

Testing Positive Semi-Definiteness via Random Submatrices

Ainesh Bakshi*

CMU

abakshi@cs.cmu.edu

Nadiia Chepurko

MIT

nadiia@mit.edu

Rajesh Jayaram*

CMU

rkjayara@cs.cmu.edu

September 21, 2020

Abstract

We study the problem of testing whether a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with bounded entries ($\|\mathbf{A}\|_\infty \leq 1$) is positive semi-definite (PSD), or ϵ -far in Euclidean distance from the PSD cone, meaning that $\min_{\mathbf{B} \succeq 0} \|\mathbf{A} - \mathbf{B}\|_F^2 > \epsilon n^2$, where $\mathbf{B} \succeq 0$ denotes that \mathbf{B} is PSD. Our main algorithmic contribution is a non-adaptive tester which distinguishes between these cases using only $\tilde{O}(1/\epsilon^4)$ queries to the entries of \mathbf{A} .¹ If instead of the Euclidean norm we considered the distance in spectral norm, we obtain the “ ℓ_∞ -gap problem”, where \mathbf{A} is either PSD or satisfies $\min_{\mathbf{B} \succeq 0} \|\mathbf{A} - \mathbf{B}\|_2 > \epsilon n$. For this related problem, we give a $\tilde{O}(1/\epsilon^2)$ query tester, which we show is optimal up to $\log(1/\epsilon)$ factors. Both our testers randomly sample a collection of principal submatrices and check whether these submatrices are PSD. Consequentially, our algorithms achieve *one-sided error*: whenever they output that \mathbf{A} is not PSD, they return a certificate that \mathbf{A} has negative eigenvalues.

We complement our upper bound for PSD testing with Euclidean norm distance by giving a $\tilde{\Omega}(1/\epsilon^2)$ lower bound for any non-adaptive algorithm. Our lower bound construction is general, and can be used to derive lower bounds for a number of spectral testing problems. As an example of the applicability of our construction, we obtain a new $\tilde{\Omega}(1/\epsilon^4)$ sampling lower bound for testing the Schatten-1 norm with a $\epsilon n^{1.5}$ gap, extending a result of Balcan, Li, Woodruff, and Zhang [BLWZ19]. In addition, our hard instance results in new sampling lower bounds for estimating the Ky-Fan Norm, and the cost of rank- k approximations, i.e. $\|\mathbf{A} - \mathbf{A}_k\|_F^2 = \sum_{i>k} \sigma_i^2(\mathbf{A})$.

*Ainesh Bakshi and Rajesh Jayaram would like to thank the partial support from the Office of Naval Research (ONR) grant N00014-18-1-2562, and the National Science Foundation (NSF) under Grant No. CCF-1815840.

¹Throughout the paper, $\tilde{O}(\cdot)$ hides $\log(1/\epsilon)$ factors.

Contents

1	Introduction	1
1.1	Our Contributions	3
1.2	Connections to Optimization, Euclidean Metrics and Linear Algebra	6
1.3	Related work	7
1.4	Technical Overview	8
1.4.1	PSD Testing with ℓ_∞ Gap	8
1.4.2	PSD Testing with ℓ_2 Gap	11
1.4.3	Lower Bounds	12
2	Preliminaries	14
3	PSD Testing with ℓ_∞ Gap	15
3.1	Warm-up: a $O(1/\epsilon^3)$ algorithm	17
3.2	Variance Bounds	19
3.3	Improving the complexity to $\tilde{O}(1/\epsilon^2)$	23
3.3.1	Case 1: Varied Subsampling and Eigenvector Switching	24
3.3.2	Case 2: Spread Negative Mass and Main Theorem	29
4	PSD Testing with ℓ_2^2 Gap	31
4.1	Analysis of the Algorithm	33
5	Lower bounds	38
5.1	Lower Bound for PSD Testing with ℓ_∞ Gap	38
5.2	Lower Bound for PSD Testing with ℓ_2 Gap	39
5.2.1	$C_{n_1+n_2}$ is Subgraph Equivalent to $C_{n_1} \oplus C_{n_2}$	45
5.3	Lower Bounds for Schatten, Ky-Fan, and Tail Error Testing	49
5.4	Lower Bound For Estimating Ky-Fan of $\mathbf{A}\mathbf{A}^T$ via Submatrices	52
6	Conclusion	53
A	Proof of Eigenvalue Identity	60

1 Introduction

Positive Semi-Definite (PSD) matrices are central objects of interest in algorithm design, and continue to be studied extensively in optimization, spectral graph theory, numerical linear algebra, statistics, and dynamical systems, among many others [VB96, WSV12, GW95, ARV09, AHK05, Ste10, ST04, DL09, Wai19, DK19, SL⁺91]. Specifically, a real-valued matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is said to be PSD if it defines a non-negative quadratic form: namely if $x^\top \mathbf{A} x \geq 0$ for all $x \in \mathbb{R}^n$. If \mathbf{A} is symmetric, this is equivalent to the eigenvalues of \mathbf{A} being non-negative. Certifying whether a matrix is PSD often provides crucial insights into the structure of metric spaces [Sch35], arises as a separation oracles in Semi-Definite Programming (SDP) [VB96], leads to faster algorithms for solving linear systems and linear algebra problems [ST04, KOSZ13, MW17, BCW19] detects existence of community structure in random graphs [SKZ14], and is used to ascertain local convexity of functions. Furthermore, testing if a matrix is PSD is also used when studying the rate of dissipation in the heat equation [ÄzÖÖ93] and the behavior of non-oscillatory, exponentially stable modes of linear differential equations [Gle94]. For these applications, in addition to testing the existence of negative eigenvalues, it is often important to provide a *certificate* that the matrix is not PSD, by exhibiting a direction in which the quadratic form is negative.

While efficient, numerically stable algorithms for computing the spectrum of a matrix have been known since Turing [Tur48], such algorithms require reading the entire matrix and incur a cubic running time in practice. Computing the eigenvalues of a matrix is often the bottleneck in applications, especially when just determining the existence of negative eigenvalues suffices. For instance, checking embeddability of a finite metric into Euclidean space, feasibility of a SDP, convexity of a function, and if specialized solvers are applicable for linear algebraic problems, all only require knowledge of whether a given matrix is PSD. The focus of this work is to study when the property of being PSD can be tested *sublinear* time and queries, without reading the entire matrix.

We approach the problem from the perspective of property testing [GGR98, Gol17], where the input matrix \mathbf{A} is promised to be either a PSD matrix, or “ ϵ -far” from PSD under an appropriate notion of distance (discussed below). Specifically, we work in the *bounded-entry model*, proposed by Balcan, Li, Woodruff, and Zhang [BLWZ19], where the input matrix has bounded entries: $\|\mathbf{A}\|_\infty \leq 1$. Boundedness is often a natural assumption in practice, and has numerous real world applications, such as recommender systems as in the Netflix Challenge [KBV09], unweighted or bounded weight graphs [Gol10, GGR98], correlation matrices, distance matrices with bounded radius, and others [LWW14, KIDP16, BLWZ19]. Further, the boundedness of entries avoids degenerate instances where an arbitrarily large entry is hidden in \mathbf{A} , thereby drastically changing the spectrum of \mathbf{A} , while being impossible to test without reading the entire matrix.

Our starting point is a simple fact: a matrix \mathbf{A} is PSD if and only if all *principal*² submatrices of \mathbf{A} are PSD. However, a much more interesting direction is: if \mathbf{A} is not PSD, what can be said about the eigenvalues of the submatrices of \mathbf{A} ? Specifically, if \mathbf{A} is far from PSD, how large of a submatrix must one sample in order to find a negative eigenvalue? Note that given a principal submatrix $\mathbf{A}_{T \times T}$ with $x^\top \mathbf{A}_{T \times T} x < 0$ for some $x \in \mathbb{R}^{|T|}$, this direction x can be used as a certificate that the input matrix is not PSD, since $y^\top \mathbf{A} y = x^\top \mathbf{A}_{T \times T} x < 0$, where y is the result of padding x with 0’s. Further,

²Recall that a principal submatrix $\mathbf{A}_{T \times T}$ for $T \subseteq [n]$ is the restriction of \mathbf{A} to the rows and columns indexed by T .

it leads us to a natural algorithm to test definiteness: sample multiple principal submatrices and compute their eigenvalues. If any are negative, then \mathbf{A} must not be PSD. Determining the query complexity of this task is the principal focus of this paper. Specifically, we ask:

Can the positive semi-definiteness of a bounded matrix be tested via the semi-definiteness of a small random submatrix?

The Testing Models. The distance from \mathbf{A} to the PSD cone is given by $\min_{\mathbf{B} \geq 0} \|\mathbf{A} - \mathbf{B}\|$, where $\|\cdot\|$ is a norm, and $\mathbf{B} \geq 0$ denotes that \mathbf{B} is PSD. To instantiate $\|\cdot\|$, we consider two natural norms over $n \times n$ matrices: the spectral norm ($\|\cdot\|_2$) and the Euclidean norm ($\|\cdot\|_F$). Perhaps surprisingly, the distance of a symmetric matrix \mathbf{A} to the PSD cone under these norms can be characterized in terms of the eigenvalues of \mathbf{A} . In particular, let $\lambda \in \mathbb{R}^n$ be the vector of eigenvalues of \mathbf{A} . Then, the spectral norm distance corresponds to the ℓ_∞ distance between λ and the positive orthant. Similarly, the squared Frobenius distance corresponds to the ℓ_2^2 distance between λ and the positive orthant.

Therefore, we will refer to the two resulting gap problems as the ℓ_∞ -gap and the ℓ_2^2 -gap, respectively. This connection between matrix norms of \mathbf{A} and vector norms of eigenvalues λ will be highly useful for the analysis of random submatrices. Next, we formally define the testing problems:

Problem 1.1 (PSD Testing with Spectral norm/ ℓ_∞ -gap). Given $\epsilon \in (0, 1]$ and a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ such that $\|\mathbf{A}\|_\infty \leq 1$, distinguish whether \mathbf{A} satisfies:

- (1) \mathbf{A} is PSD.
- (2) \mathbf{A} is ϵ -far from the PSD cone in Spectral norm: $\min_{\mathbf{B} \geq 0} \|\mathbf{A} - \mathbf{B}\|_2 = \max_{i: \lambda_i < 0} |\lambda_i(\mathbf{A})| \geq \epsilon n$.

The fact that the spectral norm distance from \mathbf{A} to the PSD cone ($\min_{\mathbf{B} \geq 0} \|\mathbf{A} - \mathbf{B}\|_2$) is equivalent to the magnitude of the smallest negative eigenvalue of \mathbf{A} is a consequence of the variational principle for eigenvalues. For general non-symmetric matrices \mathbf{A} , one can replace (2) above with the condition $x^\top \mathbf{A} x < -\epsilon n$ for some unit vector $x \in \mathbb{R}^n$, which is equivalent to (2) if \mathbf{A} is symmetric (again by the variational principle). We note that our results for the ℓ_∞ -gap hold in this more general setting.³

Next, if we instantiate $\|\cdot\|$ with the (squared) Euclidean norm, we obtain the ℓ_2^2 gap problem.

Problem 1.2 (PSD Testing with ℓ_2^2 -gap). Given $\epsilon \in (0, 1]$ and a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ such that $\|\mathbf{A}\|_\infty \leq 1$, distinguish whether \mathbf{A} satisfies:

- (1) \mathbf{A} is PSD.
- (2) \mathbf{A} is ϵ -far from the PSD cone in squared Euclidean norm:

$$\min_{\mathbf{B} \geq 0} \|\mathbf{A} - \mathbf{B}\|_F^2 = \sum_{i: \lambda_i(\mathbf{A}) < 0} \lambda_i^2(\mathbf{A}) \geq \epsilon n^2 \quad (1)$$

Note that the identity $\min_{\mathbf{B} \geq 0} \|\mathbf{A} - \mathbf{B}\|_F^2 = \sum_{i: \lambda_i(\mathbf{A}) < 0} \lambda_i^2(\mathbf{A})$ in equation 1 also follows from the variational principle for eigenvalues (see Appendix A). Similarly to the ℓ_∞ -gap, if \mathbf{A} is not symmetric

³Also note that given query access to any $\mathbf{A} \in \mathbb{R}^{n \times n}$, one can always run a tester on the symmetrization $\mathbf{B} = (\mathbf{A} + \mathbf{A}^\top)/2$, which satisfies $x^\top \mathbf{A} x = x^\top \mathbf{B} x$ for all x , with at most a factor of 2 increase in query complexity.

one can always run a tester on the symmetrization $(\mathbf{A} + \mathbf{A}^\top)/2$. Also observe that $\|\mathbf{A}\|_F^2 \leq n^2$ and $\|\mathbf{A}\|_2 \leq n$ for bounded entries matrices, hence the respective scales of n, n^2 in the two gap instances above. Notice by definition, if a symmetric matrix \mathbf{A} is ϵ -far from PSD in ℓ_∞ then \mathbf{A} is ϵ^2 -far from PSD in ℓ_2^2 . However, the converse is clearly not true, and as we will see the complexity of PSD testing with ϵ^2 - ℓ_2^2 gap is strictly harder than testing with ϵ - ℓ_∞ gap.⁴

In fact, there are several important examples of matrices which are far from the PSD cone in ℓ_2^2 , but which are not far in ℓ_∞ . For instance, if \mathbf{A} is a random matrix with bounded moments, such as a matrix with i.i.d. Rademacher $(\{1, -1\})$ or Gaussian entries, then as a consequence of Wigner's Semicircle Law \mathbf{A} will be $\Omega(1)$ -far in ℓ_2^2 distance. However, $\|\mathbf{A}\|_2 = O(\sqrt{n})$ with high probability, so \mathbf{A} will only be $O(1/\sqrt{n})$ -far in ℓ_∞ distance. Intuitively, such random instances should be very "far" from being PSD, and the ℓ_2^2 distance captures this fact.

Remark 1.3. A previous version of this work defined the gap in Problem 1.1 in full generality (without symmetry assumed) as $x^\top \mathbf{A} x < -\epsilon n$ for a unit vector x . We have since changed the presentation, as this more general definition does not clearly emphasize the connection between the Problem 1.1 and the spectral norm. The authors would like to thank an anonymous reviewer for this suggestion. We note that the results themselves remain unaffected.

1.1 Our Contributions

We now introduce our main contributions. Our algorithms for PSD testing randomly sample principal submatrices and check if they are PSD. Thus, all our algorithms have one-sided error; when \mathbf{A} is PSD, they always return PSD, and whenever our algorithms return Not PSD, they output a certificate in the form of a principal submatrix which is not PSD. In what follows, $\omega < 2.373$ is the exponent of matrix multiplication, and $\tilde{O}, \tilde{\Omega}$ notation only hide $\log(1/\epsilon)$ factors (and $\log(s)$ factors for Ky-Fan-s and residual error bounds), thus our bounds have no direct dependency on the input size n . We first state our result for the ℓ_∞ gap problem in its most general form, which is equivalent to Problem 1.1 in the special case when \mathbf{A} is symmetric.

Theorem 3.25 (ℓ_∞ -gap Upper Bound) *There is a non-adaptive sampling algorithm which, given $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$ and $\epsilon \in (0, 1)$, returns PSD if $x^\top \mathbf{A} x \geq 0$ for all $x \in \mathbb{R}^n$, and with probability $2/3$ returns Not PSD if $x^\top \mathbf{A} x \leq -\epsilon n$ for some unit vector $x \in \mathbb{R}^n$. The algorithm make $\tilde{O}(1/\epsilon^2)$ queries to the entries of \mathbf{A} , and runs in time $\tilde{O}(1/\epsilon^\omega)$.*

We demonstrate that the algorithm of Theorem 3.25 is optimal up to $\log(1/\epsilon)$ factors, even for adaptive algorithms with two-sided error. Formally, we show:

Theorem 5.1 (ℓ_∞ -gap Lower Bound) *Any adaptive or non-adaptive algorithm which solves the PSD testing problem with ϵ - ℓ_∞ gap with probability at least $2/3$, even with two-sided error and if \mathbf{A} is promised to be symmetric, must query $\tilde{\Omega}(1/\epsilon^2)$ entries of \mathbf{A} .*

Next, we present our algorithm for the ℓ_2^2 -gap problem. Our algorithm crucially relies on first running our tester for the ℓ_∞ -gap problem, which allows us to demonstrate that if \mathbf{A} is far from PSD

⁴The difference in scaling of ϵ between the ℓ_∞ and ℓ_2^2 gap definitions (ϵ is squared in the latter) is chosen for the sake of convenience, as it will become clear the two problems are naturally studied in these respective parameterizations.

in ℓ_2^2 but close in ℓ_∞ , then it must be far, under other notions of distance such as Schatten norms or residual tail error, from any PSD matrix.

Theorem 4.12 (ℓ_2^2 -gap Upper Bound) *There is a non-adaptive sampling algorithm which, given a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$ and $\epsilon \in (0, 1)$, returns PSD if \mathbf{A} is PSD, and with probability $2/3$ returns Not PSD if $\min_{\mathbf{B} \geq 0} \|\mathbf{A} - \mathbf{B}\|_F^2 \geq \epsilon n^2$. The algorithm make $\tilde{O}(1/\epsilon^4)$ queries to \mathbf{A} , and runs in time $\tilde{O}(1/\epsilon^{2\omega})$.*

We complement our upper bound by a $\tilde{\Omega}(\frac{1}{\epsilon^2})$ lower bound for PSD-testing with ϵ - ℓ_2^2 gap, which holds even for algorithms with two sided error. Our lower bound demonstrates a separation between the complexity of PSD testing with $\sqrt{\epsilon}$ - ℓ_∞ gap and PSD testing with ϵ - ℓ_2^2 -gap, showing that the concentration of negative mass in large eigenvalues makes PSD testing a strictly easier problem.

Theorem 5.12 (ℓ_2^2 -gap Lower Bound) *Any non-adaptive algorithm which solves the PSD testing problem with ϵ - ℓ_2^2 gap with probability at least $2/3$, even with two-sided error, must query $\tilde{\Omega}(1/\epsilon^2)$ entries of \mathbf{A} .*

Our lower bound is built on discrete hard instances which are “locally indistinguishable”, in the sense that the distribution of any small set of samples is completely identical between the PSD and ϵ -far cases. At the heart of the lower bound is a key combinatorial Lemma about arrangements of paths on cycle graphs (see discussion in Section 1.4.3). Our construction is highly general, and we believe will likely be useful for proving other lower bounds for matrix and graph property testing problems. Exemplifying the applicability of our construction, we obtain as an immediate corollary a new lower bound for testing the Schatten-1 norm of \mathbf{A} . Recall, that the Schatten-1 norm is defined via $\|\mathbf{A}\|_{S_1} = \sum_i \sigma_i(\mathbf{A})$, where $\sigma_1(\mathbf{A}) \geq \dots \geq \sigma_n(\mathbf{A})$ are the singular values of \mathbf{A} .

Theorem 5.19 (Schatten-1 Lower Bound) *Fix any $1/\sqrt{n} \leq \epsilon \leq 1$. Then any non-adaptive algorithm in the bounded entry model that distinguishes between*

1. $\|\mathbf{A}\|_{S_1} \geq \epsilon n^{1.5}$,
2. $\|\mathbf{A}\|_{S_1} \leq (1 - \epsilon_0)\epsilon n^{1.5}$

with probability $2/3$, where $\epsilon_0 = 1/\log^{O(1)}(1/\epsilon)$, must make at least $\tilde{\Omega}(1/\epsilon^4)$ queries to \mathbf{A} .

Note that one always has $\|\mathbf{A}\|_{S_1} \leq n^{1.5}$ in the bounded entry model ($\|\mathbf{A}\|_\infty \leq 1$), which accounts for the above scaling. Theorem 5.19 extends a lower bound of Balcan et. al. [BLWZ19], which is $\Omega(n)$ for the special case of $\epsilon, \epsilon_0 = \Theta(1)$. Thus, for the range $\epsilon = \tilde{O}(n^{-1/4})$, our lower bound is an improvement. To the best of our knowledge, Theorem 5.19 gives the first $\tilde{\Omega}(n^2)$ sampling lower bound for testing Schatten-1 in a non-degenerate range (i.e., for $\|\mathbf{A}\|_{S_1} > n$).

Remark 1.4. We note that the lower bound of [BLWZ19] is stated for a slightly different version of gap (a “ ϵ - ℓ_0 ”-gap), where either $\|\mathbf{A}\|_{S_1} \geq c_1 n^{1.5}$ for a constant c_1 , or at least ϵn^2 of the entries of \mathbf{A} must be changed (while respecting $\|\mathbf{A}\|_\infty \leq 1$) so that the Schatten-1 is larger than $c_1 n^{1.5}$. However, their lower bound construction itself satisfies the “Schatten-gap” version as stated in Theorem 5.19, where either $\|\mathbf{A}\|_{S_1} \geq c_1 n^{1.5}$, or $\|\mathbf{A}\|_{S_1} \leq c_2 n^{1.5}$ and $c_1 > c_2$ are constants. From here, it is easy to see that this gap actually implies the ℓ_0 -gap (and this is used to obtain the ℓ_0 -gap lower bound in [BLWZ19]), since if $\|\mathbf{A}\|_{S_1} \leq c_2 n^{1.5}$ then for any \mathbf{E} with $\|\mathbf{E}\|_\infty \leq 2$ and $\|\mathbf{E}\|_0 \leq \epsilon n^2$ for a small enough

constant $\epsilon < c_2^2$, we have $\|\mathbf{A} + \mathbf{E}\|_{S_1} \leq \|\mathbf{A}\|_{S_1} + \|\mathbf{E}\|_{S_1} \leq n^{1.5}(c_2 + 2\sqrt{\epsilon}) < c_1 n^{1.5}$. So Theorem 5.19 implies a lower bound of $\tilde{\Omega}(1/\epsilon^2)$ for distinguishing $\|\mathbf{A}\|_{S_1} \geq \sqrt{\epsilon} n^{1.5}$ from the case of needing to change at least $\tilde{\Omega}(\epsilon n^2)$ entries of \mathbf{A} so that $\|\mathbf{A}\|_{S_1} \geq \sqrt{\epsilon} n^{1.5}$. Thus, our lower bound *also* extends the ℓ_0 -gap version of the results of [BLWZ19] for the range $\epsilon = \tilde{O}(1/\sqrt{n})$.

In addition to Schatten-1 testing, the same lower bound construction and techniques from Theorem 5.12 also result in new lower bounds for testing the Ky-Fan s norm $\|\mathbf{A}\|_{KF(s)} = \sum_{i=1}^s \sigma_i(\mathbf{A})$, as well as the cost of the best rank- s approximation $\|\mathbf{A} - \mathbf{A}_s\|_F^2 = \sum_{i>s} \sigma_i^2(\mathbf{A})$, stated below. In the following, s is any value $1 \leq s \leq n/(\text{poly log } n)$, and c is a fixed constant.

Theorem 5.20 (Ky-Fan Lower Bound) *Any non-adaptive algorithm in the bounded entry model which distinguishes between*

1. $\|\mathbf{A}\|_{KF(s)} > \frac{c}{\log s} n$
2. $\|\mathbf{A}\|_{KF(s)} < (1 - \epsilon_0) \cdot \frac{c}{\log s} n$

with probability 2/3, where $\epsilon_0 = \Theta(1/\log^2(s))$, must query at least $\tilde{\Omega}(s^2)$ entries of \mathbf{A} .

Theorem 5.21 (Residual Error Lower Bound) *Any non-adaptive algorithm in the bounded entry model which distinguishes between*

1. $\|\mathbf{A} - \mathbf{A}_s\|_F^2 > \frac{c}{s \log s} n$
2. $\|\mathbf{A} - \mathbf{A}_s\|_F^2 < (1 - \epsilon_0) \cdot \frac{c}{s \log s} n$

with probability 2/3, where $\epsilon_0 = 1/\log^{O(1)}(s)$, must query at least $\tilde{\Omega}(s^2)$ entries of \mathbf{A} .

Our lower bound for the Ky-Fan norm complements a Ky-Fan testing lower bound of [LW16b], which is $\Omega(n^2/s^2)$ for distinguishing **1)** $\|\mathbf{A}\|_{KF(s)} < 2.1s\sqrt{n}$ from **1)** $\|\mathbf{A}\|_{KF(s)} > 2.4s\sqrt{n}$ when $s = O(\sqrt{n})$. Note their bound decreases with s , whereas ours increases, thus the two bounds are incomparable (although they match up to $\log(s)$ factors at $s = \Theta(\sqrt{n})$).⁵ We also point out that there are (not quite matching) upper bounds for both the problems of Ky-Fan norm and s -residual error testing in the bounded entry model, just based on a standard application of the Matrix Bernstein Inequality.⁶ We leave the exact query complexity of these and related testing problems for functions of singular values in the bounded entry model as subject for future work.

A Remark on the ℓ_2^2 -Gap. We note that there appear to be several key barriers to improving the query complexity of PSD testing with ℓ_2^2 -gap beyond $O(1/\epsilon^4)$, which we briefly discuss here. First, in general, to preserve functions of the squared singular values of \mathbf{A} up to error ϵn^2 , such as $\|\mathbf{A}\|_F^2 = \sum_i \sigma_i^2(\mathbf{A})$ or $\|\mathbf{A}\|_2^2 = \sigma_1^2(\mathbf{A})$, any algorithm which samples a submatrix must make $\Omega(1/\epsilon^4)$ queries (see Lemma 5.22 for estimating $\sum_{i \leq k} \sigma_i^2$ for any k). In other words, detecting ϵn^2 -sized

⁵The bound from [LW16b] is stated in the sketching model, however the entries of the instance are bounded, thus it also applies to the sampling model considered here.

⁶See Theorem 6.1.1 of [Tro15], applied to $S_k = a_{(k)}(a_{(k)})^\top$, where $a_{(k)}$ is the k -th row sampled in \mathbf{A} ; for the case of residual error, one equivalently applies matrix Bernstein inequality to estimate the head $\sum_{i \leq k} \sigma_i^2(\mathbf{A})$. These bounds can be tightened via the usage of interior Chernoff bounds [GT11].

perturbations in the spectrum of a matrix in general requires $\Omega(1/\epsilon^4)$ sized submatrices. This rules out improving the query complexity by detecting the ϵn^2 negative mass in \mathbf{A} via, for instance, testing if the sum of squares of top $k = 1/\epsilon$ singular values has $\Theta(\epsilon n^2)$ less mass than it should if \mathbf{A} were PSD (even this may require $\Omega(k^2/\epsilon^4)$ queries, see the discussion in Section 1.4.2).

