $\frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij})$ can be evaluated in $O(\mathcal{T})$ time. Given that $(\mathcal{P}, \mathcal{Q})$ is a computable pair, it follows that T_{sum} and T_{max} can be computed in polynomial time. We now show that thresholding these statistics solves the asymmetric detection problem ASD given sufficient LDPs for the LLR under each of \mathcal{Q} and \mathcal{P} . We begin with the sum test T_{sum} .

Proposition 23 (Sum Test) *Let M be an instance of $\text{ASD}(n, k, \mathcal{P}, \mathcal{Q})$ and let*

$$\tau_{\text{sum}} = -d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}) + \frac{k^2}{2n^2} \cdot d_{\text{SKL}}(\mathcal{P}, \mathcal{Q})$$

Suppose that $k \ll n$ and

$$\begin{aligned} E_{\mathcal{P}} \left(\frac{1}{2} \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \right) &= \omega(k^{-2}) \\ E_{\mathcal{Q}} \left(-\frac{2n^2 - k^2}{2n^2 - 2k^2} \cdot d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}) \right) &= \omega(n^{-2}) \\ E_{\mathcal{Q}} \left(-\frac{2n^2 - k^2}{2n^2} \cdot d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}) \right) &= \omega(n^{-2}) \end{aligned}$$

Then $\mathbb{P}_{H_0} [T_{\text{sum}}(M) \geq \tau_{\text{sum}}] \rightarrow 0$ and $\mathbb{P}_{H_1} [T_{\text{sum}}(M) < \tau_{\text{sum}}] \rightarrow 0$ as $n \rightarrow \infty$.

Proof Let $\tau' = -\frac{2n-k^2}{2n} \cdot d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})$ and note that $\tau' \leq \tau_{\text{sum}}$. Under H_0 , by a Chernoff bound we have that if $\lambda \geq 0$ then

$$\begin{aligned} \mathbb{P}_{H_0} [T_{\text{sum}}(M) \geq \tau_{\text{sum}}] &\leq \mathbb{P}_{H_0} [T_{\text{sum}}(M) \geq \tau'] \\ &= \mathbb{P}_{H_0} [\exp(n^2 \lambda \cdot T_{\text{sum}}(M)) \geq \exp(n^2 \lambda \cdot \tau')] \\ &\leq \exp(n^2 \cdot \psi_{\mathcal{Q}}(\lambda) - n^2 \lambda \cdot \tau') \end{aligned}$$

Since $\tau' \geq -d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})$, we may take $\lambda \geq 0$ so that $\lambda \cdot \tau' - \psi_{\mathcal{Q}}(\lambda)$ is arbitrarily close to $E_{\mathcal{Q}}(\tau')$. Therefore we have that

$$\mathbb{P}_{H_0} [T_{\text{sum}}(M) \geq \tau_{\text{sum}}] \leq \exp(-n^2 \cdot E_{\mathcal{Q}}(\tau')) = o(1)$$

Let $S', T' \subseteq [n]$ be the latent row and column indices of the planted part of M under H_1 . Note that $\tau_{\text{sum}} = (1 - \frac{k^2}{n^2}) \cdot \tau_1 + \frac{k^2}{n^2} \cdot \tau_2$ where $\tau_1 = -\frac{2n^2 - k^2}{2n^2 - 2k^2} \cdot d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})$ and $\tau_2 = \frac{1}{2} \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})$. Thus

under H_1 , by a Chernoff and union bound we have that if $\lambda_1, \lambda_2 \leq 0$

$$\begin{aligned}
 \mathbb{P}_{H_1} [T_{\text{sum}}(M) < \tau_{\text{sum}}] &\leq \mathbb{P}_{H_1} \left[\sum_{(i,j) \notin S' \times T'} \log \frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij}) < \left(1 - \frac{k^2}{n^2}\right) \cdot \tau_1 \right] \\
 &\quad + \mathbb{P}_{H_1} \left[\sum_{(i,j) \in S' \times T'} \log \frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij}) < \frac{k^2}{n^2} \cdot \tau_2 \right] \\
 &\leq \exp(-(n^2 - k^2)\lambda_1 \cdot \tau_1) \cdot \mathbb{E}_{H_1} \left[\exp \left(\lambda_1 \cdot \sum_{(i,j) \notin S' \times T'} \log \frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij}) \right) \right] \\
 &\quad + \exp(-k^2\lambda_2 \cdot \tau_2) \cdot \mathbb{E}_{H_1} \left[\exp \left(\lambda_2 \cdot \sum_{(i,j) \in S' \times T'} \log \frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij}) \right) \right] \\
 &= \exp((n^2 - k^2) \cdot \psi_{\mathcal{Q}}(\lambda_1) - (n^2 - k^2) \cdot \lambda_1 \cdot \tau_1) \\
 &\quad + \exp(k^2 \cdot \psi_{\mathcal{P}}(\lambda_2) - k^2 \cdot \lambda_2 \cdot \tau_2)
 \end{aligned}$$

Since $\tau_1 \leq -d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})$ and $\tau_2 \leq d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})$, we may take $\lambda_1 \leq 0$ and $\lambda_2 \leq 0$ so that $\lambda_1 \cdot \tau_1 - \psi_{\mathcal{Q}}(\lambda_1)$ is arbitrarily close to $E_{\mathcal{Q}}(\tau_1)$ and $\lambda_2 \cdot \tau_2 - \psi_{\mathcal{P}}(\lambda_2)$ is arbitrarily close to $E_{\mathcal{P}}(\tau_2)$. This yields that

$$\mathbb{P}_{H_1} [T_{\text{sum}}(M) < \tau_{\text{sum}}] \leq \exp(-(n^2 - k^2) \cdot E_{\mathcal{Q}}(\tau_1)) + \exp(-k^2 \cdot E_{\mathcal{P}}(\tau_2)) = o(1)$$

which completes the proof of the proposition. \blacksquare

Given an LDP for the LLR under \mathcal{Q} and \mathcal{P} , it also holds that T_{sum} solves the asymmetric detection problem.