The key issue at play in the above barrier appears to be the requirement of sampling submatrices. Indeed, notice for the simplest case of $\|\mathbf{A}\|_F^2$, we can easily estimate $\|\mathbf{A}\|_F^2$ to additive ϵn^2 via $O(1/\epsilon^2)$ queries to random entries of \mathbf{A} . On the other hand, if these queries must form a submatrix, then it is easy to see that $\Omega(1/\epsilon^4)$ queries are necessary, simply from the problem of estimating $\|\mathbf{A}\|_F^2$ whose rows (or columns) have values determined by a coin flip with bias either equal to $1/2$ or $1/2 + \epsilon$. On the other hand, for testing positive semi-definiteness, especially with one-sided error, the requirement of sampling a principal submatrix seems unavoidable.

In addition, a typical approach when studying spectral properties of submatrices is to first pass to a random row submatrix $\mathbf{A}_{S \times [n]}$, argue that it preserves the desired property (up to scaling), and then iterate the process on a column submatrix $\mathbf{A}_{S \times T}$. Unfortunately, these types of arguments are not appropriate when dealing with eigenvalues of \mathbf{A} , since after passing to the rectangular matrix $\mathbf{A}_{S \times [n]}$, any notion of negativity of the eigenvalues has now been lost. This forces one to argue indirectly about functions of the singular values of $\mathbf{A}_{S \times [n]}$, returning to the original difficulty described above. We leave it as an open problem to determine the exact non-adaptive query complexity of PSD testing with ℓ_2^2 -gap. For a further discussion of these barriers and open problems, see Section 6.

1.2 Connections to Optimization, Euclidean Metrics and Linear Algebra

We now describe some explicit instances where our algorithms may be useful for testing positive semi-definiteness. We emphasize that in general, the distance between \mathbf{A} and the PSD cone may be too small to verify via our testers. However, when the input matrices satisfy a non-trivial gap from the PSD cone, we can speed up some basic algorithmic primitives. The first is testing feasibility of the PSD constraint in a Semi-Definite Program (SDP) with sublinear queries and time, so long as the variable matrix has bounded entries. Importantly, our algorithms also output a separating hyperplane to the PSD cone.

Corollary 1.5 (Feasibility and Separating Hyperplanes for SDPs). *Given a SDP \mathcal{S} , let $\mathbf{X} \in \mathbb{R}^{n \times n}$ be a symmetric matrix that violates the PSD constraint for \mathcal{S} . Further, suppose $\|\mathbf{X}\|_\infty \leq 1$ and \mathbf{X} is ϵn^2 -far in entry-wise ℓ_2^2 distance to the PSD cone. Then, there exists an algorithm that queries $\tilde{O}(1/\epsilon^4)$ entries in \mathbf{X} and runs in $\tilde{O}(1/\epsilon^{2\omega})$ time, and with probability $9/10$, outputs a vector \tilde{v} such that $\tilde{v}^T \mathbf{X} \tilde{v} < 0$. Moreover, if $\lambda_{\min}(\mathbf{X}) < -\epsilon n$, then there is an algorithm yielding the same guarantee, that queries $\tilde{O}(1/\epsilon^2)$ entries in \mathbf{X} and runs in $\tilde{O}(1/\epsilon^\omega)$ time.*

While in the worst case, our algorithm may need to read the whole matrix to exactly test if \mathbf{X} is PSD, there may be applications where relaxing the PSD constraint to the convex set of matrices which are close to the PSD cone in Euclidean distance is acceptable. Moreover, our algorithm may be run as a preliminary step at each iteration of an SDP solver to check if the PSD constraint is badly violated, resulting in speed-ups by avoiding an expensive eigendecomposition of \mathbf{X} whenever our algorithm outputs a separating hyperplane [VB96].

Next, we consider the problem of testing whether an arbitrary finite metric d over n points, $x_1, \dots, x_n \in \mathbb{R}^d$ is embeddable into Euclidean Space. Testing if a metric is Euclidean has a myriad of applications, such as determining whether dimensionality reduction techniques such as Johnson-Lindenstrauss can be used [PR03], checking if efficient Euclidean TSP solvers can be applied [Aro98], and more recently, computing a low-rank approximation in sublinear time [BW18, IVWW19]. It is well known (Schoenberg's criterion [Dat10]) that given a distance matrix $\mathbf{D} \in \mathbb{R}^{n \times n}$ such that $\mathbf{D}_{i,j} = d(x_i, x_j)$, the points are isometrically embeddable into Euclidean space if and only if $\mathbf{G} = \mathbf{1} \cdot \mathbf{D}_{1,*} + \mathbf{D}_{1,*}^\top \cdot \mathbf{1}^\top - \mathbf{D} \geq 0$, where $\mathbf{D}_{1,*}$ is the first row of \mathbf{D} . Notice that embeddability is scale invariant, allowing one to scale distances to ensure boundedness. Furthermore, since our algorithms sample submatrices and check for non-positive semi-definiteness, the tester need not know this scaling in advance, and gives guarantees for distinguishing definiteness if the necessary gap is satisfied after hypothetically scaling the entries.

Corollary 1.6 (Euclidean Embeddability of Finite Metrics). *Given a finite metric d on n points $\{x_1, \dots, x_n\}$, let $\mathbf{D} \in \mathbb{R}^{n \times n}$ be the corresponding distance matrix, scaled so that $\|\mathbf{D}\|_\infty \leq 1/3$, and let $\mathbf{G} = \mathbf{1} \mathbf{D}_{1,*} + \mathbf{D}_{1,*}^\top \mathbf{1}^\top - \mathbf{D}$. Then if $\min_{\mathbf{B} \geq 0} \|\mathbf{G} - \mathbf{B}\|_F^2 \geq \epsilon n^2$, there exists an algorithm that queries $\tilde{O}(1/\epsilon^4)$ entries in \mathbf{A} and with probability $9/10$, determines the non-embeddability of $\{x_1, \dots, x_n\}$ into Euclidean space. Further, the algorithm runs in time $\tilde{O}(1/\epsilon^{2\omega})$.*

Remark 1.7. An intriguing question is to characterize geometric properties of finite metrics based on the ℓ_2^2 -distance of the Schoenberg matrix \mathbf{G} from the PSD cone. For instance, given a finite metric with Schoenberg matrix \mathbf{G} that is close to being PSD in ℓ_2^2 -distance, can we conclude that the metric has a low worst or average case distortion embedding into Euclidean space?

Remark 1.8. Since rescaling entries to be bounded only affects the gap parameter ϵ , in both of the above cases, so long as the magnitude of the entries in \mathbf{X}, \mathbf{D} do not scale with n , the running time of our algorithms is still sublinear in the input.

Finally, several recent works have focused on obtaining sublinear time algorithms for low-rank approximation when the input matrix is PSD [MW17, BCW19]. However, such algorithms only succeed when the input is PSD or close to PSD (in ℓ_2^2), and it is unknown how to verify whether these algorithms succeeded in sublinear time. Therefore, our tester can be used as a pre-processing step to determine input instances where the aforementioned algorithms provably will (or will not) succeed.

1.3 Related work

Property testing in the bounded entry model was first considered in [BLWZ19] to study the query complexity of testing spectral properties of matrices, such as stable rank (the value $\|\mathbf{A}\|_F^2 / \|\mathbf{A}\|_2^2$) and Schatten p norms. A related model, known as the *bounded row model*, where rows instead of entries are required to be bounded, was studied by Li, Wang, and Woodruff [LWW14], who gave tight bounds for testing stable rank in this model. In addition, the problem of testing the rank of a matrix from a small number of queries has been well studied [PR03, KS03, LWW14, BBG18], as well the problem of estimating the rank via a random submatrix [BH11, BZ16]. Notice that since rank is not a smooth spectral property, hiding an unbounded value in a single entry of \mathbf{A} cannot drastically alter the rank. Thus, for testing rank, the condition of boundedness is not required.

More generally, the bounded entry model is the natural sampling analogue for the *linear sketching model*, where the algorithm gets to choose a matrix $\mathbf{S} \in \mathbb{R}^{t \times n^2}$, where t is the number of “queries”, and then observes the product $\mathbf{S} \cdot \text{vec}(\mathbf{A})$, where $\text{vec}(\mathbf{A})$ is the vectorization of \mathbf{A} [LW17, LW16b, BCK⁺18, LW16a, LWW14, LNW14, BKKS19, LNW19]. The model has important applications to streaming and distributed algorithms. Understanding the query complexity of sketching problems, such as estimating spectral norms and the top singular values [AN13, LNW14, LW16b], estimating Schatten and Ky-Fan norms [LW16b, LW17, LW16a, BKKS19], estimating ℓ_p norms [AMS96, Ind06, KNW10, JW19, BJWY20], and ℓ_p sampling [MW10, JST11, JW18, JSTW19], has been a topic of intense study. For the problem of sketching *eigenvalues* (with their signs), perhaps the most related result is [AN13], which gives point-wise estimates of the top eigenvalues. Notice that linear sketching can simulate sampling by setting the rows of \mathbf{S} to be standard basis vectors, however sketching is in general a much stronger query model. Note that to apply a linear sketch, unlike in sampling, one must read all the entries of \mathbf{A} , which does not yield sublinear algorithms.

A special case of the sketching model is the *matrix-vector product* model, which has been studied extensively in the context of compressed sensing [CRT06, EK12] and sparse recovery [GI10]. Here, one chooses vectors v_1, \dots, v_k and observes the products $\mathbf{A}v_1, \dots, \mathbf{A}v_k$. Like sketching, matrix-vector product queries are a much stronger access model than sampling. Recently, in the matrix-vector product model, Sun et. al. considered testing various graph and matrix properties [SWYZ19], and Han et. al. considered approximating spectral sums and testing positive semi-definiteness [HMAS17].

Lastly, while there has been considerable work on understanding concentration of norms and singular values of random matrices, not as much is known about their *eigenvalues*. Progress in understanding the behavior of singular values of random matrices includes concentration bounds for spectral norms of submatrices [RV07, Tro08], concentration bounds for extreme singular values [GT11, Tro15, Ver10, GLSS18, KS18], non-commutative Khintchine inequalities for Schatten- p norms [LPP91, Pis09, PR17], as well as Kadison-Singer type discrepancy bounds [MSS15, KLS20, SZ20]. These random matrix concentration bounds have resulted in improved algorithms for many fundamental problems, such as low-rank approximation and regression [CW17, MW17, BW18, IVWW19, DJS⁺19] and spectral sparsification [ST11, BSST13, ACK⁺16]. However, in general, understanding behavior of negative eigenvalues of random matrices and submatrices remains largely an open problem.

1.4 Technical Overview

In this section, we describe the techniques used in our non-adaptive testing algorithms for the ℓ_∞ and more general ℓ_2^2 gap problem, as well as the techniques involved in our lower bound construction for the ℓ_2^2 -gap.

1.4.1 PSD Testing with ℓ_∞ Gap

Recall in the general statement of the ℓ_∞ -gap problem, our task is to distinguish between $\mathbf{A} \in \mathbb{R}^{n \times n}$ satisfying $x^\top \mathbf{A} x \geq 0$ for all $x \in \mathbb{R}^n$, or $x^\top \mathbf{A} x \leq -\epsilon n$ for some unit vector $x \in \mathbb{R}^n$. Since if $x^\top \mathbf{A} x \geq 0$ for all $x \in \mathbb{R}^n$ the same holds true for all principal submatrices of \mathbf{A} , it suffices to show that in the

ϵ -far case we can find a $k \times k$ principal submatrix $\mathbf{A}_{T \times T}$ such that $y^\top \mathbf{A}_{T \times T} y < 0$ for some $y \in \mathbb{R}^k$.⁷

Warmup: A $O(1/\epsilon^3)$ query algorithm. Since we know $x^\top \mathbf{A} x \leq -\epsilon n$ for some fixed x , one natural approach would be to show that the quadratic form with the *same* vector x , projected onto to a random subset $T \subset [n]$ of its coordinates, is still negative. Specifically, we would like to show that the quadratic form $Q_T(x) = x_T^\top \mathbf{A}_{T \times T} x_T$, of x with a random principal submatrix $\mathbf{A}_{T \times T}$ for $T \subset [n]$ will continue to be negative. If $Q_T(x) < 0$, then clearly $\mathbf{A}_{T \times T}$ is not PSD. Now while our algorithm does not know the target vector x , we can still analyze the concentration of the scalar random variable $Q_T(x)$ over the choice of T , and show that it is negative with good probability.

Proposition 3.7 and Lemma 3.8 (informal) *Suppose $\mathbf{A} \in \mathbb{R}^{n \times n}$ satisfies $\|\mathbf{A}\|_\infty \leq 1$ and $x^\top \mathbf{A} x \leq -\epsilon n$ where $\|x\|_2 \leq 1$. Then if $k \geq 6/\epsilon$, and if $T \subset [n]$ is a random sample of expected size k , we have $\mathbb{E}[Q_T(x)] \leq -\frac{\epsilon k^2}{4n}$ and $\text{Var}(Q_T(x)) \leq O(\frac{k^3}{n^2})$.*

By the above Proposition, after setting $k = \Theta(1/\epsilon^2)$, we have that $|\mathbb{E}[Q_T(x)]|^2 = \Omega(\text{Var}(Q_T(x)))$, and so by Chebyshev's inequality, with constant probability we will have $Q_T(x) < 0$. This results in a $k^2 = O(1/\epsilon^4)$ query tester. To improve the complexity, we could instead set $k = \Theta(1/\epsilon)$ and re-sample T for k times independently to reduce the variance. Namely, one can sample submatrices T_1, T_2, \dots, T_k , and analyze $\frac{1}{k} \sum_{i=1}^k Q_{T_i}(x)$. The variance of this sum goes down to $O(\frac{k^2}{n^2})$, so, again by Chebyshev's inequality, the average of these quadratic forms will be negative with constant probability. If this occurs, then at least one of the quadratic forms must be negative, from which we can conclude that at least one of $\mathbf{A}_{T_i \times T_i}$ will fail to be PSD, now using only $O(1/\epsilon^3)$ queries.

A Family of Hard Instances One could now hope for an even tighter analysis of the concentration of $Q_T(x)$, so that $O(1/\epsilon^2)$ total queries would be sufficient. Unfortunately, the situation is not so simple, and in fact the two aforementioned testers are tight in the query complexity for the matrix dimensions they sample. Consider the hard instance \mathbf{A} in the left of Figure 1, which is equal to the identity on the diagonal, and is zero elsewhere except for a small subset $S \subset [n]$ of $|S| = \epsilon^2 n$ rows and columns, where we have $\mathbf{A}_{S \times \bar{S}} = \mathbf{A}_{\bar{S} \times S} = -\mathbf{1}$, where \bar{S} is the complement of S . Notice that if we set $x_i^2 = 1/(2n)$ for $i \notin S$ and $x_i^2 = 1/(2\epsilon^2 n)$ for $i \in S$, then $x^\top \mathbf{A} x \leq -\epsilon n/4$. However, in order to even see a single entry from S , one must sample from at least $\Omega(1/\epsilon^2)$ rows or columns. In fact, this instance itself gives rise to a $\Omega(1/\epsilon^2)$ lower bound for any testing algorithm, even for adaptive algorithms (Theorem 5.1).

The difficulty of the above instance is that the negative mass of $x^\top \mathbf{A} x$ is hidden in only a ϵ^2 -fraction of \mathbf{A} . On the other hand, since the negative entries are so large and concentrated, one need only sample $O(1)$ entries from a single row $i \in S$ in order for $\mathbf{A}_{T \times T}$ to be non-PSD in the prior example. Thus, an algorithm for such instances would be to sample $O(1/\epsilon^2)$ principal submatrices, each of *constant* size. On the other hand, the set S could also be more spread out; namely, we could have $|S| = \alpha n$ for any $\epsilon^2 \leq \alpha \leq \epsilon$, but where each entry in $\mathbf{A}_{S \times \bar{S}}$ is set to $-\epsilon/\sqrt{\alpha}$ (see the matrix in the right side of Figure 1). If instead, we define $x_i^2 = 1/(2\alpha n)$ for $i \in S$, we still have $x^\top \mathbf{A} x < -\epsilon n/4$. However, now any submatrix $\mathbf{A}_{T \times T}$ with $|T \cap S| = 1$ must have at least $|T| \geq \alpha/\epsilon^2$ rows and columns, otherwise $\mathbf{A}_{T \times T}$ would be PSD due to the identity on the diagonal.

⁷This can be efficiently checked by computing the eigenvalues of $\mathbf{A}_{T \times T} + \mathbf{A}_{T \times T}^\top$.




Figure 1: Hard instances for ℓ_∞ testing. On the left, the negative mass is highly concentrated in $|S| = \epsilon^2 n$ rows and columns, and on the right it more spread out over $|S| = \alpha n$, where $\epsilon^2 \leq \alpha \leq \epsilon$.

The aforementioned instances suggest the following approach: query matrices at $O(\log \frac{1}{\epsilon})$ different scales of subsampling. Specifically, for each $\epsilon^2 \leq \alpha = 2^i \leq \epsilon$, we sample $\tilde{O}(\frac{\epsilon^2}{\alpha^2})$ independent $k \times k$ submatrices, each of size $k = \tilde{O}(\alpha/\epsilon^2)$, giving a total complexity of $\tilde{O}(\frac{1}{\epsilon^2})$. The analysis now proceeds by a complete characterization of the ways in which $x^\top \mathbf{A}x$ can be negative. Specifically, we prove the following: either a substantial fraction of the negative mass is hidden inside of a small set of rows and columns S with $|S| < \epsilon n$, or it is the case that $\text{Var}(\mathbf{Q}_T(x))$ is small enough so that a single $k \times k$ submatrix will already be non-PSD with good probability when $k \gtrsim 1/\epsilon$. Given this classification, it suffices to demonstrate a level of subsampling which will find a non-PSD submatrix when the negative mass is concentrated inside inside a small set S .

Eigenvector Switching. To analyze this case, ideally, one would like to demonstrate that conditioned on T intersecting S at some level of subsampling, we will have $\mathbf{Q}_T(x) < 0$ with good probability. Unfortunately, the approach of analyzing the quadratic form with respects to x will no longer be possible; in fact, $\mathbf{Q}_T(x)$ may never be negative conditioned on $|T \cap S| = 1$ (unless $|T| > 1/\epsilon$, which we cannot afford in this case). The complication arises from the fact that the coordinates of x_i in the small set S can be extremely large, and thus the diagonal contribution of $x_i^2 \mathbf{A}_{i,i}$ will dominate the quadratic form of a small submatrix. For instance, if $\mathbf{A}_{T \times T}$ is a sample with $k = |T| = O(1)$ which intersects the set S in the leftmost matrix in Figure 1, where $x_i = 1/(\epsilon\sqrt{n})$ for $i \in S$ and $x_i = 1/\sqrt{n}$ otherwise, then $\mathbf{Q}_T(x) \approx k/n - (k/\sqrt{n})x_i + \mathbf{A}_{i,i}x_i^2$, which is dominated by the diagonal term $\mathbf{A}_{i,i}x_i^2 = 1/(\epsilon^2 n)$. Thus, while $\mathbf{A}_{T \times T}$ itself is not PSD, we have that $\mathbf{Q}_T(x) > 0$.

To handle this, we must and analyze the quadratic form $\mathbf{Q}_T(\cdot)$ with respect to *another* direction y . The vector y may not even satisfy $y^\top \mathbf{A}y < 0$, however conditioned on $|T \cap S| \geq 1$, we will have $\mathbf{Q}_T(y) < 0$ with good probability. Clearly, we must scale down the large coordinates x_i for $i \in S$. However, one cannot scale too low, otherwise the negative contribution of the rows $i \in S$ would become too small. The correct scaling is then a careful balancing act between the contributions of the different portions of $\mathbf{A}_{T \times T}$. Informally, since the x_i 's for $i \in S$ make up a $|S|/n$ fraction of all coordinates, they can be as large as $x_i^2 \geq (n/|S|) \cdot (1/n)$. However, inside of the smaller submatrix $\mathbf{A}_{T \times T}$, then conditioned on $i \in T$, since $|T|$ is small x_i now makes up a larger $1/|T|$ fraction of the

submatrix, thus we should scale down x_i to only be $x_i^2 \approx |T|/n$. With this scaling in mind, we (very roughly) set $y_i^2 = (|S|/n) \cdot (|T|/n)$ if $i \in S$, and set $y_i = x_i$ otherwise. The remaining argument then requires a careful analysis of the contribution of entries of \mathbf{A} outside of S to show that the target vector y indeed satisfies $\mathbf{Q}_T(y) < 0$ with good probability conditioned on T intersecting S .

1.4.2 PSD Testing with ℓ_2 Gap

Recall in the ℓ_2 gap problem, our task is to distinguish between \mathbf{A} being PSD, and \mathbf{A} being ϵ -far in ℓ_2^2 distance from any PSD matrix, namely that $\sum_{i: \lambda_i(\mathbf{A}) < 0} \lambda_i^2(\mathbf{A}) > \epsilon n^2$. In what follows, we refer to the quantity $\sum_{i: \lambda_i(\mathbf{A}) < 0} \lambda_i^2(\mathbf{A})$ as the *negative mass* of \mathbf{A} . First observe that in the special case that we had a “large” negative eigenvalue, say $\lambda_{\min}(\mathbf{A}) = -\sqrt{\epsilon}n$, then by applying our testing algorithm for ℓ_∞ -gap, we could find a non-PSD submatrix with only $\tilde{O}(1/\epsilon)$ queries. However, in general the negative mass of \mathbf{A} may be spread out over many smaller eigenvalues. Thus, we cannot hope to apply our earlier approach for the ℓ_∞ -gap, which preserved the quadratic form $\mathbf{Q}_T(x) = x_T^\top \mathbf{A}_{T \times T} x_T$ with respects to a fixed direction x . Instead, our approach will be to show that if \mathbf{A} is ϵ -far from PSD in ℓ_2^2 , then the singular values of \mathbf{A} must be “far” from PSD, in some other notion of distance, allowing us to indirectly infer the existence of negative eigenvalues in submatrices.

PSD matrices are top-heavy, and a reduction to estimating the tail. Our first step is to show that if $\mathbf{A} \in \mathbb{R}^{n \times n}$ is PSD, then the t -“tail” of \mathbf{A} , defined as $\sum_{i>t} \sigma_i^2(\mathbf{A})$, cannot be too large. This can be derived from the following fact: if \mathbf{A} is PSD then we can bound the Schatten-1 norm of \mathbf{A} by $\|\mathbf{A}\|_{S_1} = \sum_i \sigma_i(\mathbf{A}) = \text{Tr}(\mathbf{A})$, which is at most n if $\|\mathbf{A}\|_\infty \leq 1$. This simple fact will prove highly useful, since whenever we can demonstrate that the Schatten-1 norm of a submatrix $\mathbf{A}_{T \times T}$ is larger than $|T|$, we may immediately conclude that $\mathbf{A}_{T \times T}$ is not PSD. In addition, it implies:

Proposition 4.2 (PSD matrices are top-heavy) *Fix any $n \in \mathbb{N}$, $1 \leq t \leq n$, and $\mathbf{D} \in \mathbb{R}^{n \times n}$. Then if \mathbf{D} is PSD, we have*

$$\sum_{i>t} \sigma_i(\mathbf{D})^2 \leq \frac{1}{t} (\text{Tr}(\mathbf{D}))^2$$

In particular, if \mathbf{D} has bounded entries $\|\mathbf{D}\|_\infty \leq 1$, we have $\sum_{i>t} \sigma_i(\mathbf{D})^2 \leq \frac{1}{t} n^2$.

On the other hand, suppose that \mathbf{A} is ϵ -far from PSD, and let $t > 10/\epsilon$. Then if no eigenvalue is smaller than $-\epsilon n/100$, a condition which can be checked with $\tilde{O}(1/\epsilon^2)$ queries by first running our ℓ_∞ -gap tester, then the negative mass must be spread out, and it must be the case that a substantial fraction of the negative mass of \mathbf{A} is contained in the bottom $n - t$ singular values. Specifically, we must have $\sum_{i>t} \sigma_i(\mathbf{A})^2 > (\epsilon/2)n^2$, whereas any PSD matrix \mathbf{D} would have to satisfy $\sum_{i>t} \sigma_i^2(\mathbf{D}) \leq (\epsilon/10)n^2$ by the above Proposition. Thus, after first running our ℓ_∞ tester, it will suffice to estimate the tail $\sum_{i>t} \sigma_i^2(\mathbf{A})$. Equivalently, since $\|\mathbf{A}\|_F^2 = \sum_i \sigma_i^2(\mathbf{A})$ can be efficiently estimated, it also suffices to estimate the “head” $\sum_{i \leq t} \sigma_i^2(\mathbf{A})$ to additive $O(\epsilon n^2)$.

In order to accomplish this, one could utilize the tools from random matrix concentration, such as Matrix Bernstein’s inequality [Tro15], which allows one to estimate each σ_i^2 to error ηn^2 by taking a random rectangular $O(1/\eta^2) \times O(1/\eta^2)$ sized submatrix. The error in estimating $\sum_{i \leq t} \sigma_i^2(\mathbf{A})$ is then $t \eta n^2$, thus one needs to set $\eta = O(\epsilon/t)$, giving a $O(1/\epsilon^8)$ tester with two-sided error. Using a careful bucketing analysis on the error, along with the more powerful Interior Matrix Chernoff bounds of Gittens and Tropp [GT11], one can improve this to $O(t^2/\epsilon^4) = O(1/\epsilon^6)$. However, substantial

improvements on unconditional estimation of $\sum_{i \leq t} \sigma_i^2(\mathbf{A})$ seem unlikely. In fact, we demonstrate that even for $t = 1$ (spectral norm estimation), tools such as matrix concentration inequalities which sample submatrices of \mathbf{A} , must make $\Omega(1/\epsilon^4)$ queries (Lemma 5.22), which rules out, for instance, a $o(t^2/\epsilon^4)$ upper bound for general t . Thus, instead of unconditional estimation, our main insight is to demonstrate conditions under which $\sum_{i \leq t} \sigma_i^2(\mathbf{A})$ can be efficiently estimated. When these conditions do not hold, we show that it is because the Schatten-1 norm of our sampled submatrix must be too large, from which we can deduce the existence of negative eigenvalues in our query.