Proposition 24 (Max Test) *Let M be an instance of $\text{ASD}(n, k, \mathcal{P}, \mathcal{Q})$ and suppose there is a $\tau_{\text{max}} \in (-d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}), d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}))$ with*

$$E_{\mathcal{Q}}(\tau_{\text{max}}) \geq 2 \log n + \omega(1) \quad \text{and} \quad E_{\mathcal{P}}(\tau_{\text{max}}) = \omega(1)$$

then $\mathbb{P}_{H_0} [T_{\text{max}}(M) \geq \tau_{\text{max}}] \rightarrow 0$ and $\mathbb{P}_{H_1} [T_{\text{max}}(M) < \tau_{\text{max}}] \rightarrow 0$ as $n \rightarrow \infty$.

Proof This follows from the same argument used to analyze the search test T_{search} in the proof of Proposition 30 applied with $k = 1$. \blacksquare

We now proceed to show that our algorithmic upper bounds hold for all computable pairs $(\mathcal{P}, \mathcal{Q})$ in UC-B. To do this, we will establish the following simple consequences for $(\mathcal{P}, \mathcal{Q})$ in UC-B. As mentioned in Section 3, the class UC-B is introduced in Hajek et al. (2017). The third property below is derived in Section 3 of Hajek et al. (2017).

Lemma 25 (Properties of UC-B) *Suppose that $(\mathcal{P}, \mathcal{Q})$ is in UC-A with constant $C \geq 1$. Then*

1. *It holds for all $\tau \in [-2C \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}), 0]$ that*

$$E_{\mathcal{P}}(d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) + \tau) \geq \frac{\tau^2}{4C \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})}$$

2. It holds for all $\tau \in [-2C \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P}), 2C \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P})]$ that

$$E_{\mathcal{Q}}(-d_{\text{KL}}(\mathcal{Q}||\mathcal{P}) + \tau) \geq \frac{\tau^2}{4C \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P})}$$

3. It holds that $d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) = \Theta(d_{\text{KL}}(\mathcal{Q}||\mathcal{P}))$.

Proof The fact that $(\mathcal{P}, \mathcal{Q})$ is in UC-B implies that

$$\begin{aligned} E_{\mathcal{P}}(d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) + \tau) &= \sup_{\lambda \in \mathbb{R}} \{(d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) + \tau) \cdot \lambda - \psi_{\mathcal{P}}(\lambda)\} \\ &\geq \sup_{\lambda \in [-1, 0]} \{\tau \cdot \lambda - C \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) \cdot \lambda^2\} \end{aligned}$$

Now set $\lambda = \frac{\tau}{2C \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q})} \in [-1, 0]$ and note that this implies that

$$E_{\mathcal{P}}(d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) + \tau) \geq \frac{\tau^2}{4C \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q})}$$

Similarly, we have that

$$E_{\mathcal{Q}}(-d_{\text{KL}}(\mathcal{Q}||\mathcal{P}) + \tau) \geq \sup_{\lambda \in [-1, 1]} \{\tau \cdot \lambda - C \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P}) \cdot \lambda^2\} \geq \frac{\tau^2}{4C \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P})}$$

on setting $\lambda = \frac{\tau}{2C \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P})} \in [-1, 1]$. Property 3 follows from Lemma 2 in [Hajek et al. \(2017\)](#), which shows that

$$\min\{d_{\text{KL}}(\mathcal{Q}||\mathcal{P}), d_{\text{KL}}(\mathcal{P}||\mathcal{Q})\} \geq \frac{1}{C} \cdot \max\{d_{\text{KL}}(\mathcal{Q}||\mathcal{P}), d_{\text{KL}}(\mathcal{P}||\mathcal{Q})\}$$

This implies that $d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) = \Theta(d_{\text{KL}}(\mathcal{Q}||\mathcal{P}))$. ■

We now combine these properties with Propositions 23 and 24 to show algorithmic achievability of the computational barriers shown above for SSD when $(\mathcal{P}, \mathcal{Q})$ is in UC-B.

Corollary 26 (Algorithmic Upper Bounds for UC-B) *Suppose that $(\mathcal{P}, \mathcal{Q})$ is a computable pair in UC-B. Then it follows that:*

- If $k = o(n)$, $k = \Omega(\sqrt{n})$ and $d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}) = \omega\left(\frac{n^2}{k^4}\right)$, then T_{sum} solves $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$.
- If $k = o(\sqrt{n})$ and $d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}) \geq c \cdot \log n$ for some sufficiently large constant $c > 0$, then T_{max} with $\tau_{\text{max}} = 0$ solves $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$.

Proof We begin with the first statement. By Proposition 23, it suffices to verify the lower bounds on $E_{\mathcal{P}}$ and $E_{\mathcal{Q}}$ in the statement of the proposition. To do this, we apply properties (1) and (2) in Lemma 25. Let $C \geq 1$ be the constant for which $(\mathcal{P}, \mathcal{Q})$ is in UC-B and observe that

$$\begin{aligned} E_{\mathcal{P}}\left(\frac{1}{2} \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q})\right) &\geq \frac{\left(-\frac{1}{2} \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q})\right)^2}{4C \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q})} = \frac{1}{16C} \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) = \omega(k^{-2}) \\ E_{\mathcal{Q}}\left(-\frac{2n^2 - k^2}{2n^2 - 2k^2} \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P})\right) &\geq \frac{\left(\frac{k^2}{2n^2 - 2k^2} \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P})\right)^2}{4C \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P})} = \frac{k^4}{16C(n^2 - k^2)^2} \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P}) = \omega(n^{-2}) \\ E_{\mathcal{Q}}\left(-\frac{2n^2 - k^2}{2n^2} \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P})\right) &\geq \frac{\left(\frac{k^2}{2n^2} \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P})\right)^2}{4C \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P})} = \frac{k^4}{16C \cdot n^4} \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P}) = \omega(n^{-2}) \end{aligned}$$

since $k = o(n)$, $d_{\text{KL}}(\mathcal{P}\|\mathcal{Q}) = \Theta(d_{\text{KL}}(\mathcal{Q}\|\mathcal{P}))$ and $d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}) = \omega\left(\frac{n^2}{k^4}\right)$. We now verify the second statement for $\tau_{\max} = 0$. It suffices to verify the two lower bounds on $E_{\mathcal{P}}$ and $E_{\mathcal{Q}}$ in the statement of Proposition 24. Note that