In the first case, if the t -th singular value is not too large, say $\sigma_{t+1}(\mathbf{A}) \leq 10n/t$, we show that the (re-scaled) tail $\frac{n^2}{k^2} \sum_{i > t} \sigma_i^2(\mathbf{A}_{S \times T})$ of a random rectangular matrix, where $|S| = |T| = k = O(1/\epsilon^2)$, approximates the tail of \mathbf{A} to error $O(\epsilon n^2)$. Our argument relies on splitting \mathbf{A} into head and tail pieces $\mathbf{A} = \mathbf{A}_t + \mathbf{A}_{-t}$, where \mathbf{A}_t is \mathbf{A} projected onto the top- t eigenvectors of \mathbf{A} . We demonstrate that the spectral mass of each is preserved after passing to a random row submatrix, and additionally demonstrate that $\sigma_{\max}(\mathbf{A}_{-t}) = \sigma_{t+1}(\mathbf{A})$ does not grow too much using spectral decay inequalities for random submatrices [RV07]. This forces the spectrum of $(\mathbf{A}_{-t})_{S \times [n]}$ to be well spread out, allowing us to apply interlacing inequalities to demonstrate that after adding $(\mathbf{A}_t)_{S \times [n]}$ back in, the resulting tail is still sufficiently large, and then iterate the argument when sampling columns to obtain $\mathbf{A}_{S \times T}$.

On the other hand, if $\sigma_{t+1}(\mathbf{A})$ is too large, then after moving to a random row submatrix the spectral norm of \mathbf{A}_{-t} can concentrate highly in its top eigenvalues, which can then be absorbed by the top t eigenvalues of \mathbf{A}_t , stealing too much mass from the tail. Instead, note that if $\sigma_{t+1}(\mathbf{A}) \geq 10n/t$, then the Schatten norm of \mathbf{A} must be large, namely $\sum_i \sigma_i(\mathbf{A}) > 10n$, which cannot occur if \mathbf{A} is PSD. We show that by applying Interior Eigenvalue Matrix Chernoff bounds (mentioned above), we can preserve this fact, obtaining $\frac{n}{k} \sigma_{t+1}(\mathbf{A}_{S \times T}) > 10n/t$ with good probability when $k = \Omega(1/\epsilon^2)$. If this is the case, then the Schatten norm of the submatrix will be too large: $\|\mathbf{A}_{S \times T}\|_{S_1} \geq t(10k/t) > 10k$. To obtain a certificate from this fact, we move to the larger principal submatrix $\mathbf{A}_{(S \cup T) \times (S \cup T)}$, which we show must still have large Schatten norm, from which we can infer the existence of negative eigenvalues. Similarly, in the earlier case, we show that the large tail of $\mathbf{A}_{S \times T}$ implies that the principal submatrix $\mathbf{A}_{(S \cup T) \times (S \cup T)}$ also has too large of a tail, meaning it must not be PSD.

1.4.3 Lower Bounds

As seen above, the distribution of negative mass in the matrix \mathbf{A} plays an important role in the complexity of testing if \mathbf{A} is PSD. Specifically, the problem becomes easier the more concentrated the negative mass is within a few eigenvalues. So in order to avoid a $o(1/\epsilon^2)$ upper bound from the ℓ_∞ -testing algorithm, our hard instance must have $|\lambda_{\min}(\mathbf{A})| = O(\epsilon n)$ in the ϵ -far case. On the other hand, we cannot allow the negative mass to be extremely spread out, otherwise we would have to add many more *positive* eigenvalues to avoid violating the trace constraint $|\text{Tr}(\mathbf{A})| = |\sum_i \lambda_i(\mathbf{A})| \leq n$ implied by the boundedness, creating further spectral differences between the instances. With this in mind, our hard distribution will have $1/\epsilon$ negative eigenvalues, each roughly equal to $\lambda_i(\mathbf{A}) = -\epsilon n$.

The Hard Instance. Our first insight is to construct a discrete instance, with the property that the distribution induced by observing a small sample of the “meaningful” entries of \mathbf{A} is *identical* in both cases. Specifically, we construct two distributions: \mathcal{D}_{YES} and \mathcal{D}_{NO} over $n \times n$ matrices. In both cases, \mathbf{A} will be block diagonal, with k disjoint blocks $B_1, B_2, \dots, B_k \subset [n]$, each of size $|B_i| = n/k$,

for some parameter k ; we will later set $k = \Theta(1/\epsilon)$, so our target lower bound is $\Omega(k^2)$. In \mathcal{D}_{YES} , each block $\mathbf{A}_{B_i \times B_i}$ will be PSD, whereas in \mathcal{D}_{NO} we will have $\lambda_{\min}(\mathbf{A}_{B_i \times B_i}) = -\tilde{\Theta}(n/k) \approx -\epsilon n$. The partition $B_1 \cup B_2 \cup \dots \cup B_k = [n]$ is chosen randomly, so that for any fixed set of samples, only a small fraction them will be contained inside any block $\mathbf{A}_{B_i \times B_i}$. The diagonal entries will always be fixed to 1, and all off-diagonal entries are either $\{0, 1, -1\}$. The samples $a_1, a_2, \dots, a_s \in [n] \times [n]$ of any algorithm can then be interpreted as a graph H (possibly with self-loops), where for each edge $a_r = (i, j) \in E(H)$, the algorithm learns the value $\mathbf{A}_{i,j} \in \{0, 1, -1\}$.

Now consider the algorithm which just samples a $t \times t$ principal submatrix $T \subset [n]$, so that H is a t -clique. Now in expectation $\mathbb{E}[|T \cap B_i|] = \frac{t}{k}$ for each i , however, by a balls and bins argument, as t approaches k we will obtain some blocks i with $|T \cap B_i| = \Omega(\log k / \log \log k)$. Thus, to fool this query, we must be able to “fool” cliques of size roughly $\log k$ within a block B_i . On the other hand, an algorithm could find many more entries in a block by lop-sided sampling: for instance, it could sample k^2 entries in a single column of \mathbf{A} (H is a k^2 -star), getting k entries inside a column of a block $\mathbf{A}_{B_i \times B_i}$. Thus we must also fool large star queries. It turns out that the right property to consider is the *matching number* $\nu(H)$ of the query graph H , i.e. the size of a maximum matching. Notice for a star H , we have $\nu(H) = 1$. We prove (roughly) that if within each block B_i , one can “fool” every query graph H inside $\mathbf{A}_{B_i \times B_i}$ with matching number $\nu(H) < \ell$, one would obtain a lower bound of $\Omega(k^{\frac{2(\ell-1)}{\ell}})$. Thus, it will suffice to fool all query graphs H within a block B_i with $\nu(H) \leq \log k$.

For a first step towards this, suppose that in \mathcal{D}_{YES} , we set each block independently to $\mathbf{A}_{B_i \times B_i} = vv^\top$, where $v \in \{1, -1\}^{|B_i|}$ is a random sign vector, and in \mathcal{D}_{NO} , we set $\mathbf{A}_{B_i \times B_i} = -vv^\top$ (except we fix the diagonal to be 1 in both cases). Now notice that the distribution of any individual entry $(\mathbf{A}_{B_i \times B_i})_{a,b}$ is symmetric, and identical in both \mathcal{D}_{YES} and \mathcal{D}_{NO} . Furthermore, it is not difficult to check that the distribution of a path or star query H inside of $\mathbf{A}_{B_i \times B_i}$ is also identical in both cases. On the other hand, if H contained a *triangle*, then this would not be the case, since in \mathcal{D}_{YES} one could never have a negative cycle (x, y, z) where $v_x v_y = v_y v_z = v_z v_x = -1$, whereas this could occur in \mathcal{D}_{NO} , since we could have that $-v_x v_y = -v_y v_z = -v_z v_x = -1$. Thus, roughly, to distinguish between these distributions \mathcal{D}_{YES} from \mathcal{D}_{NO} , an algorithm must sample a triangle within one of the blocks $\mathbf{A}_{B_i \times B_i}$, which one can show requires $\Omega(k^{4/3})$ queries, yielding a first lower bound.⁸

Boosting to $\Omega(k^2)$. Given the above example, we would now like to construct instances which fool H with larger and larger $\nu(H)$. In fact, our next insight is to have an even simpler structure on \mathcal{D}_{YES} and \mathcal{D}_{NO} : each of them will be a random permutation of one of two *fixed* matrices $\mathbf{D}_1, \mathbf{D}_2$ respectively. We now formalize the “fooling” condition we need. For a matrix \mathbf{B} and a query graph H , let $(\mathbf{B})_H$ denote the result of setting all entries of \mathbf{B} not in H equal to zero. Then the matrices $\mathbf{D}_1, \mathbf{D}_2$ must have the property that for any graph H with $\nu(H) \leq \log k$, if $\sigma : [m] \rightarrow [m]$ is a random permutation and $\mathbf{P}_\sigma \in \mathbb{R}^{m \times m}$ is the row permutation matrix corresponding to σ , then the distribution of $(\mathbf{P}_\sigma \mathbf{D}_1 \mathbf{P}_\sigma^\top)_H$ is identical to the distribution $(\mathbf{P}_\sigma \mathbf{D}_2 \mathbf{P}_\sigma^\top)_H$. We call this property *H-subgraph equivalence*. This implies that any algorithm which queries the edges in H inside of $\mathbf{P}_\sigma \mathbf{D}_1 \mathbf{P}_\sigma^\top$ or $\mathbf{P}_\sigma \mathbf{D}_2 \mathbf{P}_\sigma^\top$ will be unable to distinguish between them with any advantage. To obtain a lower bound, we must also have a gap between $\lambda_{\min}(\mathbf{D}_1)$ and $\lambda_{\min}(\mathbf{D}_2)$, so that their spectrum can be shifted to make one PSD and the other far. Furthermore, neither $\lambda_{\min}(\mathbf{D}_1)$ or $\lambda_{\min}(\mathbf{D}_2)$ can be too negative, otherwise by shifting we would lose boundedness of the entries.

⁸Note that $\nu(H) = 1$ for a triangle H , so the $\Omega(k^{2(\ell-1)/\ell})$ lower bound when $\nu(H) < \ell$ is actually loose here.

A priori, it is not even clear that such matrices $\mathbf{D}_1, \mathbf{D}_2$ exist, even for fixed values of $v(H)$, such as $v(H) = 5$. Our main contribution now is to demonstrate their existence for every $v(H)$. Our construction is simple, but perhaps surprisingly so. Both $\mathbf{D}_1, \mathbf{D}_2$ will be adjacency matrices; in the PSD case, we set \mathbf{D}_1 to be the cycle graph C_{2m+1} on $2m + 1 = \Theta(\log k)$ vertices, and in the ϵ -far case we set \mathbf{D}_2 to be the disjoint union of two cycles $C_{m+1} \oplus C_m$. Since one of m and $m + 1$ is even, while $2m + 1$ is odd, we will have that $\lambda_{\min}(C_{m+1} \oplus C_m) = -2$, but $\lambda_{\min}(C_{2m+1}) > -2$.⁹ To show subgraph equivalence, it suffices to show a slightly more general version of the following: for any graph H with $v(H) < m/4$, the number of subgraphs of C_{2m+1} isomorphic to H is the same as the number of subgraphs of $C_{m+1} \oplus C_m$ isomorphic to H .¹⁰ Note that if $v(H) < m/4$, then H is just a disjoint collection of paths.

Our proof of this fact is by a construction of a bijection from arrangements of H in C_{2m+1} to H in $C_{m+1} \oplus C_m$. While a seemingly simple property, some care must be taken when designing a bijection. Our mapping involves first “swapping” two paths (whose length depends on H) in C_{2m+1} , before “splitting” C_{2m+1} into two cycles of length m and $m + 1$. We direct the reader to Section 5.2.1 for further details.

Amplifying the Gap. The subgraph equivalence between C_{2m+1} and $C_{m+1} \oplus C_m$ prevents any algorithm from distinguishing between them with a small number of samples, however the gap in the minimum eigenvalue shrinks at the rate of $\Theta(1/m^2)$. Meaning, if we set $\gamma = \lambda_{\min}(C_{2m+1}) = 2 - \Theta(1/m^2)$, while the matrix $\gamma\mathbf{I} + C_{2m+1}$ is PSD and has constant sized entries, we only have $\lambda_{\min}(\gamma\mathbf{I} + C_{m+1} \oplus C_m) = -\Theta(1/m^2)$, which is not far enough from PSD. Instead, recall that we only need $m = \Omega(\log k)$ to fool all H with $v(H) \leq \log k$, but the block size which we must fill is much larger: $\mathbf{A}_{B_i \times B_i}$ has size $|B_i| = n/k$. Thus, instead of setting $m = \Theta(n/k)$ and filling all of $\mathbf{A}_{B_i \times B_i}$ with the cycles, we set $m = \Theta(\log k)$, and we amplify the spectral gap by taking the tensor product of the small graphs C_{2m+1} and $C_{m+1} \oplus C_m$ with a large, fixed matrix \mathbf{M} , so that $(\gamma\mathbf{I} + C_{2m+1}) \otimes \mathbf{M}$ has $|B_i|$ rows and columns. We prove that taking the tensor product with any fixed \mathbf{M} preserves the subgraph equivalence properties of the original matrices. From here, our lower bounds for testing PSD with ℓ_2 gap, Schatten norms, Ky fan, and the cost of the best rank- k approximation, all follow by a proper choice of \mathbf{M} . For PSD testing, we can choose $\mathbf{M} = \mathbf{1}$ to be the all 1’s matrix, and to amplify the gap in Schatten 1 norm, we can choose \mathbf{M} to be a random Rademacher matrix. Since $\mathbf{M} = \mathbf{1}$ is PSD and $\|\mathbf{M}\|_2 = \tilde{\Omega}(n/k)$, the gap is amplified to the desired $-\tilde{\Omega}(n/k)$. Finally, we remark that to obtain a lower bound for another norm, any matrix \mathbf{M} which is large in that norm may be suitable, so long as the original sub-graph equivalent matrices also have a gap in that norm. We pose it as an interesting open problem to design other pairs of matrices $\mathbf{D}_1, \mathbf{D}_2$ with different spectral gaps which have good sub-graph equivalence properties.

2 Preliminaries

We now introduce the notation and definitions that will be used consistently throughout the paper. Additional, specialized notation will be introduced as needed in their respective sections.

⁹To intuitively see why this is true, note that if m is even and $v \in \{-1, 1\}^m$ is the vector that assigns opposite signs to adjacent vertices of C_m , then we have $C_m v = -2v$. However, if m is odd, this assignment v is no longer possible.

¹⁰A more general statement is needed since H can also query for edges which do not exist in C_{2m+1} .

Specifically, our lower bound construction in Section 5 utilizes several additional pieces of notation, such as those concerning signed graphs, which are introduced at the beginning of that section.

Singular Values and Eigenvalues. We use boldface \mathbf{A} notation to denote matrices. For a $n \times d$ matrix \mathbf{A} , let $\sigma_{\max}(\mathbf{A}) = \sigma_1(\mathbf{A}) \geq \sigma_2(\mathbf{A}) \geq \dots \geq \sigma_{\min\{n,d\}}(\mathbf{A}) = \sigma_{\min}(\mathbf{A})$ denote the singular values of \mathbf{A} . If \mathbf{A} is rank r , let $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ be its singular value decomposition, where $\mathbf{U} \in \mathbb{R}^{n \times r}$, $\mathbf{V} \in \mathbb{R}^{d \times r}$ have orthonormal columns, and $\mathbf{\Sigma} \in \mathbb{R}^{r \times r}$ is a diagonal matrix with the (non-zero) singular values σ_i on the diagonal. We use $\mathbf{\Sigma}_k$ to denote the matrix $\mathbf{\Sigma}$ but with all entries but the k largest singular values removed and use $\mathbf{\Sigma}_{-k}$ to denote the matrix $\mathbf{\Sigma}$ but with all entries but the $n - k$ smallest singular values removed. Let $\mathbf{A}_k = \mathbf{U}\mathbf{\Sigma}_k\mathbf{V}^\top$ and $\mathbf{A}_{-k} = \mathbf{U}\mathbf{\Sigma}_{-k}\mathbf{V}^\top$. The matrix \mathbf{A}_k is referred to as the *truncated SVD* of \mathbf{A} , and is the best rank- k approximation to \mathbf{A} : $\|\mathbf{A} - \mathbf{A}_k\|_F^2 = \sum_{i>k} \sigma_i^2(\mathbf{A}) = \min_{\mathbf{B} \text{ rank-}k} \|\mathbf{A} - \mathbf{B}\|_F^2$. For the special case when $\mathbf{A} \in \mathbb{R}^{n \times n}$ is symmetric, we use $\mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$ to denote the Eigenvalue Decomposition of \mathbf{A} , where $\lambda_{\max}(\mathbf{A}) = \lambda_1(\mathbf{A}) \geq \lambda_2(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A}) = \lambda_{\min}(\mathbf{A})$ denote the eigenvalues of \mathbf{A} . A real-symmetric matrix \mathbf{A} is said to be *Positive Semi-Definite* (PSD) if $\lambda_{\min} \geq 0$, which is equivalent to having $x^\top \mathbf{A} x \geq 0$ for all $x \in \mathbb{R}^n$. We will utilize the Loewner ordering on symmetric matrices.

Definition 2.1 (Loewner Ordering). For symmetric matrices \mathbf{B}, \mathbf{D} , we write $\mathbf{B} \geq \mathbf{D}$ if $\mathbf{B} - \mathbf{D}$ is PSD.

Notice that if $\mathbf{B} \geq \mathbf{D}$, then by definition $x^\top \mathbf{B} x \geq x^\top \mathbf{D} x$ for all $x \in \mathbb{R}^n$. Then by an application of the Courant-Fischer variational principle for eigenvalues, we have that $\lambda_i(\mathbf{B}) \geq \lambda_i(\mathbf{D})$ for all $i \in [n]$.

Matrix Norms and Submatrices. We use the notation $\|\mathbf{A}\|_2 = \sigma_{\max}(\mathbf{A})$ to denote the spectral norm of \mathbf{A} , $\|\mathbf{A}\|_F = (\sum_{i,j} A_{i,j}^2)^{1/2} = (\sum_{i=1}^n \sigma_i^2(\mathbf{A}))^{1/2}$ to denote the Frobenius norm of \mathbf{A} . For $p \geq 1$, we write $\|\mathbf{A}\|_{S_p} = (\sum_{i=1}^n \sigma_i^p(\mathbf{A}))^{1/p}$ to denote the Schatten p -norm of \mathbf{A} , and $\|\mathbf{A}\|_{\text{KF}(p,k)} = (\sum_{i=1}^k \sigma_i^p(\mathbf{A}))^{1/p}$ to denote the (p, k) -Ky-Fan norm of \mathbf{A} . If p is not specified for a Ky-Fan norm, it is assumed to be 1, namely $\|\mathbf{A}\|_{\text{KF}(k)} = \|\mathbf{A}\|_{\text{KF}(1,k)}$. For subsets $S, T \subseteq [n]$, we denote the matrix $\mathbf{A}_{S \times T} \in \mathbb{R}^{|S| \times |T|}$ as the matrix \mathbf{A} restricted to the submatrix of the rows in S and the columns in T . If $S = T$, then the square submatrix $\mathbf{A}_{S \times T} = \mathbf{A}_{S \times S}$ is called a *principal submatrix* of \mathbf{A} . For a vector $x \in \mathbb{R}^n$ and subset $S \subseteq [n]$, we write $x_S \in \mathbb{R}^n$ to denote the vector obtained after setting equal to zero all coordinates x_i with $i \notin S$. Finally, we use the notation $\mathbf{A}_{i,*}$ to denote the i -th row of \mathbf{A} , and $\mathbf{A}_{*,i}$ to denote the i -th column of \mathbf{A} .

3 PSD Testing with ℓ_∞ Gap

In this section, we introduce our algorithm for the PSD testing problem with ℓ_∞ -gap. As discussed earlier, we consider a more general version of the ℓ_∞ gap than the definition presented in Problem 1.1, which allows one to test a notion of positive semi-definiteness which applies to non-symmetric matrices as well. Specifically, we define the *PSD* case as when $x^\top \mathbf{A} x \geq 0$ for all $x \in \mathbb{R}^n$, and the *far* case as when $x^\top \mathbf{A} x < -\epsilon n$ for a unit vector x . We note that if \mathbf{A} is symmetric, this definition is equivalent to Problem 1.1. In fact, as we will see shortly, one can always reduce the non-symmetric case to the symmetric case, so this distinction will not matter algorithmically. Formally, we solve the following problem:

Definition 3.1 (General PSD Testing with ℓ_∞ -Gap.). Fix, $\epsilon \in (0, 1]$ and let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be any matrix satisfying $\|\mathbf{A}\|_\infty \leq 1$. The goal is to distinguish between the following two cases:

- **YES Instance:** \mathbf{A} satisfies $x^\top \mathbf{A}x \geq 0$, for all $x \in \mathbb{R}^n$.
- **NO Instance:** There exists a unit vector $x \in \mathbb{R}^n$ such that $x^\top \mathbf{A}x < -\epsilon n$.

with probability at least $2/3$.

Reducing to the symmetric case In the case where \mathbf{A} is symmetric, as in Problem 1.1, the above gap instance can be restated in terms of the minimum eigenvalue of \mathbf{A} . Specifically, we are promised that either $\lambda_{\min}(\mathbf{A}) \geq 0$ or $\lambda_{\min}(\mathbf{A}) \leq -\epsilon n$. However, we now observe that one can reduce to the symmetric case with only a factor of 2 loss in the query complexity, by simply querying the symmetrization $(\mathbf{A} + \mathbf{A}^\top)/2$. First note, that for any $x \in \mathbb{R}^n$ and any matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, we have $x^\top \mathbf{A}x = x^\top \mathbf{A}^\top x$, thus for any x we have $x^\top \mathbf{A}x = x^\top \frac{\mathbf{A} + \mathbf{A}^\top}{2} x$. Thus $x^\top \mathbf{A}x \geq 0$ for all x if and only if $x^\top \frac{\mathbf{A} + \mathbf{A}^\top}{2} x \geq 0$ for all x , which occurs if and only if the matrix $\frac{\mathbf{A} + \mathbf{A}^\top}{2}$ is PSD. Similarly, we have that $x^\top \mathbf{A}x \leq -\epsilon n$ for some unit vector x if and only if $x^\top \frac{\mathbf{A} + \mathbf{A}^\top}{2} x \leq -\epsilon n$ for some unit vector x , which occurs if and only if $\lambda_{\min}(\frac{\mathbf{A} + \mathbf{A}^\top}{2}) \leq -\epsilon n$. Note also that the matrix $\frac{\mathbf{A} + \mathbf{A}^\top}{2}$ has bounded entries $\|\frac{\mathbf{A} + \mathbf{A}^\top}{2}\|_\infty \leq 1$ if $\|\mathbf{A}\|_\infty \leq 1$. Moreover, query access to $\frac{\mathbf{A} + \mathbf{A}^\top}{2}$ can be simulated via query access to $\frac{\mathbf{A} + \mathbf{A}^\top}{2}$ with a loss of at most a factor of 2 in the query complexity, by symmetrizing the queries. In fact, our algorithms will not even incur this factor of 2 loss, since all queries our algorithms make will belong to principal submatrices of \mathbf{A} . Thus, in what follows, we can restrict ourselves to the original formulation as specified in Problem 1.1, and assume our input \mathbf{A} is symmetric.

The goal of this section is now to prove the following theorem, which demonstrate the existence of a $\tilde{O}(1/\epsilon^2)$ query one-sided error tester for the above problem. In Section 5, we demonstrate that this complexity is optimal (up to $\log(1/\epsilon)$ factors), even for testers with two sided error (Theorem 5.1).

Theorem 3.2 (Query Optimal One-Sided Tester for ℓ_∞ Gap (see Theorem 3.25)). *There is an algorithm which, given \mathbf{A} with $\|\mathbf{A}\|_\infty \leq 1$ such that either $x^\top \mathbf{A}x \geq 0$ for all x (YES case), or such that $x^\top \mathbf{A}x \leq -\epsilon n$ for some $x \in \mathbb{R}^n$ with $\|x\|_2 \leq 1$ (NO case), distinguishes the two cases with probability at least $3/4$, while making at most $\tilde{O}(\frac{1}{\epsilon^2})$ queries to the entries of \mathbf{A} , and runs in time $\tilde{O}(1/\epsilon^\omega)$, where $\omega < 2.373$ is the exponent of matrix multiplication. Moreover, in the first case when $x^\top \mathbf{A}x \geq 0$ for all x , is PSD, the algorithm always correctly outputs YES.*

Algorithmic Setup First recall that if \mathbf{A} is PSD, then every *principal* submatrix $\mathbf{A}_{T \times T}$ of \mathbf{A} for $T \subseteq [n]$ is also PSD. Thus, it will suffice to query a collection of principal submatrices $\mathbf{A}_{T_1 \times T_1}, \mathbf{A}_{T_2 \times T_2}, \dots, \mathbf{A}_{T_t \times T_t}$ of \mathbf{A} , and return Not PSD if any one of them fails to be PSD. Such a algorithm then naturally has one-sided error, since if \mathbf{A} is PSD it will always return PSD. Thus, in the remainder of the section, it suffices to consider only the NO case, and demonstrate that, if \mathbf{A} satisfies $x^\top \mathbf{A}x \leq -\epsilon n$ for some unit vector $x \in \mathbb{R}^n$, then with good probability at least one of the sampled principal submatrices will fail to be PSD.