$$\begin{aligned} E_{\mathcal{P}}(0) &\geq \frac{(-d_{\text{KL}}(\mathcal{P}\|\mathcal{Q}))^2}{4C \cdot d_{\text{KL}}(\mathcal{P}\|\mathcal{Q})} = \frac{1}{4C} \cdot d_{\text{KL}}(\mathcal{P}\|\mathcal{Q}) \geq 3 \log n \\ E_{\mathcal{Q}}(0) &\geq \frac{(d_{\text{KL}}(\mathcal{Q}\|\mathcal{P}))^2}{4C \cdot d_{\text{KL}}(\mathcal{Q}\|\mathcal{P})} = \frac{1}{4C} \cdot d_{\text{KL}}(\mathcal{Q}\|\mathcal{P}) \geq 3 \log n \end{aligned}$$

for sufficiently large $c > 0$. This completes the proof of the corollary. \blacksquare

Appendix F. Statistical Limit of Submatrix Detection

In this section, we show information-theoretic lower bounds for our universal formulations of submatrix detection and provide a test statistic showing that this boundary is achievable.

F.1. Information-Theoretic Lower Bound for Submatrix Detection

Assuming that \mathcal{P} and \mathcal{Q} have finite χ^2 divergence, we can obtain the following information-theoretic lower bound for SSD with the distribution pair $(\mathcal{P}, \mathcal{Q})$. The proof uses a similar χ^2 divergence computation as in Lemma 15 and the information-theoretic lower bounds for planted dense subgraph shown in Hajek et al. (2015).

Theorem 27 *Suppose that \mathcal{P} and \mathcal{Q} are probability distributions on a measurable space (X, \mathcal{B}) where \mathcal{P} is absolutely continuous with respect to \mathcal{Q} . If $\chi^2(\mathcal{P}\|\mathcal{Q})$ is finite and satisfies that*

$$\chi^2(\mathcal{P}\|\mathcal{Q}) < \frac{1}{16e} \left(\frac{1}{n} \log \left(\frac{en}{k} \right) \wedge \frac{n^2}{k^4} \right)$$

then there is a function $\tau : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ such that $\lim_{t \rightarrow 0^+} \tau(t) = 0$ and

$$d_{\text{TV}}(\mathcal{M}_n(\mathcal{P}, \mathcal{Q}, k), \mathcal{Q}^{\otimes n \times n}) \leq \tau \left(\frac{\chi^2(\mathcal{P}\|\mathcal{Q})}{\frac{1}{n} \log \left(\frac{en}{k} \right) \wedge \frac{n^2}{k^4}} \right)$$

To prove this, we will need the following lemma of Hajek et al. (2015) bounding the moment generating function of a hypergeometric random variable squared.

Lemma 28 (Lemma 6 in Hajek et al. (2015)) *There exists a function $\tau_1 : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ satisfying that $\lim_{t \rightarrow 0^+} \tau_1(t) = 1$ such that for any $k \leq n$ the following holds: if $H \sim \text{Hypergeometric}(n, k, k)$ and $\lambda = \kappa \left(\frac{1}{k} \log \left(\frac{en}{k} \right) \wedge \frac{n^2}{k^4} \right)$ where $0 < \kappa < \frac{1}{16e}$ then*

$$\mathbb{E} [\exp(\lambda H^2)] \leq \tau_1(\kappa)$$

Using this upper bound, we now can prove the information-theoretic lower bounds for submatrix detection in Theorem 27.

Proof [Proof of Theorem 27] Let $f : X \rightarrow [0, \infty)$ be the Radon-Nikodym derivative $f = \frac{d\mathcal{P}}{d\mathcal{Q}}$. Observe that $\mathcal{M}_n(\mathcal{P}, \mathcal{Q}, k)$ can be written as the mixture

$$\mathcal{M}_n(\mathcal{P}, \mathcal{Q}, k) = \binom{n}{k}^{-1} \sum_{S \subseteq [n]: |S|=k} \mathcal{R}_S$$

where \mathcal{R}_S is the distribution of $n \times n$ matrices $M \in X^{n \times n}$ with independent entries such that $M_{ij} \sim \mathcal{P}$ if $i, j \in S$ and $M_{ij} \sim \mathcal{Q}$ otherwise. Note that \mathcal{R}_S is therefore absolutely continuous with respect to $\mathcal{Q}^{\otimes n \times n}$ with Radon-Nikodym derivative $\frac{d\mathcal{R}_S}{d\mathcal{Q}^{\otimes n \times n}}(M) = \prod_{i,j \in S} f(M_{ij})$ for each $M \in X^{n \times n}$. It follows that $\mathcal{M}_n(\mathcal{P}, \mathcal{Q}, k)$ is also absolutely continuous with respect to $\mathcal{Q}^{\otimes n \times n}$ with Radon-Nikodym derivative

$$\frac{d\mathcal{M}_n(\mathcal{P}, \mathcal{Q}, k)}{d\mathcal{Q}^{\otimes n \times n}}(M) = \binom{n}{k}^{-1} \sum_{S \subseteq [n]: |S|=k} \frac{d\mathcal{R}_S}{d\mathcal{Q}^{\otimes n \times n}}(M) = \mathbb{E}_{S \sim \mathcal{U}_{k,n}} \left[\prod_{i,j \in S} f(M_{ij}) \right]$$

for each $x \in X^{n \times n}$ where $\mathcal{U}_{k,n}$ is the uniform distribution on k -subsets of $[n]$. By Fubini's theorem,

$$\begin{aligned} & \chi^2(\mathcal{M}_n(\mathcal{P}, \mathcal{Q}, k) \parallel \mathcal{Q}^{\otimes n \times n}) + 1 \\ &= \mathbb{E}_{M \sim \mathcal{Q}^{\otimes n \times n}} \left[\left(\frac{d\mathcal{M}_n(\mathcal{P}, \mathcal{Q}, k)}{d\mathcal{Q}^{\otimes n \times n}}(M) \right)^2 \right] \\ &= \mathbb{E}_{M \sim \mathcal{Q}^{\otimes n \times n}} \left[\mathbb{E}_{S \sim \mathcal{U}_{k,n}} \left[\prod_{i,j \in S} f(M_{ij}) \right] \cdot \mathbb{E}_{T \sim \mathcal{U}_{k,n}} \left[\prod_{i,j \in T} f(M_{ij}) \right] \right] \\ &= \mathbb{E}_{S,T \sim \mathcal{U}_{k,n}} \left[\mathbb{E}_{M \sim \mathcal{Q}^{\otimes n \times n}} \left[\left(\prod_{i,j \in S} f(M_{ij}) \right) \left(\prod_{i,j \in T} f(M_{ij}) \right) \right] \right] \\ &= \mathbb{E}_{S,T \sim \mathcal{U}_{k,n}} \left[\prod_{i,j \in S \cap T} \mathbb{E}_{M_{ij} \sim \mathcal{Q}} [f(M_{ij})^2] \prod_{(i,j) \in S^2 \cup T^2 - (S \cap T)^2} \mathbb{E}_{M_{ij} \sim \mathcal{Q}} [f(M_{ij})] \right] \\ &= \mathbb{E}_{S,T \sim \mathcal{U}_{k,n}} \left[(1 + \chi^2(\mathcal{P} \parallel \mathcal{Q}))^{|S \cap T|^2} \right] \end{aligned}$$

where the last equality holds since $\mathbb{E}_{M_{ij} \sim \mathcal{Q}} [f(M_{ij})] = 1$ and $1 + \chi^2(\mathcal{P} \parallel \mathcal{Q}) = \mathbb{E}_{M_{ij} \sim \mathcal{Q}} [f(M_{ij})^2]$. Note that $H = |S \cap T| \sim \text{Hypergeometric}(n, k, k)$. Let τ_1 be the function in Lemma 28. The given bounds on $\chi^2(\mathcal{P} \parallel \mathcal{Q})$ imply that we can apply Lemma 28 with $\lambda = \chi^2(\mathcal{P} \parallel \mathcal{Q})$. Combining this with Cauchy-Schwarz and the fact that $1 + \chi^2(\mathcal{P} \parallel \mathcal{Q}) \leq \exp(\chi^2(\mathcal{P} \parallel \mathcal{Q}))$ yields that

$$\begin{aligned} 2 \cdot d_{\text{TV}}(\mathcal{M}_n(\mathcal{P}, \mathcal{Q}, k), \mathcal{Q}^{\otimes n \times n})^2 &\leq \chi^2(\mathcal{M}_n(\mathcal{P}, \mathcal{Q}, k) \parallel \mathcal{Q}^{\otimes n \times n}) \\ &= \mathbb{E}_{S,T \sim \mathcal{U}_{k,n}} \left[(1 + \chi^2(\mathcal{P} \parallel \mathcal{Q}))^{|S \cap T|^2} \right] - 1 \\ &\leq \mathbb{E} \left[\exp(H^2 \cdot \chi^2(\mathcal{P} \parallel \mathcal{Q})) \right] - 1 \\ &\leq \tau_1 \left(\frac{\chi^2(\mathcal{P} \parallel \mathcal{Q})}{\frac{1}{n} \log \left(\frac{en}{k} \right) \wedge \frac{n^2}{k^4}} \right) - 1 \end{aligned}$$

Setting $\tau = \sqrt{\frac{1}{2}(\tau_1 - 1)}$ which satisfies $\lim_{t \rightarrow 0^+} \tau(t) = 0$ completes the proof of the theorem. \blacksquare

Now using the fact that the minimum Type I+II error of a hypothesis testing problem between \mathcal{L}_0 and \mathcal{L}_1 is $1 - d_{\text{TV}}(\mathcal{L}_0, \mathcal{L}_1)$, we arrive at the following corollary providing a regime in which submatrix detection is statistically impossible.

Corollary 29 *Suppose that \mathcal{P} and \mathcal{Q} are probability distributions on a measurable space (X, \mathcal{B}) where \mathcal{P} is absolutely continuous with respect to \mathcal{Q} . If $\chi^2(\mathcal{P} \parallel \mathcal{Q})$ is finite and satisfies that*

$$\chi^2(\mathcal{P} \parallel \mathcal{Q}) = o\left(\frac{1}{k} \log\left(\frac{n}{k}\right) \wedge \frac{n^2}{k^4}\right)$$

then any test $\phi : X^{n \times n} \rightarrow \{0, 1\}$ has an asymptotic Type I+II error of at least one on $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$.

Note that this corollary implies that if $(\mathcal{P}, \mathcal{Q})$ is in UC-C and satisfies that $\chi^2(\mathcal{P} \parallel \mathcal{Q}) = O(d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}))$, then $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$ is information-theoretically impossible if

$$d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}) = o\left(\frac{1}{k} \log\left(\frac{n}{k}\right) \wedge \frac{n^2}{k^4}\right)$$

which matches the bounds in Section 3.1.

F.2. Search Test Statistic

In this section, we give a simple search test statistic showing statistical achievability. Given a computable pair of distributions $(\mathcal{P}, \mathcal{Q})$ over the measurable space (X, \mathcal{B}) and a matrix $M \in X^{n \times n}$, define

$$T_{\text{search}}(M) = \max_{S, T \subseteq [n]: |S|=|T|=k} \left(\frac{1}{k^2} \sum_{i \in S} \sum_{j \in T} \log \frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij}) \right)$$

Note that T_{search} can be computed in $O(n^{2k} \cdot \mathcal{T})$ time where the Radon-Nikodym derivative $\frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij})$ can be evaluated in $O(\mathcal{T})$ time. We now show that thresholding this statistics solves the asymmetric detection problem ASD given sufficient LDPs for the LLR under each of \mathcal{Q} and \mathcal{P} . We begin with the sum test T_{sum} .