Moreover, as shown above, it suffices to consider the case where \mathbf{A} is symmetric. In this case, we will fix x to be the eigenvector associated with the smallest eigenvalue of \mathbf{A} . Thus, in what follows, we can fix x so that $\min_{z \in \mathbb{R}^n: \|z\|_2=1} z^\top \mathbf{A}z = x^\top \mathbf{A}x = \lambda_{\min}(\mathbf{A}) = -\epsilon n$. Notice here we define ϵ to satisfy the equality $\lambda_{\min}(\mathbf{A}) = -\epsilon n$, however our algorithm need only know a lower bound $\epsilon_0 < \epsilon$ on ϵ . The reason for this is that the input parameter ϵ_0 will only effect the sizes of the random submatrices being sampled (smaller ϵ_0 increases the size). Thus, an algorithm run with parameter

ϵ_0 can be simulated by first running the algorithm with parameter $\epsilon > \epsilon_0$, and then randomly adding rows and columns to the sampled submatrices from the correct marginal distribution. Thus, there is a coupling such that the submatrices chosen by an algorithm with any input $\epsilon_0 < \epsilon$ will always contain the submatrices sampled by an algorithm given the input exactly ϵ , so if the latter algorithm sampled a non-PSD submatrix, so would the former. Thus, for the sake of analysis, we can assume that the value ϵ is known.

Throughout the following section, we will assume $1/\epsilon < c \cdot n$ for some sufficiently small constant c . Notice that if this was not the case, we would have $1/\epsilon^2 = \Omega(n^2)$, and we would be permitted to read the entire matrix \mathbf{A} , as this is within our target budget of $\tilde{O}(1/\epsilon^2)$.

3.1 Warm-up: a $O(1/\epsilon^3)$ algorithm

We first describe a $O(1/\epsilon^3)$ query algorithm for the problem of PSD testing with ℓ_∞ -gap. The general approach and results of this algorithm will be needed for the more involved $\tilde{O}(1/\epsilon^2)$ query algorithm which we shall develop in the sequel. As noted above, it suffices to analyze the NO case, where we have $x^\top \mathbf{A} x = \lambda_{\min}(\mathbf{A}) = -\epsilon n$ for a unit vector $x \in \mathbb{R}^n$. Our goal will be to analyze the random variable $Z = x_T^\top \mathbf{A}_{T \times T} x_T$, where $T \subset [n]$ is a random subset, where each $i \in [n]$ is selected independently with some probability δ . Notice that if $\delta_i \in \{0, 1\}$ is an indicator variable indicating that we sample $i \in T$, then we have $Z = x_T^\top \mathbf{A}_{T \times T} x_T = \sum_{i,j} x_i \mathbf{A}_{i,j} x_j \delta_i \delta_j$.

Now, algorithmically, we do not know the vector x . However, if we can demonstrate concentration of Z , and show that $Z < 0$ for our sampled set S , then we can immediately conclude that $\mathbf{A}_{T \times T}$ is not PSD, a fact which *can* be tested. Thus, the analysis will proceed by pretending that we did know x , and analyzing the concentration of $x_T^\top \mathbf{A}_{T \times T} x_T$. In the following section, however, we will ultimately analyze the concentration of this random variable with respects a slightly different vector than x .

We first remark that we can assume, up to a loss in a constant factor in the value of ϵ , that the diagonal of \mathbf{A} is equal to the identity.

Proposition 3.3. *We can assume $\mathbf{A}_{i,i} = 1$, for all $i \in [n]$. Specifically, by modifying \mathbf{A} so that $\mathbf{A}_{i,i} = 1$, for all $i \in [n]$, the completeness (PSD) case is preserved and the soundness (not PSD) case is preserved up to a factor of $1/2$ in the parameter ϵ .*

Proof. Every time we observe an entry $\mathbf{A}_{i,i}$ we set it equal to 1. In this new matrix, if \mathbf{A} was PSD to begin with, then \mathbf{A} will still be PSD, since this modification corresponds to adding a non-negative diagonal matrix to \mathbf{A} . If $x^\top \mathbf{A} x \leq -\epsilon n$ originally for some $x \in \mathbb{R}^n$, then $x^\top \mathbf{A} x \leq -\epsilon n + 1$ after this change, since the diagonal contributes at most $\sum_i \mathbf{A}_{i,i} x_i^2 \leq \|x\|_2^2 \leq 1$ to the overall quadratic form. Note this additive term of 1 is at most $(\epsilon n)/2$ since we can assume $\epsilon = \Omega(1/n)$. ■

We now notice that if x is the eigenvector associated with a large enough eigenvalue, the ℓ_2 mass of x cannot be too concentrated.

Proposition 3.4. *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a symmetric matrix with $\lambda_{\min}(\mathbf{A}) = -\epsilon n$, and let $x \in \mathbb{R}^n$ be the (unit) eigenvector associated with $\lambda_{\min}(\mathbf{A})$. Then we have that $\|x\|_\infty \leq \frac{1}{\epsilon \sqrt{n}}$.*

Proof. By Cauchy-Schwartz, for any $i \in [n]$:

$$|\lambda_{\min}| \cdot |x_i| = |\langle \mathbf{A}_{i,*}, x \rangle| \leq \|\mathbf{A}_{i,*}\|_2 \leq \sqrt{n}$$

from which the proposition follows using $\lambda_{\min}(\mathbf{A}) = -\epsilon n$. ■

Recall that our goal is to analyze the random variable $Z = x_T^\top \mathbf{A}_{T \times T} x_T = \sum_{i,j} x_i \mathbf{A}_{i,j} x_j \delta_i \delta_j$. To proceed, we bound the moments of Z . Our bound on these moments can be tightened as a function of the *row and column contributions* of the target vector x , which we now define.

Definition 3.5. Fix any $y \in \mathbb{R}^n$. Then for any $i \in [n]$, define the total row and column contributions of i as $\mathcal{R}_i(y) = \sum_{j \in [n] \setminus i} y_i \mathbf{A}_{i,j} y_j$ and $\mathcal{C}_i(y) = \sum_{j \in [n] \setminus i} y_j \mathbf{A}_{j,i} y_i$ respectively.

Notice from the above definition, we have $\sum_i \mathcal{R}_i(y) + \mathcal{C}_i(y) = 2 \left(y^\top \mathbf{A} y - \sum_i \mathbf{A}_{i,i} y_i^2 \right)$.

Fact 3.6. Let $x \in \mathbb{R}^n$ be the eigenvector associated with $\lambda_{\min}(\mathbf{A})$. Then we have $\mathcal{R}_i(x) + \mathcal{C}_i(x) \leq 0$ for all $i \in [n]$.

Proof. Suppose there was an i with $\mathcal{R}_i(x) + \mathcal{C}_i(x) > 0$. Then setting $z = x_{[n] \setminus i}$ we have $z^\top \mathbf{A} z = \langle x, \mathbf{A} x \rangle - (\mathcal{R}_i(x) + \mathcal{C}_i(x)) - \mathbf{A}_{i,i} x_i^2$. Recall from Proposition 3.3 that we can assume $\mathbf{A}_{i,i} = 1$ for all i , thus it follows that $z^\top \mathbf{A} z < \langle x, \mathbf{A} x \rangle$, which contradicts the optimality of x . ■

We now bound the expectation of the random quadratic form.

Proposition 3.7 (Expectation Bound). Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a matrix with $\|\mathbf{A}\|_\infty \leq 1$, and let $y \in \mathbb{R}^n$ be any vector with $\|y\|_2 \leq 1$ and $y^\top \mathbf{A} y < -\epsilon n$. Let $Z = \sum_{i,j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j$, where $\delta_1, \dots, \delta_n \sim \text{Bernoulli}(\frac{k}{n})$. Then if $k \geq 8/\epsilon$, we have $\mathbb{E}[Z] \leq -\frac{\epsilon k^2}{4n}$.

Proof. Let $c_{i,j} = \mathbf{A}_{i,j} y_i y_j$. First note, for any $i \in [n]$, the term $c_{i,j}$ is included in T with probability k/n if $i = j$, and with probability k^2/n^2 if $i \neq j$. So

$$\begin{aligned} \mathbb{E}[Z] &= \sum_{i \neq j} \frac{k^2}{n^2} c_{i,j} + \sum_{i \in [n]} \frac{k}{n} c_{i,i} \\ &= \frac{k^2}{n^2} \left(\langle y, \mathbf{A} y \rangle - \sum_{i \in [n]} \mathbf{A}_{i,i} y_i^2 \right) + \frac{k}{n} \sum_{i \in [n]} \mathbf{A}_{i,i} y_i^2 \\ &\leq -\frac{\epsilon k^2}{2n} + \left(\frac{k}{n} + \frac{k^2}{n^2} \right) \sum_{i \in [n]} y_i^2 \\ &\leq -\frac{\epsilon k^2}{2n} + \frac{2k}{n} \leq -\frac{\epsilon k^2}{4n} \end{aligned} \tag{2}$$

Where in the last inequality, we assume $k \geq 8/\epsilon$. ■

Next, we bound the variance of Z . We defer the proof of the following Lemma to Section 3.2.

Lemma 3.8 (Variance Bound). Let $\delta_1, \dots, \delta_n \sim \text{Bernoulli}(\frac{k}{n})$. Let $y \in \mathbb{R}^n$ be any vector such that $\|y\|_2 \leq 1$, $\|y\|_\infty \leq \frac{1}{\epsilon \sqrt{n}}$, and $y^\top \mathbf{A} y = -\epsilon n$, where $\mathbf{A} \in \mathbb{R}^{n \times n}$ satisfies $\|\mathbf{A}\|_\infty \leq 1$. Further suppose that $\mathcal{R}_i(y) + \mathcal{C}_i(y) \leq 0$ for each $i \in [n]$. Then, assuming $k \geq 6/\epsilon$, we have

$$\text{Var} \left[\sum_{i,j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] \leq O \left(\frac{k^3}{n^2} \right)$$

Moreover, if the tighter bound $\|y\|_\infty \leq \frac{\alpha}{\epsilon\sqrt{n}}$ holds for some $\alpha \leq 1$, we have

$$\mathbf{Var} \left[\sum_{i,j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] \leq O \left(\frac{k^2}{n^2} + \frac{\alpha k^3}{n^2} \right)$$

We note that the variance of the random quadratic form can be improved if we have tighter bounds on certain properties of the target vector y . We demonstrate this fact in the following Corollary, which we will use in Section 3.3. Note that the assumptions of Corollary 3.9 differ in several minor ways from those of Lemma 3.8. For instance, we do not require $k \geq 6/\epsilon$ (we note that this assumption was required only to simplify the expression in Lemma 3.8). Also notice that we do not bound the diagonal terms in Corollary 3.9. We defer the proof of Corollary 3.9 to Section 3.2.

Corollary 3.9 (Tighter Variance Bound). *Let $\delta_1, \dots, \delta_n \sim \text{Bernoulli}(\frac{k}{n})$. Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$ be any matrix and y a vector such that $|y^\top \mathbf{A} y| \leq c_1 \epsilon n$ for some value $c_1 > 0$, and such that $\|y\|_\infty \leq \frac{\alpha}{\epsilon\sqrt{n}}$ for some $\alpha > 0$. Let $\mathbf{Z} \in \mathbb{R}^n$ be defined by $\mathbf{Z}_i = \mathcal{R}_i(y) + C_i(y)$ for $i \in [n]$, and suppose we have $\|\mathbf{Z}\|_2^2 \leq c_2 \epsilon n$. Then we have*

$$\mathbf{Var} \left[\sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] \leq O \left(\frac{k^2}{n^2} + \frac{c_1^2 k^4 \epsilon^2}{n^2} + \frac{(c_1 + c_2) \epsilon k^3}{n^2} + \frac{\alpha^2 k^3}{n^2} \right)$$

We now observe that the variance computations from Lemma 3.8 immediately gives rise to a $O(1/\epsilon^3)$ algorithm.

Theorem 3.10. *There is a non-adaptive sampling algorithm which queries $O(\epsilon^{-3})$ entries of \mathbf{A} , and distinguishes the case that \mathbf{A} is PSD from the case that $\lambda_{\min}(\mathbf{A}) < -\epsilon n$ with probability $2/3$.*

Proof. Let $x \in \mathbb{R}^n$ be the eigenvector associated with $\lambda_{\min}(\mathbf{A}) = -\epsilon n$ (recall that we can assume equality), and let Z_1, \dots, Z_d be independent repetitions of the above process, with $k = 10/\epsilon$ and $d = 3840/\epsilon$. Let $Z = \frac{1}{d} \sum_{i=1}^d Z_i$. Then $\mathbf{Var}(Z) \leq \frac{6}{d} \frac{k^3}{n^2}$ by Lemma 3.8, where we used the bounds on $\|x\|_\infty$ from Proposition 3.4 and the property that $\mathcal{R}_i(x) + C_i(x) \leq 0$ for all i from Fact 3.6 to satisfies the assumptions of Lemma 3.8. By Chebysev's inequality:

$$\begin{aligned} \Pr \left[Z \geq -\frac{\epsilon k^2}{4n} + \frac{\epsilon k^2}{8n} \right] &\leq \left(\frac{64n^2}{\epsilon^2 k^4} \right) \left(\frac{6k^3}{dn^2} \right) \\ &\leq \frac{1}{10\epsilon k} \\ &\leq \frac{1}{100} \end{aligned} \tag{3}$$

It follows that with probability $99/100$, the average of the Z_i 's will be negative. Thus at least one of the Z_i 's must be negative, thus the submatrix corresponding to this Z_i will not be PSD. The total query complexity is $O(k^2 d) = O(\epsilon^{-3})$. ■

3.2 Variance Bounds

In this section, we provide the proofs of the variance bounds in Lemma 3.8 and Corollary 3.9. For convenience, we restate the Lemma and Corollary here before the proofs.

Lemma 3.8 Let $\delta_1, \dots, \delta_n \sim \text{Bernoulli}(\frac{k}{n})$. Let $y \in \mathbb{R}^n$ be any vector such that $\|y\|_2 \leq 1$, $\|y\|_\infty \leq \frac{1}{\epsilon\sqrt{n}}$, and $y^\top \mathbf{A} y = -\epsilon n$, where $\mathbf{A} \in \mathbb{R}^{n \times n}$ satisfies $\|\mathbf{A}\|_\infty \leq 1$. Further suppose that $\mathcal{R}_i(y) + C_i(y) \leq 0$ for each $i \in [n]$. Then, assuming $k \geq 6/\epsilon$, we have

$$\text{Var} \left[\sum_{i,j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] \leq O \left(\frac{k^3}{n^2} \right)$$

Moreover, if the tighter bound $\|y\|_\infty \leq \frac{\alpha}{\epsilon\sqrt{n}}$ holds for some $\alpha \leq 1$, we have

$$\text{Var} \left[\sum_{i,j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] \leq O \left(\frac{k^2}{n^2} + \frac{\alpha k^3}{n^2} \right)$$

Proof. Let $c_{i,j} = \mathbf{A}_{i,j} y_i y_j$. We have

$$\begin{aligned} \text{Var} \left[\sum_{i,j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] &\leq \frac{k}{n} \sum_i c_{i,i}^2 + \frac{k^2}{n^2} \sum_{i \neq j} c_{i,j}^2 + \frac{k^2}{n^2} \sum_{i \neq j} c_{i,j} c_{j,i} + \frac{k^2}{n^2} \sum_{i \neq j} c_{i,i} c_{j,j} + \frac{k^2}{n^2} \sum_{i \neq j} c_{i,i} c_{i,j} \\ &+ \frac{k^2}{n^2} \sum_{i \neq j} c_{i,i} c_{j,i} + \frac{k^3}{n^3} \sum_{i \neq j \neq u} c_{i,j} c_{u,j} + \frac{k^3}{n^3} \sum_{j \neq i \neq u} c_{i,j} c_{i,u} + \frac{k^3}{n^3} \sum_{i \neq j \neq u} c_{i,j} c_{j,u} + \frac{k^3}{n^3} \sum_{j \neq i \neq u} c_{i,j} c_{u,i} \\ &+ \frac{2k^3}{n^3} \sum_{i \neq j \neq u} c_{i,i} c_{j,u} + \frac{k^4}{n^4} \sum_{i \neq j \neq v \neq u} c_{i,j} c_{u,v} - \left(\frac{k^2}{n^2} \sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j - \frac{k}{n} \sum_i \mathbf{A}_{i,i} y_i^2 \right)^2 \end{aligned} \quad (4)$$

We first consider the last term $\frac{k^4}{n^4} \sum_{i \neq j \neq v \neq u} c_{i,j} c_{u,v} = \sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \sum_{u \neq v \neq i \neq j} y_u \mathbf{A}_{u,v} y_v$. Here $i \neq j \neq v \neq u$ means all 4 indices are distinct. Note that this term is canceled by a subset of the terms within $\left(\frac{k^2}{n^2} \sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j - \frac{k}{n} \sum_i \mathbf{A}_{i,i} y_i^2 \right)^2$. Similarly, the term $\frac{k^2}{n^2} \sum_{i \neq j} c_{i,i} c_{j,j}$ cancels. Moreover, after expanding $\left(\frac{k^2}{n^2} \sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j - \frac{k}{n} \sum_i \mathbf{A}_{i,i} y_i^2 \right)^2$, every remaining term which does not cancel with another term exactly is equal to another term in the variance above, but with an additional one (or two) factors of $\frac{k}{n}$ attached. Thus, if we can bound the remaining terms in Equation 4 by some value B , then an overall variance bound of $2 \cdot B$ will follow.

We now consider $\mathcal{T} = \left(\sum_{j \neq i \neq u} c_{i,j} c_{i,u} + \sum_{i \neq j \neq u} c_{i,j} c_{u,j} + \sum_{j \neq i \neq u} c_{i,j} c_{u,i} + \sum_{i \neq j \neq u} c_{i,j} c_{j,u} \right)$. We have

$$\begin{aligned} \sum_{i \neq j \neq u} c_{i,j} c_{i,u} &= \sum_i \sum_{j \neq i} y_i \mathbf{A}_{i,j} y_j \sum_{u \neq i \neq j} y_u \mathbf{A}_{i,u} y_u \\ \sum_{i \neq j \neq u} c_{i,j} c_{u,j} &= \sum_j \sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \sum_{u \neq i \neq j} y_u \mathbf{A}_{u,j} y_j \\ \sum_{j \neq i \neq u} c_{i,j} c_{u,i} &= \sum_i \sum_{j \neq i} y_i \mathbf{A}_{i,j} y_j \sum_{u \neq i \neq j} y_u \mathbf{A}_{u,i} y_i \\ \sum_{i \neq j \neq u} c_{i,j} c_{j,u} &= \sum_j \sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \sum_{u \neq i \neq j} y_j \mathbf{A}_{j,u} y_u \end{aligned}$$

Now for simplicity, we write $\mathcal{R}_i = \mathcal{R}_i(y)$ and $C_i = C_i(y)$ for $i \in [n]$. Then by assumption, we have $\mathcal{R}_i + C_i \leq 0$ for each i , thus $|\sum_i (\mathcal{R}_i + C_i)| = \sum_i |(\mathcal{R}_i + C_i)|$. Also note that we have $|\sum_i (\mathcal{R}_i + C_i)| = |2y^\top \mathbf{A}y - 2\sum_i \mathbf{A}_{i,i}y_i^2| \leq 4\epsilon n$. Now observe

$$\left| \left(\sum_i \sum_{j \neq i} y_i \mathbf{A}_{i,j} y_j \sum_{u \neq i \neq j} y_i \mathbf{A}_{i,u} y_u \right) - \sum_i \mathcal{R}_i^2 \right| = \sum_i \sum_{u \in [n] \setminus i} y_i^2 \mathbf{A}_{i,u}^2 y_u^2 \leq \sum_i y_i^2 \leq 1$$

And similarly

$$\begin{aligned} \left| \left(\sum_j \sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \sum_{u \neq i \neq j} y_u \mathbf{A}_{u,j} y_j \right) - \sum_j C_j^2 \right| &= \sum_j \sum_{u \in [n] \setminus j} y_u^2 \mathbf{A}_{u,j}^2 y_j^2 \leq \sum_j y_j^2 \leq 1 \\ \left| \left(\sum_i \sum_{j \neq i} y_i \mathbf{A}_{i,j} y_j \sum_{u \neq i \neq j} y_u \mathbf{A}_{u,i} y_i \right) - \sum_i \mathcal{R}_i C_i \right| &= \sum_i \sum_{u \in [n] \setminus i} y_i^2 \mathbf{A}_{i,u} \mathbf{A}_{u,i} y_u^2 \leq \sum_i y_i^2 \leq 1 \\ \left| \left(\sum_j \sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \sum_{u \neq i \neq j} y_j \mathbf{A}_{j,u} y_u \right) - \sum_j \mathcal{R}_j C_j \right| &= \sum_j \sum_{u \in [n] \setminus j} y_u^2 \mathbf{A}_{u,j} \mathbf{A}_{j,u} y_j^2 \leq \sum_j y_j^2 \leq 1 \end{aligned}$$

Taking these four equations together, we obtain $|\mathcal{T} - \sum_i (\mathcal{R}_i + C_i)^2| \leq 4$, so it will suffice to upper bound the value $\sum_i (\mathcal{R}_i + C_i)^2$ instead. First note that since $|y_i| \leq \frac{1}{\epsilon\sqrt{n}}$ for all i , so for any $i \in [n]$ we have

$$|(\mathcal{R}_i + C_i)| \leq \left| \sum_{j \neq i} y_i \mathbf{A}_{i,j} y_j \right| + \left| \sum_{j \neq i} y_j \mathbf{A}_{j,i} y_i \right| \leq \frac{1}{\epsilon\sqrt{n}} \left(\sum_j 2y_j \right) \leq \frac{2}{\epsilon\sqrt{n}} \|y\|_1 \leq \frac{2}{\epsilon}$$

Combining this bound with the fact that $\sum_i |(\mathcal{R}_i + C_i)| \leq 4\epsilon n$ from earlier, it follows that the sum $\sum_i (\mathcal{R}_i + C_i)^2$ is maximized by setting $2\epsilon^2 n$ of the terms $(\mathcal{R}_i + C_i)$ equal to the largest possible value of $(2/\epsilon)$, so that $\sum_i (\mathcal{R}_i + C_i)^2 \leq 2\epsilon^2 n (2/\epsilon)^2 = O(n)$. This yields an upper bound of $\frac{k^3}{n^3} \mathcal{T} = O(\frac{k^3}{n^2})$. Note, that in general, given the bound $\|y\|_\infty \leq \frac{\alpha}{\epsilon\sqrt{n}}$ for some value $\alpha \leq 1$, then each term $|(\mathcal{R}_i + C_i)| \leq \frac{2\alpha}{\epsilon}$. On the other hand, $\sum_i |(\mathcal{R}_i + C_i)| \leq 4\epsilon n$. Thus, once again, $\sum_i |(\mathcal{R}_i + C_i)|^2$ is maximized by setting $\Theta(\epsilon^2 n / \alpha)$ inner terms equal to $\Theta((\frac{\alpha}{\epsilon})^2)$, giving $\mathcal{T} \leq \alpha n$ for general $\alpha < 1$. Thus, for general $\alpha \leq 1$, we have $\frac{k^3}{n^3} \mathcal{T} = O(\frac{\alpha k^3}{n^2})$.

Next, we bound $\frac{k^2}{n^2} \sum_{i \neq j} c_{i,i} c_{i,j} + \frac{k^2}{n^2} \sum_{i \neq j} c_{i,i} c_{j,i}$ by $\frac{k^2}{n^2} \sum_i y_i^2 (\mathcal{R}_i + C_i)$. As shown above, $|\mathcal{R}_i + C_i| \leq 2y_i \sqrt{n}$, thus altogether we have

$$\frac{k^2}{n^2} \left(\sum_{i \neq j} c_{i,i} c_{i,j} + \sum_{i \neq j} c_{i,i} c_{j,i} \right) \leq \frac{k^2}{n^2} \sum_i y_i^3 \sqrt{n} \quad (5)$$

Using that $\|y\|_\infty \leq \frac{\alpha}{\epsilon\sqrt{n}}$ for $\alpha \leq 1$, and the fact that $\|y\|_2^2 \leq 1$, it follows that $\|y\|_3^3$ is maximized by having $\frac{n\epsilon^2}{\alpha^2}$ terms equal to $\|y\|_\infty \leq \frac{\alpha}{\epsilon\sqrt{n}}$, which gives an upper bound of $\|y\|_3^3 \leq \frac{\alpha}{\epsilon\sqrt{n}}$. Thus, we can bound the right hand side of Equation 5 by $\frac{k^2 \alpha}{n^2 \epsilon}$, which is $O(k^3/n^2)$ when $\alpha = 1$ using that $k = \Omega(1/\epsilon)$.

Now, we bound $\frac{k^3}{n^3} \sum_{i \neq j \neq u} c_{i,i} c_{j,u}$ by

$$\begin{aligned}
\frac{k^3}{n^3} \sum_{i \neq j \neq u} c_{i,i} c_{j,u} &\leq \frac{k^3}{n^3} \sum_i y_i^2 \mathbf{A}_{i,i} \sum_{j \neq u \neq i} y_j \mathbf{A}_{j,u} y_u \\
&\leq \frac{k^3}{n^3} \sum_i y_i^2 \mathbf{A}_{i,i} \sum_{j \neq u \neq i} y_j \mathbf{A}_{j,u} y_u \\
&\leq \frac{k^3}{n^3} \sum_i y_i^2 \mathbf{A}_{i,i} (\epsilon n + O(1)) \\
&\leq \frac{\epsilon k^3}{n^2} \\
&= O\left(\frac{k^2}{n^2}\right)
\end{aligned} \tag{6}$$

Also observe that $\sum_{i,j} c_{i,j}^2 \leq \sum_{i,j} y_i^2 y_j^2 = \|y\|_2^4 \leq 1$, so $\sum_{i \neq j} c_{i,j}^2 \leq \sum_{i,j} c_{i,j}^2 \leq 1$, and also $\sum_{i \neq j} c_{i,j} c_{j,i} \leq \sum_{i,j} y_i^2 y_j^2 \leq 1$, which bounds their corresponding terms in the variance by $O(k^2/n^2)$. Finally, we must bound the last term $\frac{k}{n} \sum_i c_{i,i}^2 = \frac{k}{n} \sum_i y_i^4 \mathbf{A}_{i,i}^2 \leq \frac{k}{n} \sum_i y_i^4$. Note that $|y_i| \leq 1/(\epsilon\sqrt{n})$ for each i , and $\|y\|_2 \leq 1$. Thus $\sum_i y_i^4$ is maximized when one has $\epsilon^2 n$ terms equal to $1/(\epsilon\sqrt{n})$, and the rest set to 0. So $\sum_i y_i^4 \leq \epsilon^2 n (\frac{1}{\epsilon\sqrt{n}})^4 \leq \frac{1}{\epsilon^2 n}$. In general, if $\|y\|_\infty \leq \frac{\alpha}{\epsilon\sqrt{n}}$, we have $\sum_i y_i^4 \leq \frac{\epsilon^2 n}{\alpha^2} (\frac{\alpha}{\epsilon\sqrt{n}})^4 \leq \frac{\alpha^2}{\epsilon^2 n}$. Thus we can bound $\frac{k}{n} \sum_i c_{i,i}^2$ by $O(\frac{k^3 \alpha^2}{n^2})$.