Proposition 30 (Search Test) *Let M be an instance of $\text{ASD}(n, k, \mathcal{P}, \mathcal{Q})$ and suppose there is a $\tau_{\text{search}} \in (-d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}), d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}))$ with*

$$E_{\mathcal{Q}}(\tau_{\text{search}}) \geq \frac{2}{k} \log\left(\frac{n}{k}\right) + \omega(k^{-2}) \quad \text{and} \quad E_{\mathcal{P}}(\tau_{\text{search}}) = \omega(k^{-2})$$

then $\mathbb{P}_{H_0}[T_{\text{search}}(M) \geq \tau_{\text{search}}] \rightarrow 0$ and $\mathbb{P}_{H_1}[T_{\text{search}}(M) < \tau_{\text{search}}] \rightarrow 0$ as $n \rightarrow \infty$.

Proof By a union bound and Chernoff bound, we have that for if $\lambda \geq 0$ then

$$\begin{aligned}
 & \mathbb{P}_{H_0} [T_{\text{search}}(M) \geq \tau_{\text{search}}] \\
 &= \sum_{S, T \subseteq [n]: |S|=|T|=k} \mathbb{P}_{H_0} \left[\sum_{i \in S} \sum_{j \in T} \log \frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij}) \geq k^2 \cdot \tau_{\text{search}} \right] \\
 &\leq \binom{n}{k}^2 \cdot \mathbb{P}_{H_0} \left[\exp \left(\lambda \cdot \sum_{i,j=1}^k \log \frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij}) \right) \geq \exp(\lambda \cdot k^2 \cdot \tau_{\text{search}}) \right] \\
 &\leq \exp \left(2 \log \binom{n}{k} + k^2 \cdot \psi_{\mathcal{Q}}(\lambda) - k^2 \lambda \cdot \tau_{\text{search}} \right)
 \end{aligned}$$

Since $\tau_{\text{search}} \in (-d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}), d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}))$, we may take $\lambda \geq 0$ so that $\lambda \cdot \tau_{\text{search}} - \psi_{\mathcal{Q}}(\lambda)$ is arbitrarily close to $E_{\mathcal{Q}}(\tau_{\text{search}})$. This implies that

$$\mathbb{P}_{H_0} [T_{\text{search}}(M) \geq \tau_{\text{search}}] \leq \exp \left(2k \cdot \log \binom{n}{k} - k^2 \cdot E_{\mathcal{Q}}(\tau_{\text{search}}) \right) = o(1)$$

since $\binom{n}{k} \leq \left(\frac{n}{k}\right)^k$. Let $S', T' \subseteq [n]$ be the latent row and column indices of the planted part of M under H_1 . Now it follows that for $\lambda \leq 0$ we have that

$$\begin{aligned}
 \mathbb{P}_{H_1} [T_{\text{search}}(M) < \tau_{\text{search}}] &= \mathbb{P}_{H_1} \left[\sum_{i \in S'} \sum_{j \in T'} \log \frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij}) < k^2 \cdot \tau_{\text{search}} \right] \\
 &\leq \mathbb{P}_{H_1} \left[\exp \left(\lambda \cdot \sum_{i \in S'} \sum_{j \in T'} \log \frac{d\mathcal{P}}{d\mathcal{Q}}(M_{ij}) \right) > \exp(\lambda \cdot k^2 \cdot \tau_{\text{search}}) \right] \\
 &\leq \exp \left(k^2 \cdot \psi_{\mathcal{P}}(\lambda) - k^2 \lambda \cdot \tau_{\text{search}} \right)
 \end{aligned}$$

Again, since $\tau_{\text{search}} \in (-d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}), d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}))$, we may take $\lambda \leq 0$ so that $\lambda \cdot \tau_{\text{search}} - \psi_{\mathcal{P}}(\lambda)$ is arbitrarily close to $E_{\mathcal{P}}(\tau_{\text{search}})$. Therefore

$$\mathbb{P}_{H_1} [T_{\text{search}}(M) < \tau_{\text{search}}] \leq \exp(-E_{\mathcal{P}}(\tau_{\text{search}})) = o(1)$$

which proves the desired result. \blacksquare

We now show that the search test solves $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$ for UC-B in $(\mathcal{P}, \mathcal{Q})$ in the parameter regime described in Section 3.1.

Corollary 31 (Statistically Achievability for UC-B) *Suppose that $(\mathcal{P}, \mathcal{Q})$ is a computable pair in UC-B. If $d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}) \geq \frac{c}{k} \log \left(\frac{n}{k}\right)$ for a sufficiently large constant $c > 0$, then T_{search} with $\tau_{\text{search}} = 0$ solves $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$.*

Proof It suffices to verify the lower bounds on $E_{\mathcal{P}}$ and $E_{\mathcal{Q}}$ in Proposition 30. Since $(\mathcal{P}, \mathcal{Q})$ is in UC-B, by Lemma 25 we have that

$$\begin{aligned}
 E_{\mathcal{P}}(0) &\geq \frac{(-d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}))^2}{4C \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})} = \frac{1}{4C} \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \geq \frac{3}{k} \log \left(\frac{n}{k}\right) \\
 E_{\mathcal{Q}}(0) &\geq \frac{(d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}))^2}{4C \cdot d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})} = \frac{1}{4C} \cdot d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}) \geq \frac{3}{k} \log \left(\frac{n}{k}\right)
 \end{aligned}$$

Since $\frac{1}{k} \log \binom{n}{k} = \omega(k^{-2})$, applying Proposition 30 now proves the corollary. \blacksquare

Appendix G. The Universality Classes UC-A, UC-B and UC-C

G.1. Universality Classes UC-A and UC-B

The universality class UC-B is discussed at length in Section 2.1 and 3 of Hajek et al. (2017), which introduces it as Assumption 2 in the context of their information-theoretic lower bounds for general submatrix recovery. They provide a means to check whether a pair $(\mathcal{P}, \mathcal{Q})$ belonging to an exponential family is in UC-B in Appendix B of Hajek et al. (2017). The discussion below shows that this method also checks membership in UC-A and UC-C.