Altogether, this gives

$$\begin{aligned}
\text{Var} \left[\sum_{i,j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] &\leq O\left(\frac{k^2}{n^2} + \frac{\alpha k^3}{n^2} + \frac{\alpha k^2}{n^2 \epsilon} + \frac{\alpha^2 k^3}{n^2}\right) \\
&= O\left(\frac{k^2}{n^2} + \frac{\alpha k^3}{n^2} + \frac{\alpha^2 k^3}{n^2}\right)
\end{aligned} \tag{7}$$

which is $O(k^3/n^2)$ in general (where $\alpha \leq 1$), where we assume $k \geq 6/\epsilon$ throughout. \blacksquare

Corollary 3.9 *Let $\delta_i \in \{0, 1\}$ be an indicator random variable with $\mathbb{E}[\delta_i] = k/n$. Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$ be any matrix and y a vector such that $|y^\top \mathbf{A} y| \leq c_1 \epsilon n$ for some value $c_1 > 0$, and such that $\|y\|_\infty \leq \frac{\alpha}{\epsilon\sqrt{n}}$ for some $\alpha > 0$. Let $\mathbf{Z} \in \mathbb{R}^n$ be defined by $\mathbf{Z}_i = \mathcal{R}_i(y) + C_i(y)$ for $i \in [n]$, and suppose we have $\|\mathbf{Z}\|_2^2 \leq c_2 \epsilon n$. Then we have*

$$\text{Var} \left[\sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] \leq O\left(\frac{k^2}{n^2} + \frac{c_1^2 k^4 \epsilon^2}{n^2} + \frac{(c_1 + c_2) \epsilon k^3}{n^2} + \frac{\alpha^2 k^3}{n^2}\right)$$

Proof. We proceed as in Lemma 3.8, except that we may remove the terms with $c_{i,j}$ for $i = j$, yielding

$$\begin{aligned} \mathbf{Var} \left[\sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] &\leq \frac{k^2}{n^2} \sum_{i \neq j} c_{i,j}^2 + \frac{k^2}{n^2} \sum_{i \neq j} c_{i,j} c_{j,i} + \frac{k^3}{n^3} \sum_{i \neq j \neq u} c_{i,j} c_{u,j} \\ &+ \frac{k^3}{n^3} \sum_{j \neq i \neq u} c_{i,j} c_{i,u} + \frac{k^3}{n^3} \sum_{i \neq j \neq u} c_{i,j} c_{j,u} + \frac{k^3}{n^3} \sum_{j \neq i \neq u} c_{i,j} c_{u,i} \\ &+ \frac{k^3}{n^3} \sum_{i \neq j \neq u} c_{i,i} c_{j,u} + \frac{k^4}{n^4} \sum_{i \neq j \neq v \neq u} c_{i,j} c_{u,v} - \left(\frac{k^2}{n^2} \sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \right)^2 \end{aligned} \quad (8)$$

As in Lemma 3.8, we can cancel the term $\frac{k^4}{n^4} \sum_{i \neq j \neq v \neq u} c_{i,j} c_{u,v}$ with a subterm of $-\left(\frac{k^2}{n^2} \sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j\right)^2$, and bound the remaining contribution of $-\left(\frac{k^2}{n^2} \sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j\right)^2$ by individually bounding the other terms in the sum.

First, we can similarly bound the last term by $c_1^2 \epsilon^2 k^4 / n^2$ as needed. Now when bounding

$$\mathcal{T} = \left(\sum_{j \neq i \neq u} c_{i,j} c_{i,u} + \sum_{i \neq j \neq u} c_{i,j} c_{u,j} + \sum_{j \neq i \neq u} c_{i,j} c_{u,i} + \sum_{i \neq j \neq u} c_{i,j} c_{j,u} \right)$$

we first observe that in the proof of Lemma 3.8, we only needed a bound on $\|\mathbf{Z}\|_2^2$ to give the bound on \mathcal{T} . So by assumption, $\|\mathbf{Z}\|_2^2 \leq c_2 \epsilon n$, which gives a total bound of $\frac{c_2 k^3 \epsilon}{n^2}$ on $\frac{k^3}{n^3} \mathcal{T}$.

Also, we bound we bound $\frac{k^3}{n^3} \sum_{i \neq j \neq u} c_{i,i} c_{j,u}$ by

$$\begin{aligned} \frac{k^3}{n^3} \sum_{i \neq j \neq u} c_{i,i} c_{j,u} &\leq \frac{k^3}{n^3} \sum_i y_i^2 \mathbf{A}_{i,i} \sum_{j \neq u \neq i} y_j \mathbf{A}_{j,u} y_u \\ &\leq \frac{k^3}{n^3} \sum_i y_i^2 \mathbf{A}_{i,i} \sum_{j \neq u \neq i} y_j \mathbf{A}_{j,u} y_u \\ &\leq \frac{k^3}{n^3} \sum_i y_i^2 \mathbf{A}_{i,i} (c_1 \epsilon n + O(1)) \\ &\leq \frac{\epsilon c_1 k^3}{n^2} \end{aligned} \quad (9)$$

which is within our desired upper bound. Finally observe that $\sum_{i,j} c_{i,j}^2 \leq \sum_{i,j} y_i^2 y_j^2 = \|y\|_2^4 \leq 1$, so $\sum_{i \neq j} c_{i,j}^2 \leq \sum_{i,j} c_{i,j}^2 \leq 1$, and also $\sum_{i \neq j} c_{i,j} c_{j,i} \leq \sum_{i,j} y_i^2 y_j^2 \leq 1$, which bounds their corresponding terms in the variance by $O(k^2/n^2)$, which completes the proof. \blacksquare

3.3 Improving the complexity to $\tilde{O}(1/\epsilon^2)$

We now demonstrate how to obtain an improved sample complexity of $\tilde{O}(1/\epsilon^2)$ using different scales of sub-sampling, as well as a careful “eigenvector switching” argument. As before, we can

assume that \mathbf{A} is symmetric, and $x = \arg \min_{v \in \mathbb{R}^n, \|v\|_2 \leq 1} v^\top \mathbf{A} v$ is the smallest eigenvector of \mathbf{A} , so that $\langle x, \mathbf{A} x \rangle = \lambda_{\min}(\mathbf{A}) = -\epsilon n$. Also recall that our algorithms will not need to explicitly know the value $\epsilon = \min_{v \in \mathbb{R}^n, \|v\|_2 \leq 1} v^\top \mathbf{A} v / n$, only a lower bound on it, since when run on smaller ϵ our algorithm only samples larger submatrices. By Proposition 3.4, we have $\|x\|_\infty \leq \frac{1}{\epsilon \sqrt{n}}$. We now partition the coordinates of x into *level-sets*, such that all the coordinates x_i^2 within a level set have magnitudes that are close to each other.

Definition 3.11. Given (\mathbf{A}, x) , where x is as defined above, define the base level set S as $S = \{i \in [n] : |x_i|^2 \leq \frac{100}{\epsilon n}\}$, and let $T_a = \{i \in [n] : \frac{100 \cdot 2^{a-1}}{\epsilon n} \leq |x_i|^2 \leq \frac{100 \cdot 2^a}{\epsilon n}\}$ for an integer $a \geq 1$.

We now break the analysis into two possible cases. In the first case, the coordinates in one of the sets T_a contributed a substantial fraction of the “negativeness” of the quadratic form $x^\top \mathbf{A} x$, for some a sufficiently large. Since the sets T_a can become smaller as a increases while still contributing a large fraction of the negative mass, this case can be understood as the negativeness of $x^\top \mathbf{A} x$ being highly concentrated in a small fraction of the matrix, which we must then find to determine that \mathbf{A} is not PSD. In the second case, no such contributing T_a exists, and the negative mass is spread out more evenly across the terms in the quadratic form $x^\top \mathbf{A} x$. If this is the case, we will show that our variance bounds from the prior section can be made to obtain a proof that a single, large sampled principal submatrix $T \subset [n]$ will satisfy $x_T^\top \mathbf{A}_{T \times T} x_T$ will be negative with non-negligible probability.

Formally, we define the two cases as follows:

Case 1: We have $x_S \mathbf{A} x_{T_a} + x_{T_a} \mathbf{A} x_S \leq -\frac{\epsilon n}{10 \log(1/\epsilon)}$ for some a such that $2^a \geq 10^6 \zeta^3$, for some $\zeta = \Theta(\log^2(1/\epsilon))$ with a large enough constant.

Case 2: The above does not hold; namely, we have $x_S \mathbf{A} x_{T_a} + x_{T_a} \mathbf{A} x_S > -\frac{\epsilon n}{10 \log(1/\epsilon)}$ for every $2^a \geq 10^6 \zeta^3$.

3.3.1 Case 1: Varied Subsampling and Eigenvector Switching

In this section, we analyze the **Case 1**, which specifies that $x_S \mathbf{A} x_{T_a} + x_{T_a} \mathbf{A} x_S \leq -\frac{\epsilon n}{10 \log(1/\epsilon)}$ for some T_a such that $2^a \geq 10^6 \zeta^3$, where $\zeta = \Theta(\log^2(1/\epsilon))$ is chosen with a sufficiently large constant. Recall here that $x \in \mathbb{R}^n$ is the (unit) eigenvector associated with $\lambda_{\min}(\mathbf{A}) = -\epsilon n$. We now fix this value a associated with T_a . In order to find a principal submatrix $\mathbf{A}_{T \times T}$ that is not PSD for some sampled subset $T \subset [n]$, we will need to show that $T \cap T_a$ intersects in at least one coordinate.

As discussed in Section 1.4.1, we will need to switch our analysis from x to a different vector y , in order to have $y_T^\top \mathbf{A}_{T \times T} y_T < 0$ with non-negligible probability conditioned on $|T \cap T_a| \geq 1$. To construct the appropriate vector y , we will first proceed by proving several propositions which bound how the quadratic form $x^\top \mathbf{A} x$ changes as we modify or remove some of the coordinates of x . For the following propositions, notice that by definition of T_b , using the fact that $\|x\|_2^2 \leq 1$, we have that $|T_b| \leq \frac{\epsilon n}{1002^{b-1}}$ for any $b \geq 1$, which in particular holds for $b = a$.

Proposition 3.12. Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ satisfy $\|\mathbf{A}\|_\infty \leq 1$. Let $S, T \subset [n]$, and let $v \in \mathbb{R}^n$ be any vector such that $\|v\|_2 \leq 1$. Then $|v_S^\top \mathbf{A} v_T| \leq \sqrt{|S| \cdot |T|}$

Proof. We have $|v_S^\top \mathbf{A} v_T| = |\sum_{i \in S} \sum_{j \in T} v_i \mathbf{A}_{i,j} v_j| \leq \sum_{i \in S} |v_i| \sum_{j \in T} |v_j| \leq \sum_{i \in S} |v_i| \|v_T\|_1 \leq \|v_S\|_1 \|v_T\|_1 \leq \sqrt{|S||T|}$ as needed. ■

Proposition 3.13. Let $\mathbf{A} \in \mathbb{R}^{n \times m}$ satisfy $\|\mathbf{A}\|_\infty \leq 1$ for any n, m . and let $v \in \mathbb{R}^n, u \in \mathbb{R}^m$ satisfy $\|u\|_2^2, \|v\|_2^2 \leq 1$. Then

$$\sum_{j=1}^m \left(\sum_{i=1}^n v_i \mathbf{A}_{i,j} u_j \right)^2 \leq n$$

Proof. We have $(\sum_{i=1}^n v_i \mathbf{A}_{i,j} u_j)^2 \leq u_j^2 (\sum_{i=1}^n |v_i|)^2 \leq u_j^2 \|v\|_1^2$, so the sum can be bounded by $\sum_{j=1}^m u_j^2 \|v\|_1^2 = \|v\|_1^2 \|u\|_2^2 \leq \|v\|_1^2 \leq n$ as needed. ■

Proposition 3.14. Let x be as defined above. Then we have $|\langle x_S, \mathbf{A} x_S \rangle| \leq 10\epsilon n$.

Proof. Suppose $\langle x_S, \mathbf{A} x_S \rangle = C\epsilon n$ for a value C with $|C| > 10$. Note that $|\langle x_{[n] \setminus S}, \mathbf{A} x_{[n] \setminus S} \rangle| \leq \frac{\epsilon n}{100}$ by Proposition 3.12, using that $|[n] \setminus S| = |\cup_{b \geq 1} T_b| \leq \frac{\epsilon n}{100}$ (here we use the fact that at most $\frac{\epsilon n}{100}$ coordinates of a unit vector can have squared value larger than $\frac{100}{\epsilon n}$). If $C > 0$, then we must have that $(\langle x_S, \mathbf{A} x_{[n] \setminus S} \rangle + \langle x_{[n] \setminus S}, \mathbf{A} x_S \rangle) \leq -(C + 99/100)\epsilon n$ for us to have that $\langle x, \mathbf{A} x \rangle = -\epsilon n$ exactly. Thus if C is positive and larger than 10, it would follow that by setting $v = x_S/2 + x_{[n] \setminus S}$, we would obtain a vector v with $\|v\|_2 \leq 1$ such that v has smaller quadratic form with \mathbf{A} than x , namely with $v^\top \mathbf{A} v \leq -(C + 99/100)\epsilon n/2 + \epsilon C n/4 + n\epsilon/100 < -\epsilon n$ using that $C > 10$, which contradicts the optimality of x as the eigenvector for $\lambda_{\min}(\mathbf{A})$. Furthermore, if $C < -10$, then $x_S^\top \mathbf{A} x_S < -10\epsilon$, which again contradicts the optimality of x . ■

Now recall that the total row and column contributions of i are defined as $\mathcal{R}_i(x) = \sum_{j \in [n] \setminus i} x_i \mathbf{A}_{i,j} x_j$ and $\mathcal{C}_i(x) = \sum_{j \in [n] \setminus i} x_j \mathbf{A}_{j,i} x_i$ respectively. In the remainder of the section, we simply write $\mathcal{R}_i = \mathcal{R}_i(x)$ and $\mathcal{C}_i = \mathcal{C}_i(x)$. We now define the contribution of i within the set $S \subset [n]$.

Definition 3.15. Let $S \subset [n]$ be as defined above. Then for any $i \in [n]$, define the row and column contributions of i within S as $\mathcal{R}_i^S = \sum_{j \in S \setminus i} x_i \mathbf{A}_{i,j} x_j$ and $\mathcal{C}_i^S = \sum_{j \in S \setminus i} x_j \mathbf{A}_{j,i} x_i$ respectively.

Observe from the above definition, we have $\sum_{i \in T_a} (\mathcal{R}_i^S + \mathcal{C}_i^S) = (x_S \mathbf{A} x_{T_a} + x_{T_a} \mathbf{A} x_S) \leq -\frac{\epsilon n}{10 \log(1/\epsilon)}$, where the inequality holds by definition of Case 1.

Proposition 3.16. We have $\sum_{i \in S} (\mathcal{R}_i^S + \mathcal{C}_i^S)^2 \leq 1601\epsilon n$.

Proof. Let $z^S, z, z^- \in \mathbb{R}^{|S|}$ be vectors defined for $i \in S$ as $z_i^S = \mathcal{R}_i^S + \mathcal{C}_i^S$, $z_i = \mathcal{R}_i + \mathcal{C}_i$, and $z^- = z - z^S$. Notice that our goal is to bound $\|z^S\|_2^2$, which by triangle inequality satisfies $\|z^S\|_2^2 \leq 2(\|z\|_2^2 + \|z^-\|_2^2)$. First note that

$$\begin{aligned} \|z^-\|_2^2 &= \sum_{i \in S} \left(\sum_{j \notin S} x_i \mathbf{A}_{i,j} x_j + \sum_{j \notin S} x_j \mathbf{A}_{j,i} x_i \right)^2 \\ &\leq 2 \sum_{i \in S} \left(\sum_{j \notin S} x_i \mathbf{A}_{i,j} x_j \right)^2 + 2 \sum_{i \in S} \left(\sum_{j \notin S} x_j \mathbf{A}_{j,i} x_i \right)^2 \end{aligned} \tag{10}$$

Using that $[n] \setminus S < \epsilon n/100$, we have by Proposition 3.13 that $\sum_{i \in S} \left(\sum_{j \notin S} x_i \mathbf{A}_{i,j} x_j \right)^2 \leq \epsilon n/100$, so $\|z^-\|_2^2 \leq \epsilon n/25$.

We now bound $\|z\|_1 = \sum_{i \in S} |\mathcal{R}_i + C_i|$. By Fact 3.6, we have $\mathcal{R}_i + C_i \leq 0$ for all $i \in [n]$, which means that $\|z\|_1 \leq \sum_{i \in [n]} |\mathcal{R}_i + C_i| = |2\langle x, \mathbf{A}x \rangle - 2 \sum_{i \in [n]} \mathbf{A}_{i,i}(x_i)^2| \leq 2\epsilon n$. Next, we bound $\|z\|_\infty$. Notice that $|\mathcal{R}_i + C_i| = 2|x_i \langle \mathbf{A}_{i,*}, x \rangle - \mathbf{A}_{i,i}(x_i)^2| \leq 2\epsilon n(x_i)^2 + 2\mathbf{A}_{i,i}(x_i)^2 < 4\epsilon n(x_i)^2$, using that $\mathbf{A}x = -\epsilon nx$, since x is an eigenvector of \mathbf{A} . Since for $i \in S$, it also follows that $(x_i)^2 \leq \frac{100}{\epsilon n}$, thus $\|z\|_\infty \leq 400$ as needed. It follows that $\|z\|_2^2$ is maximized by having $2\epsilon n/400$ coordinates equal to 400, giving $\|z\|_2^2 \leq 2\epsilon n/400(400)^2 = 800\epsilon n$. It follows then that $\|z^S\|_2^2 \leq 1601\epsilon n$ as needed. ■

Eigenvector Setup: We now define the “target” direction $y \in \mathbb{R}^n$ which we will use in our analysis for **Case 1**. First, we will need the following definitions. Let

$$D_a^p = \left\{ t \in T_a : -\frac{2^{p+1}2^a}{\log(1/\epsilon)} \leq \mathcal{R}_t^S + C_t^S \leq -\frac{2^p 2^a}{\log(1/\epsilon)} \right\}$$

Define the *fill* β of T_a as the value such that $\beta = 2^{-p}$ where $p \geq 1$ is the smallest value of p such that $-|D_a^p|(\frac{2^p 2^a}{\log(1/\epsilon)}) \leq -\frac{\epsilon n}{40 \log^2(1/\epsilon)}$. Note that at least one such p for $1 \leq p \leq \log(1/\epsilon)$ must exist. Let $T_a^* = D_a^p$ where $\beta = 2^{-p}$. Observe that $x_S \mathbf{A} x_{T_a^*} + x_{T_a^*} \mathbf{A} x_S \leq -\frac{\epsilon n}{40 \log^2(1/\epsilon)}$. Finally, we define our target “eigenvector” y as

$$y = x_S + \zeta \beta (2^{-a} \cdot x_{T_a}) \quad (11)$$

where $\zeta = \Theta(\log^2(1/\epsilon))$ is as above, and we also define our target submatrix subsampling size as $\lambda = \frac{2000\beta^2 \zeta^2 \log(1/\epsilon)}{2^a \epsilon}$. First, we prove that for a random submatrix $\mathbf{A}_{T \times T}$, where $i \in [n]$ is sampled and added to T with probability λ/n , we have that $y^\top \mathbf{A}_{T \times T} y$ is negative in expectation conditioned on $|T \cap T_a^*| \geq 1$.

Lemma 3.17. Suppose we are in case 1 with T_a contributing such that $2^a \geq 10^6 \zeta^3$. Let δ_i be an indicator variable that we sample coordinate i , with $\mathbb{E}[\delta_i] = \frac{\lambda}{n}$ and $\lambda = \frac{2000\beta^2 \zeta^2 \log(1/\epsilon)}{2^a \epsilon}$. Then if $y = x_S + \zeta \beta (2^{-a} x_{T_a})$ where $\zeta \geq 100 \log^2(1/\epsilon)$, and if $t \in T_a^*$, then

$$\mathbb{E} \left[\sum_{i,j \in S \cup \{t\}} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \mid \delta_t = 1 \right] \leq -\frac{50\zeta\lambda}{n}$$

Proof. First observe $\mathbb{E}[\sum_{i \in S} y_i^2 \mathbf{A}_{i,i} \delta_i] \leq \frac{\lambda}{n} \|x\|_2^2 \leq \frac{\lambda}{n}$ since x is a unit vector. Note that $y_S = x_S$ by construction, so we can use Proposition 3.14 to bound $|\langle y_S, \mathbf{A} y_S \rangle|$ by $10\epsilon n$, which gives

$$\begin{aligned} \mathbb{E} \left[\sum_{i,j \in S \cup \{t\}} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \mid \delta_t = 1 \right] &\leq \frac{\lambda^2}{n^2} \left(\langle y_S, \mathbf{A} y_S \rangle - \sum_{i \in S} y_i^2 \mathbf{A}_{i,i} \delta_i \right) + \frac{\lambda}{n} + \frac{\lambda}{n} \sum_{i \in S} (\mathbf{A}_{t,i} + \mathbf{A}_{i,t}) y_i y_t + y_t^2 \\ &\leq \frac{20\lambda^2 \epsilon}{n} + \frac{\lambda}{n} \left(1 + \frac{\zeta \beta}{2^a} (\mathcal{R}_t^S + C_t^S) \right) + \left(\frac{\zeta \beta}{2^a} \right)^2 \left(\frac{1002^a}{\epsilon n} \right) \\ &\leq \frac{20\lambda^2 \epsilon}{n} + \frac{\lambda}{n} \left(1 + \frac{\zeta \beta}{2^a} (\mathcal{R}_t^S + C_t^S) \right) + \frac{100\zeta^2 \beta^2}{2^a \epsilon n} \end{aligned} \quad (12)$$

Now by definition of $i \in T_a^*$, we have $(\mathcal{R}_t^S + C_t^S) \leq -\frac{2^a}{\beta \log(1/\epsilon)}$. Thus

$$\mathbb{E} \left[\sum_{i,j \in S \cup \{t\}} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \mid \delta_t = 1 \right] \leq \frac{20\lambda^2 \epsilon}{n} + \frac{\lambda}{n} \left(1 - \frac{\zeta}{\log(1/\epsilon)} \right) + \frac{100\zeta^2 \beta^2}{2^a \epsilon n} \quad (13)$$

Setting $\zeta > 100 \log^2(1/\epsilon)$, we first note that $\frac{\lambda}{n} \left(1 - \frac{\zeta}{\log(1/\epsilon)} \right) \leq -\frac{99\lambda\zeta}{100n \log(1/\epsilon)}$. Since

$$\lambda = \frac{2000\beta^2 \zeta^2 \log(1/\epsilon)}{2^a \epsilon} \leq \frac{2000\beta^2}{2^6 \zeta \epsilon} \leq \frac{1}{\zeta \epsilon}$$

it follows that $\frac{20\lambda^2 \epsilon}{n} \leq \frac{20\lambda}{n\zeta} \leq \frac{20\lambda}{n} < \frac{\lambda\zeta}{5 \log(1/\epsilon)n}$. Thus $\frac{10\lambda^2 \epsilon}{n} - \frac{99\lambda\zeta}{100n \log(1/\epsilon)} \leq -\frac{3\lambda\zeta}{4n \log(1/\epsilon)}$. So we can simply and write

$$\begin{aligned} \mathbb{E} \left[\sum_{i,j \in S \cup \{t\}} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \mid \delta_t = 1 \right] &\leq -\frac{3\lambda\zeta}{4n \log(1/\epsilon)} + \frac{100\zeta^2 \beta^2}{2^a \epsilon n} \\ &= -\frac{1500\beta^2 \zeta^3}{2^a \epsilon n} + \frac{100\zeta^2 \beta^2}{2^a \epsilon n} \\ &\leq -\frac{1400\beta^2 \zeta^3}{2^a \epsilon n} \\ &\leq -\frac{50\zeta \lambda}{n} \end{aligned} \quad (14)$$

as desired. ■

Lemma 3.18. *Let δ_i be an indicator variable with $\mathbb{E}[\delta_i] = \lambda/n$. Then*

$$\Pr \left[\left| \sum_{i,j \in S} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j - \mathbb{E} \left[\sum_{i,j \in S} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] \right| \geq C \frac{\lambda}{n} \right] \leq \frac{49}{50}$$

Where $C > 0$ is some constant.

Proof. We can apply Corollary 3.9, where we can set the values of c_1, c_2 to be bounded by constants by the results of Proposition 3.14 and 3.16, and by definition of the set S we can set $\alpha \leq \sqrt{\epsilon}$ for the α in Corollary 3.9, and using that $\lambda \leq O(1/\epsilon)$, we obtain:

$$\text{Var} \left[\sum_{i \neq j \in S} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] \leq \frac{C\lambda^2}{n^2}$$

for some constant C . Now note that $\mathbb{E}[\sum_{i \neq j \in S} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j] \leq \frac{30\lambda^2 \epsilon}{n} \leq O(\frac{\lambda}{n})$, thus by Chebyshev's, with probability 99/100, we have $|\sum_{i \neq j \in S} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j| \leq O(\frac{\lambda}{n})$. Moreover, note that $\sum_i \mathbf{A}_{i,i} y_i^2 \delta_i$ can be assumed to be a positive random variable using that $\mathbf{A}_{i,i} = 1$, and note the expectation of this variable is $\frac{\lambda}{n}$, and is at most $100\lambda/n$ with probability 99/100. Thus $|\sum_{i \in S} y_i^2 \mathbf{A}_{i,i} \delta_i - \mathbb{E}[\sum_{i \in S} y_i^2 \mathbf{A}_{i,i}]| \leq 100\lambda/n$. By a union bound, we have:

$$\Pr \left[\left| \sum_{i,j \in S} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right| \geq C \frac{\lambda}{n} \right] \leq \frac{49}{50}$$

Where $C = 150$. ■

Lemma 3.19. Fix any $t \in T_a^*$. Then

$$\Pr \left[\left| \sum_{i \in S} y_t (y_i \mathbf{A}_{i,t} + \mathbf{A}_{t,i} y_i) \delta_i - \mathbb{E} \left[\sum_{i \in S} y_t (y_i \mathbf{A}_{i,t} + \mathbf{A}_{t,i} y_i) \delta_i \right] \right| \geq \frac{10\lambda}{n} \right] \leq \frac{1}{100}$$

Proof. By independence of the δ_i 's

$$\begin{aligned} \text{Var} \left(\sum_{i \in S} (y_i \mathbf{A}_{i,t} y_t + y_t \mathbf{A}_{t,i} y_i) \delta_i \right) &\leq \frac{\lambda}{n} \sum_i (y_i \mathbf{A}_{i,t} y_t + y_t \mathbf{A}_{t,i} y_i)^2 \\ &\leq \frac{\lambda}{n} \sum_i 2y_t^2 y_i^2 \\ &\leq \frac{2\lambda}{n} (\zeta \beta 2^{-a})^2 \frac{1002^a}{\epsilon n} \\ &\leq \frac{2\lambda}{n} \left(\frac{100\zeta^2 \beta^2}{2^a \epsilon n} \right) \\ &\leq \frac{2\lambda^2}{5n^2} \end{aligned} \tag{15}$$

Now since $\mathbb{E}[\sum_{i \in S} (y_i \mathbf{A}_{i,t} y_t + y_t \mathbf{A}_{t,i} y_i) \delta_i] \leq -\frac{\lambda}{n} \frac{\zeta}{\log(1/\epsilon)} \leq -\frac{100\lambda}{n}$, the desired result follows by Chebyshev's inequality. ■

Lemma 3.20. Fix any $t \in T_a^*$. Then

$$\Pr \left[\sum_{i,j \in \text{SU}\{t\}} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \leq \frac{-25\zeta\lambda}{n} \mid \delta_t = 1 \right] \geq 24/25$$

Proof. Conditioned on $\delta_t = 1$, we have

$$\begin{aligned} &\left| \left[\sum_{i,j \in \text{SU}\{t\}} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] - \mathbb{E} \left[\sum_{i,j \in \text{SU}\{t\}} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \mid \delta_t = 1 \right] \right| \\ &\leq \left| \sum_{i \in S} y_t (y_i \mathbf{A}_{i,t} + \mathbf{A}_{t,i} y_i) \delta_i - \mathbb{E} \left[\sum_{i \in S} y_t (y_i \mathbf{A}_{i,t} + \mathbf{A}_{t,i} y_i) \delta_i \right] \right| \\ &\quad + \left| \sum_{i,j \in S} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j - \mathbb{E} \left[\sum_{i,j \in S} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \right] \right| \\ &\leq C \frac{\lambda}{n} \end{aligned} \tag{16}$$

for some constant $C \leq 200$, where the last fact follows from Lemmas 3.18 and 3.19 with probability $24/25$. Since $\mathbb{E} \left[\sum_{i,j \in \text{SU}\{t\}} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \mid \delta_t = 1 \right] \leq -\frac{50\zeta\lambda}{n}$ by Lemma 3.17, by scaling ζ by a sufficiently large constant, the result follows. ■

Theorem 3.21. Suppose we are in case 1 with T_a contributing such that $2^a \geq 10^6 \zeta^3$. Then there is an algorithm that queries at most $O(\frac{\log^7(1/\epsilon)}{\epsilon^2})$ entries of \mathbf{A} , and finds a principal submatrix of \mathbf{A} which is not PSD with probability at least $9/10$ in the NO case. The algorithm always returns YES on a YES instance.

Proof. By the above, we just need to sample a expected size $O(\lambda^2)$ submatrix from the conditional distribution of having sampled at least one entry from T_a^* . Since $|T_a^*| \geq \beta/10 \frac{\epsilon n}{2^a \log(1/\epsilon)}$, and since $\lambda = \Theta(\frac{\beta^2 \zeta^2 \log(1/\epsilon)}{2^a \epsilon})$, we see that this requires a total of k samples of expected size $O(\lambda^2)$, where

$$\begin{aligned} k &= (n/|T_a^*|)/\lambda \leq \left(\frac{2^a 10 \log(1/\epsilon)}{\beta \epsilon}\right) \left(\frac{2^a \epsilon}{\beta^2 \zeta^2 \log(1/\epsilon)}\right) \\ &\leq 10 \frac{2^{2a}}{\beta^3 \zeta^2} \end{aligned} \tag{17}$$

Thus the total complexity is $O(k\lambda^2)$, and we have

$$\begin{aligned} k\lambda^2 &\leq 10 \frac{2^{2a}}{\beta^3 \zeta^2} \left(\frac{\beta^4 \zeta^4 \log^2(1/\epsilon)}{2^{2a} \epsilon^2}\right) \\ &\leq 10 \frac{\beta \zeta^2 \log^2(1/\epsilon)}{\epsilon^2} \\ &= O\left(\frac{\zeta^2 \log^2(1/\epsilon)}{\epsilon^2}\right) \end{aligned} \tag{18}$$

we use the fact that we can set $\zeta = O(\log^2(1/\epsilon))$. Finally, note that we do not know β or 2^a , but we can guess the value of λ in powers of 2, which is at most $O(\frac{\zeta^2}{\epsilon^2})$, and then set k to be the value such that $k\lambda^2$ is within the above allowance. This blows up the complexity by a $\log(1/\epsilon)$ factor to do the guessing. ■

3.3.2 Case 2: Spread Negative Mass and Main Theorem

In the prior section, we saw that if the quadratic form $x^T \mathbf{A} x$ satisfies the condition for being in Case 1, we could obtain a $\tilde{O}(1/\epsilon^2)$ query algorithm for finding a principal submatrix $A_{T \times T}$ such that $y^T \mathbf{A}_{T \times T} y < 0$ for some vector y . Now recall that $S = \{i \in [n] : |x_i|^2 \leq \frac{1}{\epsilon n}\}$, and let $T_a = \{i \in [n] : \frac{1002^{a-1}}{\epsilon n} \leq |x_i|^2 \leq \frac{1002^a}{\epsilon n}\}$ for $a \geq 1$. Recall that the definition of Case 1 was that $x_S^T \mathbf{A} x_{T_a} + x_{T_a}^T \mathbf{A} x_S \leq -\epsilon n / (10 \log(1/\epsilon))$ for some $2^a \geq 10^6 \zeta^3$. In this section, we demonstrate that if this condition does not hold, then we will also obtain a $\tilde{O}(1/\epsilon^2)$ query algorithm for the problem.

Thus, suppose now that we are in Case 2; namely that $x_S \mathbf{A} x_{T_a} + x_{T_a} \mathbf{A} x_S > -\epsilon n / (10 \log(1/\epsilon))$ for all $2^a \geq 10^6 \zeta^3$. Now let $T^+ = \cup_{2^a > 10^6 \zeta^3} T_a$ and let $T^- = \cup_{2^a \leq 10^6 \zeta^3} T_a$. Let $S^* = S \cup T^-$. We now observe an important fact, which states that if we are not in Case 1, then $x_{S^*} \mathbf{A} x_{S^*}$ contributes a substantial fraction of the negativeness in the quadratic form.

Fact 3.22. *Suppose we are in Case 2: meaning that $x_S^T \mathbf{A} x_{T_a} + x_{T_a} \mathbf{A} x_S > -\epsilon n / (10 \log(1/\epsilon))$ for all $2^a \geq 10^6 \zeta^3$. Then we have $x_{S^*}^T \mathbf{A} x_{S^*} \leq -\epsilon n / 2$.*

Proof. Notice that this implies that $x_S^T \mathbf{A} x_{T^+} + x_{T^+} \mathbf{A} x_S \geq -\epsilon n / 10$, since there are at most $\log(1/\epsilon)$ level sets included in T^+ by Proposition 3.4. Note since the contribution of $|x_{T^+}^T \mathbf{A} x_{T^+}| \leq -\epsilon n 10^{-6} / \zeta^3$ and $|x_{T^-} \mathbf{A} x_{T^+} + x_{T^+} \mathbf{A} x_{T^-}| \leq \sqrt{|T^-| |T^+|} \leq \epsilon n / 100$ by Proposition 3.12. Thus if $x^T \mathbf{A} x \leq -\epsilon n$ to begin

with, it follows that we must have

$$\begin{aligned}
x_{S^*}^\top \mathbf{A} x_{S^*} &\leq x^\top \mathbf{A} x - ((x_S^\top \mathbf{A} x_{T^+} + x_{T^+}^\top \mathbf{A} x_S) - (x_{T^+}^\top \mathbf{A} x_{T^+}) - (x_{T^-}^\top \mathbf{A} x_{T^+} + x_{T^+}^\top \mathbf{A} x_{T^-})) \\
&\leq -\epsilon n + \epsilon n/10 + \epsilon n 10^{-6}/\zeta^3 \epsilon n/100 \\
&< -\epsilon n/2
\end{aligned} \tag{19}$$

■

We now proceed by analyzing the result of sampling a principal submatrix from the quadratic form $x_{S^*}^\top \mathbf{A} x_{S^*}$, which by the prior fact is already sufficiently negative. Specifically, we will demonstrate that the variance of the standard estimator from Lemma 3.8, and specifically Corollary 3.9, is already sufficiently small to allow for a single randomly chosen $O(1/\epsilon) \times O(1/\epsilon)$ principal submatrix of \mathbf{A} to have negative quadratic form with x_{S^*} with good probability. In order to place a bound on the variance of this estimator and apply Corollary 3.9, we will need to bound the row and column contributions of the quadratic form $x_{S^*}^\top \mathbf{A}_{S^* \times S^*} x_{S^*}$, which we now formally define.

Definition 3.23. For $i \in [n]$, define the row and column contributions of i within S^* as $\mathcal{R}_i^* = \sum_{j \in S^* \setminus i} x_i \mathbf{A}_{i,j} x_j$ and $C_i^* = \sum_{j \in S^* \setminus i} x_j \mathbf{A}_{j,i} x_i$ respectively.

Recall that the *total* row and column contributions of i are defined via $\mathcal{R}_i = \sum_{j \in [n] \setminus i} x_i \mathbf{A}_{i,j} x_j$ and $C_i = \sum_{j \in [n] \setminus i} x_j \mathbf{A}_{j,i} x_i$ respectively, and recall that we have $\mathcal{R}_i + C_i \leq 0$ for all $i \in [n]$ by Fact 3.6

Proposition 3.24. We have $\sum_{i \in S^*} (\mathcal{R}_i^* + C_i^*)^2 \leq 10^9 \cdot \zeta^3 \epsilon n$.

Proof. The proof proceeds similarly to Proposition 3.16. Let $z^*, z, z^- \in \mathbb{R}^{|S^*|}$ be defined for $i \in S^*$ via $z_i^* = \mathcal{R}_i^* + C_i^*$, $z_i = \mathcal{R}_i + C_i$, and $z^- = z - z^*$. Notice that our goal is to bound $\|z^*\|_2^2$, which by triangle inequality satisfies $\|z^*\|_2^2 \leq 2(\|z\|_2^2 + \|z^-\|_2^2)$. First note that

$$\begin{aligned}
\|z^-\|_2^2 &= \sum_{i \in S^*} \left(\sum_{j \notin S^*} x_i \mathbf{A}_{i,j} x_j + \sum_{j \notin S^*} x_j \mathbf{A}_{j,i} x_i \right)^2 \\
&\leq 2 \sum_{i \in S^*} \left(\sum_{j \notin S^*} x_i \mathbf{A}_{i,j} x_j \right)^2 + 2 \sum_{i \in S^*} \left(\sum_{j \notin S^*} x_j \mathbf{A}_{j,i} x_i \right)^2
\end{aligned} \tag{20}$$

Using that $|[n] \setminus S^*| < \epsilon n/100$, we have by Proposition 3.13 that $\sum_{i \in S^*} \left(\sum_{j \notin S^*} x_i \mathbf{A}_{i,j} x_j \right)^2 \leq \epsilon n/100$, so $\|z^-\|_2^2 \leq \epsilon n/25$.

We now bound $\|z\|_1 = \sum_{i \in S} |\mathcal{R}_i + C_i|$. Recall that we have $\mathcal{R}_i + C_i \leq 0$ for all $i \in [n]$, which means that $\|z\|_1 \leq \sum_{i \in [n]} |\mathcal{R}_i + C_i| = |2\langle x, \mathbf{A}x \rangle - 2 \sum_{i \in [n]} \mathbf{A}_{i,i} (x_i)^2| \leq 2\epsilon n$. Next, we bound $\|z\|_\infty$. Notice that $|\mathcal{R}_i + C_i| = 2|x_i \mathbf{A}_{i,*} x| = 2\epsilon n (x_i)^2 - 2\mathbf{A}_{i,i} (x_i)^2 < 4\epsilon n (x_i)^2$, using that $\mathbf{A}x = -\epsilon n x$, since x is an eigenvector of \mathbf{A} . Since $i \in S^*$, by definition we have $(x_i)^2 \leq \frac{100 \cdot 10^6 \cdot \zeta^3}{\epsilon n}$, thus $\|z\|_\infty \leq 100 \cdot 10^6 \cdot \zeta^3$. It follows that $\|z\|_2^2$ is maximized by having $2\epsilon n / (10^8 \cdot \zeta^3)$ coordinates equal to $10^8 \cdot \zeta^3$, giving $\|z\|_2^2 \leq 2\epsilon n / (10^8 \cdot \zeta^3) (10^8 \cdot \zeta^3)^2 = 2 \cdot 10^8 \cdot \zeta^3 \epsilon n$. It follows then that $\|z\|_2^2 \leq \cdot 10^9 \cdot \zeta^3 \epsilon n$ as needed. ■

Theorem 3.25. There is an algorithm which, given \mathbf{A} with $\|\mathbf{A}\|_\infty \leq 1$ such that either $x^\top \mathbf{A} x \geq 0$ for all $x \in \mathbb{R}^n$ (YES Case), or $x^\top \mathbf{A} x \leq -\epsilon n$ for some $x \in \mathbb{R}^n$ with $\|x\|_2 \leq 1$ (NO Case), distinguishes the two cases

with probability $3/4$ using at most $\tilde{O}(\frac{1}{\epsilon^2})$ queries, and running in time $\tilde{O}(1/\epsilon^\omega)$, where $\omega < 2.373$ is the exponent of fast matrix multiplication. Moreover, in the YES case the, the algorithm always outputs YES (with probability 1), and in the NO case, the algorithm returns a certificate in the form of a principal submatrix which is not PSD.

Proof. By Theorem 3.21 which handles Case 1, we can restrict ourselves to Case 2. Using Fact 3.22 as well as Proposition 3.24, can apply Corollary 3.9 with the vector $y = x_{S^*}$, setting $c_1 = \Theta(1)$ and $c_2 = \Theta(\zeta^3)$, and $\alpha = O(\sqrt{\epsilon}\zeta^3) = O(\sqrt{\epsilon}\log^6(1/\epsilon))$, to obtain that

$$\text{Var}\left[\sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j\right] \leq O(\log^{12}(1/\epsilon) \frac{k^2}{n^2})$$

where $k = \tilde{O}(1/\epsilon)$. Since by Proposition 3.7 and Fact 3.22, we have $\mathbb{E}[\sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j] \leq \frac{k^2}{4n^2} \langle x_{S^*}, \mathbf{A} x_{S^*} \rangle \leq -\frac{\epsilon k^2}{8n}$, it follows that by repeating the sampling procedure $O(\log^{12}(1/\epsilon))$, by Chebyshev's we will have that at least one sample satisfies $\sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \leq -\frac{\epsilon k^2}{4n}$ with probability $99/100$.

Now note that this random variable does not take into account the diagonal. Thus, it will suffice to bound the contribution of the random variable $\sum_{i \in [n]} \delta_i \mathbf{A}_{i,i} (y_i)^2 \tilde{O}((1/\epsilon)/n)$. First observe that $\mathbb{E}[\sum_{i \in [n]} \delta_i \mathbf{A}_{i,i} (y_i)^2] = \frac{k}{n}$. The proof proceeds by a simple bucketing argument; let $\Lambda_i = \{i \in S^* \mid \frac{2^i}{n} \leq (y_i)^2 \leq \frac{2^{i+1}}{n}\}$, and for a single $k \times k$ sampled submatrix, let $T \subset [n]$ be the rows and columns that are sample. Note that $\mathbb{E}[|T \cap \Lambda_i|] \leq k 2^{-i}$, since $|\Lambda_i| \leq 2^{-i}$. Note also that $|\Lambda_i| = 0$ for every i such that $2^i \geq \frac{100^8 \zeta^3}{\epsilon}$ by definition of S^* and the fact that y is zero outside of S^* . Then by Chernoff bounds we have that with probability $\Pr[|T \cap \Lambda_i| > \log(1/\epsilon) \max\{k 2^{-i}, 1\}] \leq 1 - \frac{\epsilon^{10}}{C}$ for some constant C for our choosing. We can then union bound over all $O(\log(1/\epsilon))$ sets Λ_i , to obtain

$$\sum_{i \in [n]} \delta_i \mathbf{A}_{i,i} (y_i)^2 \leq \sum_{i: 2^i \leq \frac{100^8 \zeta^3}{\epsilon}} \frac{2^{i+1}}{n} |T \cap \Lambda_i| \leq \sum_{i: 2^i \leq \frac{100^8 \zeta^3}{\epsilon}} \frac{2}{n} \log(1/\epsilon) \max\{k, 2^i\}$$

with probability at least $1 - \frac{\epsilon^9}{C}$. Setting $k = \Theta(\log^6(1/\epsilon)/\epsilon)$, we have that $\sum_{i \in [n]} \delta_i \mathbf{A}_{i,i} (y_i)^2 \leq \sum_{i: 2^i \leq \frac{100^8 \zeta^3}{\epsilon}} (2/n) \log(1/\epsilon) k = O(\log^2(1/\epsilon) k/n)$. Thus we can condition on $\sum_{i \in [n]} \delta_i \mathbf{A}_{i,i} (y_i)^2 = O(\log^2(1/\epsilon) k/n)$ for all $\tilde{O}(1)$ repetitions of sampling a submatrix. Since at least one sampled submatrix satisfied $\sum_{i \neq j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j \leq -\frac{\epsilon k^2}{4n}$, and since $k = \Theta(\log^6(1/\epsilon)/\epsilon)$, this demonstrates that at least one sampled submatrix will satisfy $\sum_{i,j} y_i \mathbf{A}_{i,j} y_j \delta_i \delta_j < -\frac{\epsilon k^2}{8n}$ as needed in the NO instance. The resulting query complexity is then $O(\log^2(1/\epsilon) k^2) = O(\frac{\log^{24}(1/\epsilon)}{\epsilon^2}) = \tilde{O}(\frac{1}{\epsilon^2})$ as desired. Finally, for runtime, notice that the main computation is computing the eigenvalues of a $k \times k$ principal submatrix, for $k = \tilde{O}(1/\epsilon)$, which can be carried out in time $\tilde{O}(1/\epsilon^\omega)$ [DDHK07, BVKS19]. ■

4 PSD Testing with ℓ_2^2 Gap

Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a symmetric matrix with eigenvalues $\lambda_{\max} = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n = \lambda_{\min}$. In this section, we consider the problem of testing positive semi-definiteness with an ℓ_2^2 gap. Formally, the problem statement is as follows.

Definition 4.1 (PSD Testing with ℓ_2^2 -Gap.). Fix, $\epsilon \in (0, 1]$ and let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a symmetric matrix satisfying $\|\mathbf{A}\|_\infty \leq 1$, with the promise that either

- **YES Instance:** \mathbf{A} is PSD.

- **NO Instance:** \mathbf{A} is ϵ -far from PSD in ℓ_2^2 , meaning that $\min_{\mathbf{B} \geq 0} \|\mathbf{A} - \mathbf{B}\|_F^2 = \sum_{i: \lambda_i < 0} \lambda_i^2 = \epsilon n^2$.

The PSD Testing problem with ℓ_2^2 -gap is to design an algorithm which distinguish these two cases with probability at least $2/3$, using the minimum number of queries possible to the entries of \mathbf{A} .

Our algorithm for this problem will query a principal submatrix $\mathbf{A}_{S \times S}$ and return PSD if $\mathbf{A}_{S \times S}$ is PSD, otherwise it will return not PSD. Since all principal submatrices of PSD matrices are PSD, we only need show that if \mathbf{A} is ϵ -far from PSD, then we can find a non-PSD principal submatrix with small size. Note again that this implies that our algorithm will be one-sided. Thus, in the following, we can focus on the case where \mathbf{A} is ϵ -far from PSD. We begin by stating two fundamental observations, which, along with an application of our algorithm from Section 3, will allow us to reduce the problem of PSD testing with ℓ_2 gap to the problem of testing certain functions of the singular values of \mathbf{A} .

Proposition 4.2 (PSD matrices are top heavy). *Fix any $n \in \mathbb{N}$, $1 \leq k \leq n$, and $\mathbf{D} \in \mathbb{R}^{n \times n}$. Then if \mathbf{D} is PSD, we have*

$$\sum_{i > k} \sigma_i(\mathbf{D})^2 \leq \frac{1}{k} (\text{Tr}(\mathbf{D}))^2$$

In particular, if \mathbf{D} has bounded entries $\|\mathbf{D}\|_\infty \leq 1$, we have $\sum_{i > k} \sigma_i(\mathbf{D})^2 \leq \frac{1}{k} n^2$.

Proof. We first show that $\sigma_k(\mathbf{D}) \leq k^{-1} \text{Tr}(\mathbf{D})$. To see this, suppose $\sigma_k(\mathbf{D}) > k^{-1} \text{Tr}(\mathbf{D})$. Then because \mathbf{D} is PSD, we would have $\sum_i \sigma_i = \sum_i \lambda_i = \text{Tr}(\mathbf{A}) > k \cdot k^{-1} \text{Tr}(\mathbf{D})$, a contradiction. Thus, $\sigma_i(\mathbf{D}) \leq k^{-1} \text{Tr}(\mathbf{D})$ for all $i \geq k$. Using this and the bound $\sum_{i > k} \sigma_i(\mathbf{D}) \leq \text{Tr}(\mathbf{D})$, it follows that the quantity $\sum_{i > k} \sigma_i(\mathbf{D})^2$ is maximized by having k singular values equal to $\text{Tr}(\mathbf{D})/k$, yielding $\sum_{i > k} \sigma_i(\mathbf{D})^2 \leq k \cdot (\text{Tr}(\mathbf{D})/k)^2 = k^{-1} (\text{Tr}(\mathbf{D}))^2$ as needed. ■

Proposition 4.3. *Let $\mathbf{D} \in \mathbb{R}^{n \times n}$ be a symmetric matrix such that $\|\mathbf{D}\|_\infty \leq 1$, and let $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$ be its singular values. Suppose \mathbf{D} is at least ϵ -far in L_2 from PSD, so that $\sum_{i: \lambda_i(\mathbf{D}) < 0} \lambda_i^2(\mathbf{D}) \geq \epsilon n^2$, and suppose further that $\min_i \lambda_i(\mathbf{D}) > -\frac{1}{2k} n$ for any $k \geq \frac{2}{\epsilon}$. Then we have*

$$\sum_{i > k} \sigma_i^2(\mathbf{D}) > \frac{\epsilon}{2} n^2$$

Proof. Let $W \subseteq [n]$ be the set of values $i \in [n]$ such that $\lambda_i < 0$. Let $W' \subseteq [n]$ be the set of values $i \in [n]$ such that $\sigma_i < \frac{1}{2k} n$. By assumption: $\sum_{i \in W'} \sigma_i^2 \geq \sum_{i \in W} \lambda_i^2 \geq \epsilon n^2$. Now $\sum_{i \in W'} \sigma_i^2 = \sum_{i \in W', i \leq k} \sigma_i^2 + \sum_{i \in W', i > k} \sigma_i^2$, so the fact that $|\sigma_i| \leq (1/2k)n$ for every $i \in W'$, we have that $\sum_{i \in W', i \leq k} \sigma_i^2 \leq k(n/(2k))^2 = n^2/4k < \epsilon n^2/2$. Thus we must have $\sum_{i \in W', i > k} \sigma_i^2 > \epsilon n^2/2$, giving

$$\begin{aligned} \sum_{i > k} \sigma_i^2 &\geq \sum_{i \in W', i > k} \sigma_i^2 \\ &> \epsilon n^2/2 \end{aligned} \tag{21}$$

as required. ■

4.1 Analysis of the Algorithm

Our analysis will require several tools, beginning with the following interlacing lemma.

Lemma 4.4 (Dual Lidskii Inequality, [Tao11] Chapter 1.3). *Let $\mathbf{M}_1, \mathbf{M}_2$ be $t \times t$ symmetric Matrices, and fix $1 \leq i_1 < i_2 < \dots < i_k \leq n$. Then we have*

$$\sum_{j=1}^k \lambda_{i_j}(\mathbf{M}_1 + \mathbf{M}_2) \geq \sum_{j=1}^k \lambda_{i_j}(\mathbf{M}_1) + \sum_{j=1}^k \lambda_{n-j+1}(\mathbf{M}_2)$$

We will also need the following result of Rudelson and Vershynin [RV07] on the decay of spectral norms of random submatrices.