Recall that a random variable X is sub-Gaussian if there are $a, b \in \mathbb{R}$ with $b > 0$ such that $\log \mathbb{E}[e^{\lambda X}] \leq a + \mathbb{E}[X] \cdot \lambda + b\lambda^2$ for all $\lambda \in \mathbb{R}$. Given a computable pair $(\mathcal{P}, \mathcal{Q})$, let $L(x) = \log \frac{d\mathcal{P}}{d\mathcal{Q}}(x)$ denote its LLR. The inequalities

$$\begin{aligned} \psi_{\mathcal{P}}(\lambda) - d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \cdot \lambda &\leq C \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \cdot \lambda^2 \quad \text{for all } \lambda \in [-1, 0] \\ \psi_{\mathcal{Q}}(\lambda) + d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}) \cdot \lambda &\leq C \cdot d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P}) \cdot \lambda^2 \quad \text{for all } \lambda \in [-1, 1] \end{aligned}$$

defining UC-B are exactly the inequalities required for the two distributions $L(X)$ where $X \sim \mathcal{P}$ and $L(X)$ where $X \sim \mathcal{Q}$ to be sub-Gaussian, but only required to hold for λ in restricted intervals. Thus UC-B is weaker than sub-Gaussianity of L under \mathcal{P} and \mathcal{Q} . We now observe that it similarly holds that UC-A is weaker than sub-Gaussianity of L under \mathcal{P} . If the sub-Gaussianity inequality holds for the following interval

$$\psi_{\mathcal{P}}(\lambda) - d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \cdot \lambda \leq C \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \cdot \lambda^2 \quad \text{for } \lambda \in [0, \Theta(\log n)]$$

then the same argument showing Property 1 in Lemma 25 shows that

$$E_{\mathcal{P}}((\lambda + 1) \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})) = \Omega(d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \cdot (\log n)^2)$$

for $\lambda = \Theta(\log n)$. The convexity of $E_{\mathcal{P}}$ and the fact that $E_{\mathcal{P}}(d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})) = 0$ implies that if n is large enough so that $n^\epsilon \geq \lambda + 1$, then

$$\frac{E_{\mathcal{P}}(n^\epsilon \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}))}{(n^\epsilon - 1) \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})} \geq \frac{E_{\mathcal{P}}((\lambda + 1) \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}))}{\lambda \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})} = \Omega(\log n)$$

and thus the condition needed for UC-A holds. It is also shown in Hajek et al. (2017) that $(\mathcal{P}, \mathcal{Q})$ with bounded LLR are in UC-B.

Lemma 32 (Lemma 1 in Hajek et al. (2017)) *If $|L| \leq B$ for some constant $B > 0$, then $(\mathcal{P}, \mathcal{Q})$ is in UC-B with constant $C = e^{5B}$.*

We now show that the three pairs of interest \mathcal{D}_{BC} , \mathcal{D}_{SP} and \mathcal{D}_{GP} introduced in Section 3 are in UC-A and UC-B. In Sections 2.1 and 3 of Hajek et al. (2017), it is shown that all three of these pairs lie in UC-B. Thus it suffices to verify that they lie in UC-A. Consider \mathcal{D}_{BC} where $\mathcal{P} = \mathcal{N}(\mu, 1)$ and $\mathcal{Q} = \mathcal{N}(0, 1)$. As shown in Section 2.1 of Hajek et al. (2017), we have that

$$E_{\mathcal{P}}(\theta) = \frac{1}{8} \left(\mu - \frac{2\theta}{\mu} \right)^2$$

Suppose that $\theta = n^\epsilon \cdot d_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) = \frac{1}{2} n^\epsilon \cdot \mu^2$ by a standard formula for the KL divergence between two Gaussians. Then it follows that

$$E_{\mathcal{P}}(\theta) = \frac{1}{8} (\mu - n^\epsilon \cdot \mu)^2 = \Theta(n^{2\epsilon} \mu^2)$$

Since $n^{2\epsilon} = \omega(\log n)$, it follows that \mathcal{D}_{BC} is in UC-A. In the following, we use several computations d_{KL} computations to be carried in Section G.2. As shown in Hajek et al. (2017), if $\mathcal{P} = \text{Bern}(p)$ and $\mathcal{Q} = \text{Bern}(q)$ then

$$E_{\mathcal{P}}(\theta) = D(\alpha \parallel p) \quad \text{where } \alpha = \frac{\theta + \log \frac{1-q}{1-p}}{\log \frac{p(1-q)}{q(1-p)}}$$

where $D(\cdot \parallel \cdot)$ is the binary entropy function. Letting $\theta = D(p \parallel q) + \tau$ yields that

$$E_{\mathcal{P}}(D(p \parallel q) + \tau) = D \left(p + \tau \left(\log \frac{p(1-q)}{q(1-p)} \right)^{-1} \parallel p \right)$$

When $p = cq = cn^{-\alpha}$ for some $c > 1$, it is not difficult to verify that if τ is such that $\tau = \Theta(n^\epsilon \cdot D(p \parallel q)) = \Theta(n^{\epsilon-\alpha})$ then

$$E_{\mathcal{P}}(D(p \parallel q) + \tau) = \Theta(n^{\epsilon-\alpha} \log n)$$

and thus \mathcal{D}_{sp} is in UC-A. Furthermore, if $p = n^{-\alpha} + \Theta(n^{-\gamma})$ and $q = n^{-\alpha}$ where $\gamma > \alpha > 0$, then $D(p \parallel q) = \Theta(n^{\alpha-2\gamma})$. Observe that $\log \frac{p(1-q)}{q(1-p)} = \Theta(n^{\alpha-\gamma})$. Now taking $\tau = \Theta(n^\epsilon \cdot D(p \parallel q)) = \Theta(n^{\epsilon+\alpha-2\gamma})$ yields that

$$E_{\mathcal{P}}(D(p \parallel q) + \tau) = \Theta \left(\frac{\tau^2}{p} \left(\log \frac{p(1-q)}{q(1-p)} \right)^{-2} \right) = \Theta(n^{2\epsilon+\alpha-2\gamma})$$

and since $n^\epsilon = \omega(\log n)$, it follows that \mathcal{D}_{GP} is in UC-A. We conclude this section by generalizing these computations to show that if a pair $(\mathcal{P}, \mathcal{Q})$ has bounded LLR then it is also in UC-A. The proof is similar to that of Lemma 1 in Hajek et al. (2017).