Proposition 4.5 ([RV07]). *Let $\mathbf{A} \in \mathbb{R}^{n \times m}$ be a rank r matrix with maximum Euclidean row norm bounded by M , in other words $\max_i |(\mathbf{A}\mathbf{A}^\top)_{i,i}| \leq M$. Let $Q \subset [n]$ be a random subset of rows of \mathbf{A} with expected cardinality q . Then there is a fixed constant $\kappa \geq 1$ such that*

$$\mathbb{E}[\|\mathbf{A}_{Q \times [m]}\|_2] \leq \kappa(\sqrt{\delta}\|\mathbf{A}\|_2 + \sqrt{\log q}M)$$

Finally, we will need a generalized Matrix Chernoff bound for the interior eigenvalues of sums of random matrices, which was derived by Gittens and Tropp [GT11].

Theorem 4.6 (Interior Eigenvalue Matrix Chernoff, Theorem 4.1 of [GT11]). *Consider a finite sequence $\{\mathbf{X}_j\}$ of independent, random, positive-semidefinite matrices with dimension m , and assume that $\|\mathbf{X}_j\|_2 \leq L$ for some value L almost surely. Given an integer $k \leq n$, define*

$$\mu_k = \lambda_k \left(\sum_j \mathbb{E}[\mathbf{X}_j] \right)$$

then we have the tail inequalities

$$\begin{cases} \Pr \left[\lambda_k(\sum_j \mathbf{X}_j) \geq (1 + \delta)\mu_k \right] \leq (n - k + 1) \cdot \left[\frac{e^\delta}{(1 + \delta)^{1 + \delta}} \right]^{\mu_k/L} & \text{for } \delta > 0 \\ \Pr \left[\lambda_k(\sum_j \mathbf{X}_j) \leq (1 - \delta)\mu_k \right] \leq k \cdot \left[\frac{e^{-\delta}}{(1 - \delta)^{1 - \delta}} \right]^{\mu_k/L} & \text{for } \delta \in [0, 1) \end{cases} \quad (22)$$

The Algorithm. Our first step is to run the ℓ_∞ -gap algorithm of Section 3 with $\epsilon_0 = \frac{2}{k}$, where we set $k = \frac{2 \cdot 400^2 \kappa^4}{\epsilon}$, where $\kappa \geq 1$ is the constant in Proposition 4.5. This allows us to assume that $\lambda_i \geq -\epsilon_0 n / 1000 \geq -\frac{1}{2k}n$ for all i , otherwise we have a $\tilde{O}(1/\epsilon^2)$ -query algorithm from the Section 3, and since our target complexity is $\tilde{O}(1/\epsilon^4)$, we can safely disregard the cost of running this algorithm in parallel. We begin by demonstrating that the Frobenius norm of $\mathbf{S}\mathbf{A}$ is preserved (up to scaling), where \mathbf{S} is a random row sampling matrix with sufficiently many rows.

Proposition 4.7. Let $\mathbf{M} \in \mathbb{R}^{m \times m}$. Fix $t \geq 1$ and let \mathbf{S} be a row sampling matrix which samples each row of \mathbf{M} with probability $p = \frac{t}{m}$, and let $\mathbf{S} \in \mathbb{R}^{t_0 \times m}$ be a row sampling matrix drawn from this distribution, where $\mathbb{E}[t_0] = t$. Then we have

$$\mathbb{E}\left[\frac{1}{p}\text{Tr}(\mathbf{S}\mathbf{M}\mathbf{S}^\top)\right] = \sum_i \lambda_i(\mathbf{M}) = \text{Tr}(\mathbf{M})$$

and

$$\text{Var}\left(\frac{1}{p}\text{Tr}(\mathbf{S}\mathbf{M}\mathbf{S}^\top)\right) \leq \frac{m}{t} \sum_i \mathbf{M}_{i,i}^2$$

Proof. For $i \in [m]$, let $\delta_i \in \{0,1\}$ indicate that we sample row i . We have $\mathbb{E}[\text{Tr}(\mathbf{S}\mathbf{M}\mathbf{S}^\top)] = \frac{1}{p}\mathbb{E}[\sum_{i=1}^n \delta_i \mathbf{M}_{i,i}] = \text{Tr}(\mathbf{M})$. Moreover,

$$\begin{aligned} \text{Var}\left(\frac{1}{p}\text{Tr}(\mathbf{S}\mathbf{M}\mathbf{S}^\top)\right) &\leq \frac{1}{p^2} \sum_{i=1}^n \delta_i \mathbf{M}_{i,i} - (\text{Tr}(\mathbf{M}))^2 \\ &\leq \sum_{i \neq j} \mathbf{M}_{i,i} \mathbf{M}_{j,j} + \frac{1}{p} \sum_i \mathbf{M}_{i,i}^2 - (\text{Tr}(\mathbf{M}))^2 \\ &\leq \frac{1}{p} \sum_i \mathbf{M}_{i,i}^2 \end{aligned} \tag{23}$$

as stated. ■

We now fix $t = \Theta(\log(1/\epsilon)/\epsilon^2)$, and draw row independent sampling matrices \mathbf{S}, \mathbf{T} with an expected t rows. Let $S, T \subset [n]$ be the rows and columns sampled by $\mathbf{S}, \mathbf{T}^\top$ respectively. We then compute $\mathbf{Z} = \mathbf{S}\mathbf{A}\mathbf{T}^\top$ with an expected $O(t^2)$ queries. Finally, we query the principal submatrix $\mathbf{A}_{(S \cup T) \times (S \cup T)}$, and test whether $\mathbf{A}_{(S \cup T) \times (S \cup T)}$ is PSD. Clearly if \mathbf{A} is PSD, so is $\mathbf{A}_{(S \cup T) \times (S \cup T)}$, so it suffices to analyze the NO case, which we do in the remainder.

Lemma 4.8. Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be ϵ -far from PSD with $\|\mathbf{A}\|_\infty \leq 1$. Then let $\mathbf{Z} = \mathbf{S}\mathbf{A}\mathbf{T}^\top$ be samples as described above, so that \mathbf{Z} has an expected $t = \Theta(\log(1/\epsilon)/\epsilon^2)$ rows and columns, where t is scaled by a larger enough constant, and let $k = \frac{2 \cdot 400^2 \kappa^4}{\epsilon}$, where $\kappa \geq 1$ is the constant in Proposition 4.5. Suppose further that $\sigma_{k+1}(\mathbf{A}) \leq 10n/k$. Then with probability $19/20$, we have

$$\frac{n^2}{t^2} \sum_{i>k} \sigma_i^2(\mathbf{Z}) > \epsilon n^2 / 16$$

Proof. Now write $\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}^\top$, $\mathbf{A}_k = \mathbf{U}\mathbf{\Lambda}_k\mathbf{V}^\top$, $\mathbf{A}_{-k} = \mathbf{U}\mathbf{\Lambda}_{-k}\mathbf{V}^\top$. Then $\mathbf{A} = \mathbf{A}_k + \mathbf{A}_{-k}$, and the rows of \mathbf{A}_k are orthogonal to the rows of \mathbf{A}_{-k} . Note that this implies that $\|\mathbf{A}_{i,*}\|_2^2 = \|(\mathbf{A}_k)_{i,*}\|_2^2 + \|(\mathbf{A}_{-k})_{i,*}\|_2^2$ for each $i \in [n]$ by the Pythagorean theorem, and since $\|\mathbf{A}\|_\infty \leq 1$ we have $\|(\mathbf{A}_{-k})_{i,*}\|_2^2 \leq n$.

Now set $\mathbf{M}_1 = \mathbf{S}\mathbf{A}_k\mathbf{A}_k^\top\mathbf{S}^\top$, and $\mathbf{M}_2 = \mathbf{S}\mathbf{A}_{-k}\mathbf{A}_{-k}^\top\mathbf{S}^\top$. Notice that $\mathbf{M}_1 + \mathbf{M}_2 = \mathbf{S}(\mathbf{A}_k\mathbf{A}_k^\top + \mathbf{A}_{-k}\mathbf{A}_{-k}^\top)\mathbf{S}^\top = \mathbf{S}\mathbf{A}\mathbf{A}^\top\mathbf{S}^\top$, using the fact that the rows and columns of \mathbf{A}_k are orthogonal to the rows and columns (respectively) of \mathbf{A}_{-k} . Let $p = \frac{t}{n}$ be the row sampling probability. Now suppose $\|(\mathbf{A}_{-k})\|_F^2 = \alpha n^2$. Note that we have shown that $\alpha > \epsilon/2$. By Proposition 4.7, we have $\mathbb{E}[\text{Tr}(\mathbf{M}_2)/p] = \sum_{i>k} \sigma_i^2(\mathbf{A}) = \alpha n^2 > \epsilon n^2 / 2$ for some $\alpha \geq \epsilon/2$, where the last inequality follows from Proposition 4.3. Moreover, we have

$$\begin{aligned}
\text{Var}\left(\frac{1}{p}\text{Tr}(\mathbf{M}_2)\right) &\leq \frac{1}{p} \sum_i (\mathbf{M}_2)_{i,i}^2 \\
&= \frac{1}{p} \sum_i \|(\mathbf{A}_{-k})_{i,*}\|_2^4
\end{aligned} \tag{24}$$

It follows that since each row satisfies $\|(\mathbf{A}_{-k})_{i,*}\|_2^2 \leq n$ and $\|(\mathbf{A}_{-k})\|_F^2 = \alpha n^2$, the quantity $\sum_i \|(\mathbf{A}_{-k})_{i,*}\|_2^4$ is maximized having αn rows with squared norm equal to n . This yields

$$\begin{aligned}
\text{Var}\left(\frac{1}{p}\text{Tr}(\mathbf{M}_2)\right) &\leq \frac{1}{p} \sum_i 2\alpha n \cdot n^2 \\
&\leq 2 \frac{\alpha n^4}{t} \\
&\leq \frac{\alpha^2}{100^2} n^4
\end{aligned} \tag{25}$$

Where in the last line, we used that $t > \frac{4 \cdot 100^2}{\epsilon_0} \geq \frac{2 \cdot 100^2}{\alpha}$. Then by Chebyshev's inequality, with probability 99/100, we have $\frac{1}{p}\text{Tr}(\mathbf{M}_2) > \alpha n^2 - (\alpha/10)n^2 = (9/10)\alpha n^2 \geq (9/20)\epsilon n^2$. Call this event \mathcal{E}_1 , and condition on it now. Next, by Proposition 4.5, since $\sigma_{k+1}(\mathbf{A}) \leq 10n/k$ we have $\mathbb{E}[\|\mathbf{SA}_{-k}\|_2] \leq \kappa(10\sqrt{tn}/k + \sqrt{2\log(1/\epsilon)}\sqrt{n}) < 20\kappa\sqrt{tn}/k$. Then by Markovs, we have $\|\mathbf{SA}_{-k}\|_2^2 = \|\mathbf{M}_2\|_2 \leq 200^2 \kappa^2 tn/k^2$ with probability 99/100, which we condition on now, and call this event \mathcal{E}_2 . Then by the Dual Lidskii inequality 4.4, we have

$$\begin{aligned}
\frac{1}{p} \sum_{j>k} \lambda_j(\mathbf{M}_1 + \mathbf{M}_2) &\geq \frac{1}{p} \left(\sum_{j>k} \lambda_j(\mathbf{M}_2) \right) \\
&\geq \frac{1}{p} (\text{Tr}(\mathbf{M}_2) - k\|\mathbf{M}_2\|_2) \\
&\geq (9/20)\epsilon n^2 - 200^2 \kappa^2 n^2/k \\
&\geq \epsilon n^2/4
\end{aligned} \tag{26}$$

using that $k > \frac{2 \cdot 400^2 \kappa^4}{\epsilon}$. Now let $\mathbf{W} = \frac{1}{\sqrt{p}}(\mathbf{SA})^\top$, and note that we took the transpose, so \mathbf{W} has n rows and t_1 columns, where $\mathbb{E}[t_1] = t$. Now by Chernoff bounds, with probability 99/100 we have $t_2 \leq 2t$; call this event \mathcal{E}_3 and condition on it now. The above demonstrates that $\frac{1}{p} \sum_{j>k} \lambda_j(\mathbf{M}_1 + \mathbf{M}_2) = \sum_{j>k} \lambda_j^2(\mathbf{W}) \leq \epsilon n^2/4$. Now note that $\sigma_{k+1}(\mathbf{W}) = \frac{1}{\sqrt{p}}(\sigma_{k+1}(\mathbf{SA}_k + \mathbf{SA}_{-k})) < \frac{1}{\sqrt{p}}\|\mathbf{SA}_{-k}\|_2 \leq 200\kappa n/k$, where we used the Weyl inequality for singular values: namely that for any two matrices \mathbf{A}, \mathbf{B} and value i , $|\sigma_i(\mathbf{A} + \mathbf{B}) - \sigma_i(\mathbf{A})| \leq \|\mathbf{B}\|_2$, and using that \mathbf{SA}_k is rank at most k , so $\sigma_{k+1}(\mathbf{SA}_k) = 0$.

Now draw a random row sampling matrix \mathbf{T} with an expected t rows, and write $\mathbf{N}_1 = \mathbf{TW}_k \mathbf{W}_k^\top \mathbf{T}$ and $\mathbf{N}_2 = \mathbf{TW}_{-k} \mathbf{W}_{-k}^\top \mathbf{T}$, and note again that $\mathbf{N}_1 + \mathbf{N}_2 = \mathbf{TW} \mathbf{W}^\top \mathbf{T}$. Moreover, the rows of \mathbf{W}_k live in a subspace orthogonal to the rows of \mathbf{W}_{-k} , so again by the Pythagorean theorem and boundedness of the entries in \mathbf{A} , we have $\|(\mathbf{W}_{-k})_{i,*}\|_2^2 \leq \frac{1}{p} t_1 \leq 2n$ for all $i \in [n]$. Then by Proposition 4.7, we have

$\mathbb{E}[\text{Tr}(\mathbf{N}_2)/p] = \|\mathbf{W}_{-k}\|_F^2 = \alpha n^2 \geq \epsilon n^2/4$, and

$$\begin{aligned} \text{Var}\left(\frac{1}{p}\text{Tr}(\mathbf{N}_2)\right) &\leq \frac{1}{p} \sum_{i=1}^n \|(\mathbf{W}_{-k})_{i,*}\|_2^4 \\ &\leq \frac{1}{p} n^3 \\ &\leq \frac{1}{t} n^4 \\ &\leq \frac{\epsilon^2}{100^2} n^4 \end{aligned} \tag{27}$$

Then by Chebyshev's inequality, with probability 99/100, we have $\frac{1}{p}\text{Tr}(\mathbf{N}_2) > \epsilon n^2/4 - (\epsilon/10)n^2 = \epsilon n^2/8$. Call this event \mathcal{E}_4 , and condition on it now. Now as shown above, we have $\|\mathbf{W}_{-k}\|_2 \leq 200\kappa n/k$, thus by Proposition 4.5 we have $\mathbb{E}[\|\mathbf{T}\mathbf{W}_{-k}\|_2] \leq \kappa(200\kappa\sqrt{tn}/k + 4\sqrt{\log(1/\epsilon)}\sqrt{n}) \leq 400\kappa^2\sqrt{tn}$, again where we take $t = \Theta(\frac{\log(1/\epsilon)}{\epsilon^2})$ with a large enough constant. Then by Markov's inequality, with probability 99/100 we have $\|\mathbf{N}_2\|_2 \leq 400^2\kappa^4 n^2/k^2$, and again by the Dual Lidskii inequality 4.4, we have

$$\begin{aligned} \frac{1}{p} \sum_{j>k} \lambda_j(\mathbf{N}_1 + \mathbf{N}_2) &\geq \frac{1}{p} \left(\sum_{j>k} \lambda_j(\mathbf{N}_2) \right) \\ &\geq \frac{1}{p} (\text{Tr}(\mathbf{N}_2) - k\|\mathbf{N}_2\|_2) \\ &\geq \epsilon n^2/8 - 400^2\kappa^4 n^2/k \\ &\geq \epsilon n^2/16 \end{aligned} \tag{28}$$

Using that $k \geq \frac{2 \cdot 400^2 \kappa^4}{\epsilon}$. Note moreover that

$$\frac{1}{p} \sum_{j>k} \lambda_j(\mathbf{N}_1 + \mathbf{N}_2) = \frac{1}{p} \sum_{j>k} \sigma_j^2(\mathbf{T}\mathbf{W}) = \frac{1}{p^2} \sum_{j>k} \sigma_j^2(\mathbf{S}\mathbf{A}^\top \mathbf{T}^\top)$$

Using that $\mathbf{A} = \mathbf{A}^\top$ so that $\mathbf{Z} = \mathbf{S}\mathbf{A}^\top \mathbf{T}^\top$ we conclude that $\frac{1}{p^2} \sum_{i>k} \sigma_i^2(\mathbf{Z}) = \frac{n^2}{t^2} \sum_{i>k} \sigma_i^2(\mathbf{Z}) > \epsilon n^2/16$ as desired. Note that we conditioned on \mathcal{E}_i for $i = 1, 2, 3, 4, 5$, each of which held with probability 99/100, thus the result holds with probability 19/20 by a union bound. ■

We will now address the case where $\sigma_k(\mathbf{A}) > 10n/k$.

Lemma 4.9. *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be ϵ -far from PSD with $\|\mathbf{A}\|_\infty \leq 1$. Then let $\mathbf{Z} = \mathbf{S}\mathbf{A}^\top$ be samples as described above, so that \mathbf{Z} has an expected $t = \Theta(\log(1/\epsilon)/\epsilon^2)$ rows and columns, where t is scaled by a larger enough constant, and let $k = \frac{2 \cdot 400^2 \kappa^4}{\epsilon}$, where $\kappa \geq 1$ is the constant in Proposition 4.5. Suppose further that $\sigma_k(\mathbf{A}) > 10n/k$. Then with probability 49/50, we have*

$$\frac{n}{t} \sigma_k(\mathbf{Z}) \geq 8n/k$$

Proof. The proof is by application of Theorem 4.6 twice. We first generate a random row sampling matrix \mathbf{S} with an expected t rows, and bound $\lambda_k((\mathbf{SA})^\top \mathbf{SA}) = \sigma_k^2(\mathbf{SA})$. Let $\mathbf{X}_j \in \mathbb{R}^{n \times n}$ be a random variable such that $\mathbf{X}_j = \mathbf{A}_{(j)}^\top \mathbf{A}_{(j)}$, where $\mathbf{A}_{(j)}$ is the j -th row of \mathbf{A} that was sampled in \mathbf{S} . Then $\sum_j \mathbf{X}_j = (\mathbf{SA})^\top \mathbf{SA}$, and $\mathbb{E}[\mathbf{X}_j] = \frac{t}{n} \sum_{j=1}^n \mathbf{A}_j^\top \mathbf{A}_j = \frac{t}{n} \mathbf{AA}^\top$, where \mathbf{A}_j is the j -th row of \mathbf{A} . Moreover, note that $\|\mathbf{X}_j\|_2 \leq \max_i \|\mathbf{A}_{i,*}\|_2^2 \leq n$ for all j , by the boundedness of \mathbf{A} . Thus note that $\mu_k = \lambda_k((t/n)\mathbf{A}^\top \mathbf{A}) \geq (t/n)100n^2/k^2 = \frac{100tn}{k^2}$. Thus by the Interior Matrix Chernoff Bound 4.6, we have that for some constant c :

$$\begin{aligned} \Pr\left[\lambda_k((\mathbf{SA})^\top \mathbf{SA}) \leq .9\mu_k\right] &\leq k \cdot c^{\mu_k/L} \\ &\leq k \cdot c^{\frac{100tn}{k^2} \cdot \frac{1}{n}} \\ &\leq k \cdot e^{-100 \log(k)} \\ &\leq 1/1000 \end{aligned} \tag{29}$$

Where we use $t = \Theta(\frac{\log(1/\epsilon)}{\epsilon^2})$ with a large enough constant. Also condition on the fact that \mathbf{S} has at most $2t$ rows, which holds with probability $999/1000$. Call the union of the above two event \mathcal{E}_1 , which holds with probability $99/100$, and condition on it now. Given this, we have $\sigma_k^2(\mathbf{SA}) \geq \frac{90tn}{k^2}$. Now again, let $\mathbf{Y}_j = (\mathbf{SA})_{(j)}(\mathbf{SA})_{(j)}^\top$, where $(\mathbf{SA})_{(j)}$ is the j -th column of \mathbf{SA} sampled by the column sampling matrix \mathbf{T} . Let $\mathbf{M} = (\mathbf{SA})^\top$. Then again we have $\|\mathbf{Y}_j\|_2 \leq 2t$, using that \mathbf{SA} has at most $2t$ rows, and each entry is bounded by 1. Moreover, $\sum_j \mathbf{Y}_j = \mathbf{TMM}^\top \mathbf{T}^\top$. We also have $\lambda_k(\mathbb{E}[\sum_j \mathbf{Y}_j]) = \lambda_k(\frac{t}{n} \mathbf{MM}^\top) > \frac{90t^2}{k^2}$. Applying the Interior Matrix Chernoff Bound again, we have that for some constant c :

$$\begin{aligned} \Pr\left[\lambda_k(\mathbf{T}(\mathbf{SA})^\top (\mathbf{SA})\mathbf{T}^\top) \leq .9\mu_k\right] &\leq k \cdot c^{\mu_k/L} \\ &\leq k \cdot c^{\frac{90t^2}{k^2} \cdot \frac{1}{2t}} \\ &\leq k \cdot e^{-100 \log(k)} \\ &\leq 1/1000 \end{aligned} \tag{30}$$

Call the above event \mathcal{E}_2 . Conditioned on $\mathcal{E}_\infty \cup \mathcal{E}_2$, which hold together with probability $49/50$, we have that $\sigma_k(\mathbf{SAT}^\top) \leq .9\sqrt{\frac{90t^2}{k^2}} > 8t/k$. Since $\mathbf{Z} = \mathbf{SAT}^\top$, we have $\frac{n}{t}\sigma_k(\mathbf{Z}) > 8n/k$ as needed. \blacksquare

Theorem 4.10. Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be ϵ -far from PSD with $\|\mathbf{A}\|_\infty \leq 1$. Then if $S, T \subset [n]$ are random subsets with expected each size $t = O(\log(1/\epsilon)/\epsilon^2)$, then with probability $9/10$ the principal submatrix $\mathbf{A}_{(S \cup T) \times (S \cup T)}$ is not PSD.

Proof. First, by Chernoff bounds, with probability $99/100$ we have $|S \cup T| \leq |S| + |T| \leq 4t$, which we call \mathcal{E}_1 and condition on now. First, consider the case that $\sigma_k(\mathbf{A}) \leq 10n/k$, where $k = \frac{2 \cdot 400^2 \kappa^4}{\epsilon}$. Then by Lemma 4.8, with probability $19/20$, we have that $\sum_{i>k} \sigma_i^2(\mathbf{A}_{S \times T}) > \epsilon t^2/16$. Now we first prove the following claim:

Claim 4.11. Let $\mathbf{Z} \in \mathbb{R}^{n \times m}$ be any matrix, and let $\tilde{\mathbf{Z}}$ be a rectangular submatrix of \mathbf{Z} . for any Let $\mathbf{Z}_k, \tilde{\mathbf{Z}}_k$ be the truncated SVD of $\mathbf{Z}, \tilde{\mathbf{Z}}$ respectively, for any $1 \leq k \leq \min\{n, m\}$. Then we have

$$\|\mathbf{Z} - \mathbf{Z}_k\|_F^2 \geq \|\tilde{\mathbf{Z}} - \tilde{\mathbf{Z}}_k\|_F^2$$

Proof. Note that $\|\mathbf{Z} - \mathbf{Z}_k\|_F^2 \geq \|\tilde{\mathbf{Z}}_k - \mathbf{Z}'_k\|_F^2$, where \mathbf{Z}'_k is the matrix \mathbf{Z}_k restricted to the submatrix containing $\tilde{\mathbf{Z}}$. But $\tilde{\mathbf{Z}}_k$ is the *best* rank- k approximation to $\tilde{\mathbf{Z}}$, so $\|\tilde{\mathbf{Z}} - \tilde{\mathbf{Z}}_k\|_F^2 = \min_{\mathbf{B} \text{ rank-}k} \|\tilde{\mathbf{Z}} - \mathbf{B}\|_F^2 \leq \|\tilde{\mathbf{Z}}_k - \mathbf{Z}'_k\|_F^2$, using the fact that a submatrix of a rank- k matrix is at most rank k . ■

It follows that $\|\mathbf{A}_{(S \cup T) \times (S \cup T)} - (\mathbf{A}_{(S \cup T) \times (S \cup T)})_k\|_F^2 = \sum_{j>k} \sigma_j^2(\mathbf{A}_{(S \cup T) \times (S \cup T)}) \geq \sum_{j>k} \sigma_j^2(\mathbf{A}_{S \times T}) > \epsilon t^2 / 16 > \epsilon |S \cup T|^2 / 256$. But note that if $\mathbf{A}_{(S \cup T) \times (S \cup T)}$ was PSD, then we would have $\sum_{j>k} \sigma_j^2(\mathbf{A}_{S \times T}) \leq \frac{16}{k} t^2$, which is a contradiction since $k = \frac{2 \cdot 400^2 \kappa^4}{\epsilon} > \frac{100^2}{\epsilon}$.