Lemma 33 *If $|L| \leq B$ for some constant $B > 0$, then $(\mathcal{P}, \mathcal{Q})$ is in UC-A.*

Proof First note that if $\lambda \in [0, \lambda_{\max}]$ then

$$\psi_{\mathcal{P}}''(\lambda) = \frac{\mathbb{E}_{\mathcal{P}}[L^2 \cdot \exp(\lambda L)] \cdot \psi_{\mathcal{P}}(\lambda) - \mathbb{E}_{\mathcal{P}}[L \cdot \exp(\lambda L)]^2}{\psi_{\mathcal{P}}(\lambda)^2} \leq \frac{\mathbb{E}_{\mathcal{P}}[L^2 \cdot \exp(\lambda L)]}{\mathbb{E}_{\mathcal{P}}[\exp(\lambda L)]} \leq e^{2B\lambda_{\max}} \cdot \mathbb{E}_{\mathcal{P}}[L^2]$$

As in Hajek et al. (2017), let $\phi(x) = e^x - x - 1$ and note that if $|x| \leq B$ then $\frac{1}{2}e^{-B}x^2 \leq \phi(x) \leq \frac{1}{2}e^Bx^2$ since ϕ is nonnegative, convex and satisfies $\phi(0) = \phi'(0) = 0$ and $\phi''(x) = e^x \in [e^{-B}, e^B]$ if $|x| \leq B$. Therefore we have that

$$\mathbb{E}_{\mathcal{P}}[L^2] = \mathbb{E}_{\mathcal{Q}}[L^2 \exp(L)] \leq e^B \cdot \mathbb{E}_{\mathcal{Q}}[L^2] \leq 2e^{2B} \cdot \mathbb{E}_{\mathcal{Q}}[\phi(L)] = 2e^{2B} \cdot d_{\text{KL}}(\mathcal{Q} \parallel \mathcal{P})$$

Applying Lemma 32 yields that $(\mathcal{P}, \mathcal{Q})$ is in UC-B and thus $d_{\text{KL}}(\mathcal{Q}||\mathcal{P}) \leq c \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q})$ for some $c > 0$ by Property 3 in Lemma 25. Thus $\psi''_{\mathcal{P}}(\lambda) \leq 2c \cdot e^{2B(\lambda_{\max}+1)} \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q})$ for all $\lambda \in [0, \lambda_{\max}]$. Combining this with $\psi'_{\mathcal{P}}(0) = d_{\text{KL}}(\mathcal{P}||\mathcal{Q})$ and $\psi_{\mathcal{P}}(0) = 0$ yields that

$$\psi_{\mathcal{P}}(\lambda) \leq d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) \cdot \lambda + c \cdot e^{2B(\lambda_{\max}+1)} \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) \cdot \lambda^2 \quad \text{for all } \lambda \in [0, \lambda_{\max}]$$

This inequality implies that

$$E_{\mathcal{P}}(d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) + \tau) \geq \sup_{\lambda \in [0, \lambda_{\max}]} \left\{ \tau \cdot \lambda - c \cdot e^{2B(\lambda_{\max}+1)} \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) \cdot \lambda^2 \right\} \geq \frac{\tau^2}{4c \cdot e^{2B(\lambda_{\max}+1)} d_{\text{KL}}(\mathcal{P}||\mathcal{Q})}$$

where the last inequality holds as long as

$$\frac{\tau}{2c \cdot e^{2B(\lambda_{\max}+1)} d_{\text{KL}}(\mathcal{P}||\mathcal{Q})} \leq \lambda_{\max}$$

Now take $\tau = (n^\epsilon - 1) \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q})$ and $\lambda_{\max} + 1 = \frac{1}{2B}[\epsilon \log n - \log \log n] + c'$ for some constant $c' > 0$. It follows that

$$\frac{\tau}{2c \cdot e^{2B(\lambda_{\max}+1)} d_{\text{KL}}(\mathcal{P}||\mathcal{Q})} \leq \frac{e^{-2Bc'}}{2c} \cdot \log n$$

which is at most λ_{\max} for a large enough choice of $c' = c'(B, c) > 0$. Furthermore, substituting this pair (τ, λ_{\max}) into the inequality above yields that $E_{\mathcal{P}}(n^\epsilon \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q})) = \Omega(n^\epsilon \cdot d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) \cdot \log n)$, which implies that $(\mathcal{P}, \mathcal{Q})$ is in UC-A. \blacksquare

G.2. Universality Class UC-C

The condition for $(\mathcal{P}, \mathcal{Q})$ to be in UC-C is also weaker than sub-Gaussianity of the LLR. Observe that if the sub-Gaussian inequality

$$\psi_{\mathcal{Q}}(\lambda) + d_{\text{KL}}(\mathcal{Q}||\mathcal{P}) \cdot \lambda \leq C \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P}) \cdot \lambda^2$$

holds for $\lambda = 2$, then $\chi^2(\mathcal{P}||\mathcal{Q}) = \psi_{\mathcal{Q}}(2) \leq (4C - 2) \cdot d_{\text{KL}}(\mathcal{Q}||\mathcal{P})$ and $(\mathcal{P}, \mathcal{Q})$ is in UC-C. In light of this, taking $\lambda_{\max} \geq 2$ in the proof of Lemma 33 also shows that $|L| \leq B$ implies that $(\mathcal{P}, \mathcal{Q})$ is in UC-C. For the sake of the exposition, we now verify directly that the three pairs \mathcal{D}_{BC} , \mathcal{D}_{SP} and \mathcal{D}_{GP} are in UC-C by explicitly computing their KL and χ^2 divergences.

(\mathcal{D}_{BC}) If $\mathcal{P} = \mathcal{N}(\mu, 1)$ and $\mathcal{Q} = \mathcal{N}(0, 1)$ where $\mu = n^{-\alpha}$ for some $\alpha > 0$, then we have that

$$\begin{aligned} \chi^2(\mathcal{P}||\mathcal{Q}) &= \frac{1}{2} \left(e^{\mu^2} - 1 \right) = \Theta(\mu^2) \\ d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}) &= d_{\text{KL}}(\mathcal{Q}||\mathcal{P}) + d_{\text{KL}}(\mathcal{P}||\mathcal{Q}) = \mu^2 \end{aligned}$$

as $\mu \rightarrow 0$, by well-known formulas for these divergences.

(\mathcal{D}_{SP}) If $\mathcal{P} = \text{Bern}(p)$ and $\mathcal{Q} = \text{Bern}(q)$ where $p = cq = cn^{-\alpha}$ for some constant $c > 1$ and $\alpha > 0$, then we have that

$$\begin{aligned}\chi^2(\mathcal{P}\|\mathcal{Q}) &= \frac{(p-q)^2}{q(1-q)} = \frac{(c-1)^2}{1-n^{-\alpha}} \cdot n^{-\alpha} = \Theta(n^{-\alpha}) \\ d_{\text{KL}}(\mathcal{Q}\|\mathcal{P}) &= -n^{-\alpha} \log c - (1-n^{-\alpha}) \log \left(1 - \frac{(c-1)n^{-\alpha}}{1-n^{-\alpha}} \right) \\ &= -n^{-\alpha} \log c - (1-n^{-\alpha}) \cdot \left(-\frac{(c-1)n^{-\alpha}}{1-n^{-\alpha}} + O(n^{-2\alpha}) \right) \\ &= (c-1-\log c)n^{-\alpha} + O(n^{-2\alpha}) = \Theta(n^{-\alpha}) \\ d_{\text{KL}}(\mathcal{P}\|\mathcal{Q}) &= (c^{-1}-1-\log c^{-1})cn^{-\alpha} + O(n^{-2\alpha}) = \Theta(n^{-\alpha})\end{aligned}$$

(\mathcal{D}_{GP}) If $\mathcal{P} = \text{Bern}(p)$ and $\mathcal{Q} = \text{Bern}(q)$ where $p = q + \Theta(n^{-\gamma})$ and $q = n^{-\alpha}$ for some constants $\gamma > \alpha > 0$, then we have that

$$\begin{aligned}\chi^2(\mathcal{P}\|\mathcal{Q}) &= \frac{(p-q)^2}{q(1-q)} = \Theta(n^{-2\gamma+\alpha}) \\ d_{\text{KL}}(\mathcal{Q}\|\mathcal{P}) &= -q \log \left(1 + \frac{p-q}{q} \right) - (1-q) \log \left(1 - \frac{p-q}{1-q} \right) \\ &= -q \left(\frac{p-q}{q} - \left(\frac{p-q}{q} \right)^2 + O\left(\frac{(p-q)^3}{q^3} \right) \right) \\ &\quad - (1-q) \left(-\frac{p-q}{1-q} - \left(\frac{p-q}{1-q} \right)^2 + O\left(\frac{(p-q)^3}{(1-q)^3} \right) \right) \\ &= \frac{(p-q)^2}{q(1-q)} + O(n^{-3\gamma+2\alpha}) = \Theta(n^{-2\gamma+\alpha}) \\ d_{\text{KL}}(\mathcal{P}\|\mathcal{Q}) &= \Theta(n^{-2\gamma+\alpha})\end{aligned}$$

These computations verify that all three pairs $(\mathcal{P}, \mathcal{Q})$ are in the universality class UC-C.

Appendix H. Further Questions

This work leaves a number of questions about submatrix detection and planted clique reductions unresolved. The following is an overview of some of these problems.

- **Weaker Universality Assumptions:** Can our required lower bound on $E_{\mathcal{P}}$ in our main computational lower bounds be relaxed? In other words, is there a reduction from planted clique or another conjecturally hard average-case problem to the general submatrix detection problem for a wider universality class of $(\mathcal{P}, \mathcal{Q})$?
- **Other Possible Computational Phase Diagrams:** Outside of our universality classes, are there any natural universality classes $(\mathcal{P}, \mathcal{Q})$ with different phase diagrams that can be characterized through average-case reductions? One example of a pair $(\mathcal{P}, \mathcal{Q})$ outside of our universality classes that we do not show hardness for is $\mathcal{P} = \text{Bern}(p)$ and $\mathcal{Q} = \text{Bern}(q)$




Figure 7: Computational and statistical barriers in the recovery variant of $\text{SSD}(n, k, \mathcal{P}, \mathcal{Q})$ under regularity assumptions on $(\mathcal{P}, \mathcal{Q})$. The red region is conjectured to be computationally hard but no PC reductions showing this hardness are known. Axes are parameterized as $k = \tilde{\Theta}(n^\beta)$ and $d_{\text{SKL}}(\mathcal{P}, \mathcal{Q}) = \tilde{\Theta}(n^{-\alpha})$.

where $p = n^{-\alpha}$ and $q = n^{-\beta}$ where $\beta > \alpha$. The graph variant of this submatrix detection problem corresponds to the log-density regime of planted dense subgraph and seems to obey a completely different phase diagram. Algorithms and conjectured hardness for this problem are discussed in [Bhaskara et al. \(2010\)](#); [Chlamtac et al. \(2012\)](#); [Chlamtáč et al. \(2017\)](#); [Chlamtáč and Manurangsi \(2018\)](#).

- Computational Lower Bounds for Submatrix Recovery:** Through similar detection-recovery reductions as in Section 10 of [Brennan et al. \(2018\)](#), our computational lower bounds for submatrix detection yields computational lower bounds for the general recovery variant. However, T_{sum} does not translate into a natural recovery algorithm and the computational barrier for recovery appears to be different from that of submatrix detection. This has left a region of the phase diagram with an unknown computational complexity. Semidefinite programming algorithms for recovery under regularity assumptions on $(\mathcal{P}, \mathcal{Q})$ were analyzed in [Hajek et al. \(2016b\)](#) meeting the polynomial time threshold shown in Figure 7. In a distributionally robust sub-Gaussian variant of submatrix recovery, planted clique lower bounds were shown by [Cai et al. \(2017\)](#). The known and open regions of the phase diagram for recovery are shown in Figure 7.