Now consider the case that $\sigma_k(\mathbf{A}) > 10n/k$. Then by Lemma 4.9, we have $\sigma_k(\mathbf{A}_{S \times T}) \geq 8t/k$ with probability at least 49/50. Then $\|\mathbf{A}_{S \times T}\|_{S_1} \geq \sum_{i=1}^k \sigma_i(\mathbf{A}_{S \times T}) \geq 8t$. Using the fact that the Schatten norm of a matrix is always at least as large as the Schatten norm of any submatrix (this follows from the fact that the singular values of the submatrix are point-wise dominated by the larger matrix, see Theorem 1 [Tho72]), we have $\|\mathbf{A}_{(S \cup T) \times (S \cup T)}\|_{S_1} \geq 8t$. But note that if $\mathbf{A}_{(S \cup T) \times (S \cup T)}$ was PSD, then we would have $\|\mathbf{A}_{(S \cup T) \times (S \cup T)}\|_{S_1} = \text{Tr}(\mathbf{A}_{(S \cup T) \times (S \cup T)}) \leq |S \cup T| \leq 4t$, which is a contradiction. This completes the proof of the theorem. ■

Theorem 4.12. Fix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$. There is a non-adaptive sampling algorithm that, with probability 9/10, correctly distinguishes the case that \mathbf{A} is PSD from the case that \mathbf{A} is ϵ -far from PSD in ℓ_2 , namely that $\sum_{i: \lambda_i(\mathbf{A}) < 0} \frac{\lambda_i^2(\mathbf{A})}{n^2} \geq \epsilon$. The algorithm queries a total of $O(\frac{\log^2(1/\epsilon)}{\epsilon^4})$ entries of \mathbf{A} , and always correctly classifies \mathbf{A} as PSD if \mathbf{A} is indeed PSD. Moreover, the algorithm runs in time $\tilde{O}(1/\epsilon^{2\omega})$, where $\omega < 2.373$ is the exponent of fast matrix multiplication.

Proof. We first apply the algorithm of Section 3 with $\epsilon_0 = \frac{2}{k}$, which as discussed allows us to assume that $\lambda_i \geq -\epsilon_0 n / 1000 \geq -\frac{1}{2k} n$ for all i . The cost of doing so is $\tilde{O}(1/\epsilon^2)$ queries, and this algorithm also yields one-sided error as desired. The remainder of the theorem follows directly from Theorem 4.10, using that all principal submatrices of PSD matrices are PSD. Finally, for runtime, notice that the main computation is computing the eigenvalues of a $k \times k$ principal submatrix, for $k = \tilde{O}(1/\epsilon^2)$, which can be carried out in time $\tilde{O}(1/\epsilon^{2\omega})$ [DDHK07, BVKS19]. ■

5 Lower bounds

5.1 Lower Bound for PSD Testing with ℓ_∞ Gap

We begin by demonstrating a $O(1/\epsilon^2)$ lower bound for the problem of testing positive semi-definiteness with an ℓ_∞ gap. Our lower bound holds even when the algorithm is allowed to adaptively sample entries of \mathbf{A} .

Theorem 5.1. Any adaptive or non-adaptive algorithm which receives query access to $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$, and distinguishes with probability at least 2/3 whether

- \mathbf{A} is PSD.

- $x^T \mathbf{A} x < -\epsilon n$ for some unit vector $x \in \mathbb{R}^n$ and $\epsilon \in (0, 1)$

must make $\Omega(1/\epsilon^2)$ queries to A .

Proof. We construct two distributions $\mathcal{D}_1, \mathcal{D}_2$ over matrices, and draw the input A from the mixture $(\mathcal{D}_1 + \mathcal{D}_2)/2$. \mathcal{D}_1 is supported on one matrix: the zero matrix $\mathbf{0}^{n \times n}$, which is PSD. Now set $t = 2\epsilon^2 n$ and let $B \in \mathbb{R}^{n \times n}$ be the matrix given by

$$\mathbf{B} = \begin{bmatrix} 0 & -\mathbf{1}^{n-t \times t} \\ -\mathbf{1}^{t \times n-t} & -\mathbf{1}^{t \times t} \end{bmatrix}$$

Where $-\mathbf{1}^{n \times m}$ is the $n \times m$ matrix consisting of a -1 in each entry. Now let $x \in \mathbb{R}^{n \times n}$ be defined by $x_i = 1$ for $i = 1, 2, \dots, n-t$, and let $x_j = 1/\epsilon$ for $j > n-t$. Then note that $x^T \mathbf{B} x < -\frac{1}{\epsilon} \cdot 2\epsilon^2 n^2 < -\epsilon n \|x\|_2^2$, thus \mathbf{B} is ϵ -far from PSD in ℓ_∞ gap. To sample $A \sim \mathcal{D}_1$, we set $\mathbf{A} = \mathbf{P}_\Sigma \mathbf{B} \mathbf{P}_\sigma^T$, where \mathbf{P}_σ is a randomly drawn permutation matrix, namely $\sigma \sim S_n$ uniformly at random. Notice that to distinguish $A \sim \mathcal{D}_1$ from $A \sim \mathcal{D}_2$, the algorithm must read a non-zero entry. By Yao's min-max principle, we can assume that there is a deterministic algorithm that solves the problem with probability $2/3$ over the randomness of the distribution. Fix any $k < 1/(100\epsilon^2)$, and let s_1, s_2, \dots, s_k be the adaptive sequence of entries it would sample if $A_{s_i} = 0$ for each $i = 1, 2, \dots, k$. Then the probability that any of the s_i 's land in a row or a column of $\mathbf{A} = \mathbf{P}_\Sigma \mathbf{B} \mathbf{P}_\sigma^T$ with non-zero entries is at most $1/50$. Thus with probability $49/50$ under input from $A \sim \mathcal{D}_2$, the algorithm will output the same value had A been the all zero matrix. Thus the algorithm succeeds with probability at most $51/100$ when A is drawn from the mixture, demonstrating that $\Omega(1/\epsilon^2)$ samples are required for probability $2/3$ of success. ■

5.2 Lower Bound for PSD Testing with ℓ_2 Gap

We now present our main lower bound for PSD testing. Our result relies on the construction of explicit graphs with gaps in their spectrum, which have the property that they are indistinguishable given only a small number of queries to their adjacency matrices. Our lower bound is in fact a general construction, which will also result in lower bounds for testing the Schatten 1 norm, Ky-Fan norm, and cost of the best rank k approximation.

Roadmap In the following, we will first introduce the notation and theory required for the section, beginning with the notion of subgraph equivalence of matrices. We then construct our hard distributions $\mathcal{D}_1, \mathcal{D}_2$, and prove our main conditional results, Lemma 5.8, which demonstrates a lower bound for these hard distributions conditioned on the existence of certain pairs of subgraph equivalent matrices. Finally, we prove the existence of such matrices, which is carried out in the following Section 5.2.1. Putting these pieces together, we obtain our main lower bound in Theorem 5.12.

Preliminaries and Notation In the following, it will be useful to consider *signed graphs*. A signed graph Σ is a pair $(|\Sigma|, s)$, where $|\Sigma| = (V, E)$ is a simple graph, called the *underlying graph*, and $s : E \rightarrow \{1, -1\}$ is the *sign function*. We will sometimes abbreviate the signs equivalently as $\{+, -\}$. We will write E^+, E^- to denote the set of positive and negative edges. If Σ is a signed graph, we will often write $\Sigma = (V(\Sigma), E(\Sigma))$, where $E(\Sigma)$ is a set of *signed edges*, so $E(\Sigma) \subset \binom{V(\Sigma)}{2} \times \{+, -\}$ with the

property that for each $e \in \binom{V(\Sigma)}{2}$, at most one of $(e, +), (e, -)$ is contained in $E(\Sigma)$. For a signed graph G on n vertices, let $\mathbf{A}_G \in \{1, 0, -1\}^{n \times n}$ be its adjacency matrix, where $(\mathbf{A}_G)_{i,j}$ is the sign of the edge $e = (v_i, v_j)$ if $e \in E(G)$, and is 0 otherwise.

For a graph H , let $\|H\|$ denote the number of vertices in H . For any simple (unsigned) graph G , let \bar{G} be the signed graph obtained by having $E^+(\bar{G}) = E(G)$, and $E^-(\bar{G}) = \binom{V}{2} \setminus E(G)$. In other words, \bar{G} is the complete signed graph obtained by adding all the edges in the complement of G with a negative sign, and giving a positive sign the edges originally in G . We remark that the negation of the adjacency matrix of \bar{G} is known as the *Seidel matrix* of G . In what follows, we will often not differentiate between a signed graph G and its signed adjacency matrix \mathbf{A}_G . For graphs G, H , let $G \oplus H$ denote the disjoint union of two graphs G, H . We will assume familiarity with basic group theory. For groups G, H , we write $H \leq G$ if H is a subgroup of G . For a set T , let 2^T denote the power set of T . Throughout, let S_n denote the symmetric group on n letters. For two signed graphs Σ, H , let $\mathcal{F}_H(\Sigma) = \{G = (V(\Sigma), E(G)) \mid E(G) \subseteq E(\Sigma), G \cong H\}$ be the set of signed subgraphs of Σ isomorphic to H . For a permutation $\sigma \in S_n$, we write $\mathbf{P}_\sigma \in \mathbb{R}^{n \times n}$ to denote the row permutation matrix associated with σ . For $k \geq 3$, let C_k denote the cycle graph on k vertices.

For signed graphs G, H , a signed graph isomorphism (or just isomorphism) is a graph isomorphism that preserve the signs of the edges. For any set $U \subset [n] \times [n]$ and matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, we write \mathbf{A}_U to denote the matrix obtained by setting the entries $(\mathbf{A}_U)_{i,j} = A_{i,j}$ for $(i, j) \in U$, and $(\mathbf{A}_U)_{i,j} = 0$ otherwise. A set $U \subset [n] \times [n]$ is called symmetric if $(i, j) \in U \iff (j, i) \in U$. We call U simple if it does not contain any elements of the form (i, i) . We will sometimes refer to a simple symmetric U by the underlying simple undirected graph induced U .

Subgraph Equivalence We now formalize the indistinguishability property which we will require. For matrices \mathbf{A}, \mathbf{B} , when thought of as adjacency matrices of graphs, this property can be thought of as a more general version of “locally indistinguishability”, in the sense that, for any small subgraph H of \mathbf{A} , there is a *unique* subgraph of \mathbf{B} that is isomorphic to H . The following definition is more general, in the sense that a subgraph can also have “zero valued edges”, corresponding to the fact that an algorithm can learn of the non-existence of edges, as well as their existence.

Definition 5.2 (Sub-graph Equivalence). Fix any family \mathcal{U} of symmetric subsets $\mathcal{U} = \{U_i\}_i \in 2^{[n] \times [n]}$, and let $\Gamma \leq S_n$ be a subgroup of the symmetric group on n letters. Let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$. Then we say that \mathbf{A} is (\mathcal{U}, Γ) -subgraph isomorphic to \mathbf{B} , and write $\mathbf{A} \cong_{\mathcal{U}, \Gamma} \mathbf{B}$, if for every $U_i \in \mathcal{U}$ there is a bijection $\psi_i : \Gamma \rightarrow \Gamma$ such that

$$(\mathbf{P}_\sigma \mathbf{A} \mathbf{P}_\sigma^T)_{U_i} = (\mathbf{P}_{\psi_i(\sigma)} \mathbf{B} \mathbf{P}_{\psi_i(\sigma)}^T)_{U_i}$$

for all $\sigma \in \Gamma$. If G, H are two signed graphs on n vertices with adjacency matrices $\mathbf{A}_G, \mathbf{A}_H$, then we say that G is (\mathcal{U}, Γ) -subgraph equivalent to H , and write $G \cong_{\mathcal{U}, \Gamma} H$, if $\mathbf{A}_G \cong_{\mathcal{U}, \Gamma} \mathbf{A}_H$.

Note we do not require the U_i 's to be simple in the above definition. At times, if $\Gamma = S_n$, then we may omit Γ and just write $G \cong_{\mathcal{U}} H$ or $\mathbf{A} \cong_{\mathcal{U}} \mathbf{B}$.

Example 5.3. Let G, H be arbitrary graphs on n vertices, and let each $U_i = \{U_i\}$ be a simple graph consisting of a single edge. Then $G \cong_{\mathcal{U}, S_n} H$ if and only if $|E(G)| = |E(H)|$.

Example 5.4. Let G, H be arbitrary graphs on n vertices, and let $\mathcal{U} = \{U_i\}$ be a single graph, where U_i is a triangle on any three vertices. Then $G \cong_{\mathcal{U}, S_n} H$ if and only if the number of induced subgraphs on three vertices that are triangles, wedges, and single edges, are each the same in G as in H .

In what follows, we will consider graphs that are \mathcal{U} subgraph isomorphic, for a certain family of classes \mathcal{U} , which we now define. In what follows, recall that the matching number $\nu(G)$ of a graph G is the size of a maximum matching in G , or equivalently the maximum size of any subset of pairwise vertex disjoint edges in G .

Definition 5.5. For $1 \leq t \leq n$, let \mathcal{U}_n^t be the set of all undirected, possibly non-simple graphs U_i on n vertices, with the property that after removing all self-loops, U_i does not contains any set of t vertex disjoint edges. Equivalently, after removing all self-loops from U_i , the matching number $\nu(U_i)$ of U_i is less than t .

In other words, \mathcal{U}_n^t is the set of graphs with no set of t pair-wise non-adjacent edges e_1, \dots, e_t such that each e_i is not a self loop. Notice by the above definition that $\mathcal{U}_n^t \subset \mathcal{U}_n^{t+1}$. We will also need the following definition.

Definition 5.6. For any $n, m \leq 1$, let $\Gamma_{n,m} \leq S_{nm}$ be the subgroup defined $\Gamma_{n,m} = \{\sigma \in S_{nm} \mid \sigma(i, j) = (\pi(i), j), \pi \in S_n\}$, where the tuple $(i, j) \in [n] \times [m]$ indexes into $[nm]$ in the natural way.

Notice in particular, if $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $\mathbf{D} \in \mathbb{R}^{m \times m}$, then we have

$$\{\mathbf{P}_\sigma(\mathbf{A} \otimes \mathbf{D})\mathbf{P}_\sigma^T \mid \sigma \in \Gamma_{n,m}\} = \{(\mathbf{P}_\pi \otimes \mathbf{I}_m)(\mathbf{A} \otimes \mathbf{D})(\mathbf{P}_\pi \otimes \mathbf{I}_m)^T \mid \pi \in S_n\}$$

Note also by elementary properties of Kronecker products, we have $(\mathbf{P}_\pi \otimes \mathbf{I}_m)(\mathbf{A} \otimes \mathbf{D})(\mathbf{P}_\pi \otimes \mathbf{I}_m)^T = (\mathbf{P}_\pi \mathbf{A} \mathbf{P}_\pi^T) \otimes \mathbf{D}$. For such a $\sigma \in \Gamma_{n,m}$, we write $\sigma = \pi \otimes \text{id}$, where $\pi \in S_n$

Lemma 5.7. Fix any $t, m \geq 1$, and let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$ be matrices with $\mathbf{A} \cong_{\mathcal{U}_n^t, S_n} \mathbf{B}$, where \mathcal{U}_n^t is defined as above, and let $\mathbf{T} \in \mathbb{R}^{m \times m}$ be any matrix. Then $\mathbf{A} \otimes \mathbf{T} \cong_{\mathcal{U}_{nm}^t, \Gamma_{n,m}} \mathbf{B} \otimes \mathbf{T}$, where $\Gamma \leq S_{nm}$ is as defined above.¹¹

Proof. Fix any $U'_i \in \mathcal{U}_{nm}^t$. Note that every edge of U'_i corresponds to a unique edge of a graph on n vertices. This can be seen as every edge of U'_i is of the form $((i_1, j_1), (i_2, j_2))$ where $i_1, i_2 \in [n]$, $j_1, j_2 \in [m]$, which corresponds to the edge $(i_1, i_2) \in [n] \times [n]$. So let $U_i \subset [n] \times [n]$ be the set of all such edges induced by the edges of U'_i . Observe, of course, that many distinct edges of U'_i could result in the same edge of U_i . We claim that $U_i \in \mathcal{U}_n^t$. Suppose this was not the case, and let $e_1, \dots, e_t \in U_i$ be vertex disjoint non-self loop edges, where $e_\ell = (i_\ell, j_\ell)$, $i_\ell \neq j_\ell$. Then for each $\ell \in [t]$, there must be at least one edge $e'_\ell \in U'_i$ such that $e'_\ell = ((i_\ell, a_\ell)(j_\ell, b_\ell)) \in U'_i$, and we can fix e'_ℓ to be any such edge. Then since each vertex $i_\ell \in [n]$ occured in at most one edge of e_1, \dots, e_t by assumption, it follows that each vertex $(i_\ell, j_\ell) \in [n] \times [m]$ occurs at most once in e'_1, \dots, e'_t , which contradicts the fact that the $U'_i \in \mathcal{U}_{nm}^t$.

Now that we have $U_i \in \mathcal{U}_n^t$, since $\mathbf{A} \cong_{\mathcal{U}_n^t, S_n} \mathbf{B}$ we have a bijection function $\psi_i : S_n \rightarrow S_n$ such that $(\mathbf{P}_\pi \mathbf{A} \mathbf{P}_\pi^T)_{U_i} = (\mathbf{P}_{\psi_i(\pi)} \mathbf{B} \mathbf{P}_{\psi_i(\pi)}^T)_{U_i}$. We now define the mapping $\hat{\psi}_i : \Gamma_{n,m} \rightarrow \Gamma_{n,m}$ by $\hat{\psi}_i(\pi \otimes \text{id}) = \psi_i(\pi) \otimes \text{id}$, and show that it satisfies the conditions of Definition 5.2. Now note that each $\sigma = \pi \otimes \text{id} \in \Gamma_{n,m}$ satisfies $\mathbf{P}_\sigma = \mathbf{P}_\pi \otimes \mathbf{I}$, and so $\mathbf{P}_\sigma(\mathbf{A} \otimes \mathbf{T})\mathbf{P}_\sigma^T = \mathbf{P}_\pi \mathbf{A} \mathbf{P}_\pi^T \otimes \mathbf{T}$.

¹¹Note that this fact extends naturally to tensoring with rectangular matrices \mathbf{T} .

We now claim that for any $U'_i \in \mathcal{U}_{nm}^t$, if we construct $U_i \in \mathcal{U}_n^t$ as above, we have that for any matrix $\mathbf{Z} \in \mathbb{R}^{n \times n}$ the non-zero entries of $(\mathbf{Z})_{U_i} \otimes \mathbf{T}$ contain the non-zero entries of $(\mathbf{Z} \otimes \mathbf{T})_{U'_i}$. As a consequence, if $(\mathbf{Z})_{U_i} \otimes \mathbf{T} = (\mathbf{Y})_{U_i} \otimes \mathbf{T}$ for some other matrix $\mathbf{Y} \in \mathbb{R}^{n \times n}$, we also have $(\mathbf{Z} \otimes \mathbf{T})_{U'_i} = (\mathbf{Y} \otimes \mathbf{T})_{U'_i}$. But the claim in question just follows from the construction of U'_i , since for every entry $((i_1, j_1), (i_2, j_2)) \in U'_i$ we added the entry $(i_1, i_2) \in U_i$. Now since we have that $(\mathbf{P}_\pi \mathbf{A} \mathbf{P}_\pi^T)_{U_i} = (\mathbf{P}_{\psi_i(\pi)} \mathbf{B} \mathbf{P}_{\psi_i(\pi)}^T)_{U'_i}$, we also obtain

$$(\mathbf{P}_\pi \mathbf{A} \mathbf{P}_\pi^T)_{U_i} \otimes \mathbf{T} = (\mathbf{P}_{\psi_i(\pi)} \mathbf{B} \mathbf{P}_{\psi_i(\pi)}^T)_{U'_i} \otimes \mathbf{T}$$

which as just argued implies that

$$(\mathbf{P}_\pi \mathbf{A} \mathbf{P}_\pi^T \otimes \mathbf{T})_{U'_i} = (\mathbf{P}_{\psi_i(\pi)} \mathbf{B} \mathbf{P}_{\psi_i(\pi)}^T \otimes \mathbf{T})_{U'_i}$$

Since $(\mathbf{P}_\pi \mathbf{A} \mathbf{P}_\pi^T \otimes \mathbf{T})_{U'_i} = (\mathbf{P}_\sigma (\mathbf{A} \otimes \mathbf{T}) \mathbf{P}_\sigma^T)_{U'_i}$ and $(\mathbf{P}_{\psi_i(\pi)} \mathbf{B} \mathbf{P}_{\psi_i(\pi)}^T \otimes \mathbf{T})_{U'_i} = (\mathbf{P}_{\hat{\psi}(\sigma)} (\mathbf{B} \otimes \mathbf{T}) \mathbf{P}_{\hat{\psi}(\sigma)}^T)_{U'_i}$, it follows that $\mathbf{A} \otimes \mathbf{T} \cong_{\mathcal{U}_{nm}^t, \Gamma_{n,m}} \mathbf{B} \otimes \mathbf{T}$ as required. ■

The Hard Instance. We now describe now distributions, $\mathcal{D}_1, \mathcal{D}_2$, supported on $n \times n$ matrices \mathbf{A} and parameterized by a value $k \geq 1$, such that distinguishing \mathcal{D}_1 from \mathcal{D}_2 requires $\Omega(k^2)$ samples. The distributions are parameterized by *three* matrices, $(\mathbf{B}, \mathbf{D}, \mathbf{Z})$, which are promised to satisfy the properties that $\mathbf{B}, \mathbf{D} \in \mathbb{R}^{d \times d}$ with $\mathbf{B} \cong_{\mathcal{U}_{d,S_d}^t} \mathbf{D}$ for some $t \leq d$, and $\mathbf{Z} \in \mathbb{R}^{m \times m}$, where $m = n/(dk)$. Also define $\tilde{\mathbf{B}} = \mathbf{B} \otimes \mathbf{Z}$, $\tilde{\mathbf{D}} = \mathbf{D} \otimes \mathbf{Z}$. We now define the distribution. We first define \mathcal{D}_1 . In \mathcal{D}_1 , we select a random partition of $[n]$ into L_1, \dots, L_k , where each $|L_i| = n/k$ exactly. Then for each $i \in [k]$, we select a uniformly random $\sigma_i \in \Gamma_{d,m}$ and set $\mathbf{A}_{L_i \times L_i} = \mathbf{P}_{\sigma_i} \tilde{\mathbf{B}} \mathbf{P}_{\sigma_i}^T$, and the remaining elements of \mathbf{A} are set to 0. In \mathcal{D}_2 , we perform the same procedure, but set $\mathbf{A}_{L_i \times L_i} = \mathbf{P}_{\sigma_i} \tilde{\mathbf{D}} \mathbf{P}_{\sigma_i}^T$. So if $\mathbf{A} \sim \frac{\mathcal{D}_1 + \mathcal{D}_2}{2}$, then \mathbf{A} is block-diagonal, with each block having size n/k . We first demonstrate that for any matrices $(\mathbf{B}, \mathbf{D}, \mathbf{Z})$ satisfying the above properties, distinguishing these distributions requires $\Omega(k^2)$ samples. We assume in the following that dk divides n , which will be without loss of generality since we can always embed a small instance of the lower bound with size n' such that $n/2 < n - dk \leq n' \leq n$, and such that dk divides n .

Lemma 5.8. *Fix any $1 \leq k, d \leq n$. Let $(\mathbf{B}, \mathbf{D}, \mathbf{Z})$ be any three matrices such that $\mathbf{B}, \mathbf{D} \in \mathbb{R}^{d \times d}$, $\mathbf{B} \cong_{\mathcal{U}_{d,S_d}^t} \mathbf{D}$ where $t = \log k$, and $\mathbf{Z} \in \mathbb{R}^{m \times m}$, where $m = n/(dk)$. Then any non-adaptive sampling algorithm which receives $\mathbf{A} \sim \frac{\mathcal{D}_1 + \mathcal{D}_2}{2}$ where the distributions are defined by the tuple $(\mathbf{B}, \mathbf{D}, \mathbf{Z})$ as above, and distinguishes with probability at least $2/3$ whether \mathbf{A} was drawn from \mathcal{D}_1 or \mathcal{D}_2 must sample $\Omega(k^2)$ entries of \mathbf{A} .*

Proof. We show that any algorithm cannot distinguish \mathcal{D}_1 from \mathcal{D}_2 with probability greater than $2/3$ unless it makes at least $\ell > C \cdot k^2$ queries, for some constant $C > 0$. So suppose the algorithm makes at most $C \cdot k^2/100$ queries in expectation and is correct with probability $2/3$. Then by Markov's there is a algorithm that always makes at most $\ell = Ck^2$ queries which is correct with probability $3/5$. By Yao's min-max principle, there is a deterministic algorithm making this many queries which is correct with probability $3/5$ over the distribution $\frac{\mathcal{D}_1 + \mathcal{D}_2}{2}$. So fix this algorithm, which consists of a single subset $U \subset [n] \times [n]$ with $|U| = \ell$.

We now generate the randomness used to choose the partition L_1, \dots, L_k of $[n]$. Let $U_i = U \cap L_i \times L_i = \{(i, j) \in U \mid i, j \in L_i\}$. Let \mathcal{E}_i be the event that $U_i \in \mathcal{U}_{md}^t$. We first bound $\Pr[\neg \mathcal{E}_i]$,

where the probability is over the choice of the partition $\{L_i\}_{i \in [k]}$. For $\neg \mathcal{E}_i$, there must be t pairwise vertex disjoint non-self loop edges $e_1, \dots, e_t \in U$ such that $e_j = (a_j, b_j)$ and $a_j, b_j \in L_i$. In other words, we must have $2t$ distinct vertices $a_1, b_1, \dots, a_t, b_t \in L_i$. For a fixed vertex $v \in [n]$, this occurs with probability $1/k$, and the probability that another $u \in [n] \cap L_i$ conditioned on $v \in [n]$ is strictly less than $1/k$ as we have $|L_i| = n/k$ exactly. Thus, the probability that all $2t$ vertices are contained in L_i can then be bounded $\frac{1}{k^{2t}}$. Now there are $\binom{|U|}{t} \leq \ell^t$ possible choices of vertex disjoint edges $e_1, \dots, e_t \in U$ which could result in \mathcal{E}_i failing to hold, thus $\Pr[\mathcal{E}_i] \geq 1 - \frac{\ell^t}{k^{2t}}$ and by a union bound

$$\begin{aligned} \Pr[\cap_{i=1}^k \mathcal{E}_i] &\geq 1 - \frac{\ell^t}{k^{2t-1}} \\ &\geq 1 - \frac{C^t k^{2t}}{k^{2t-1}} \\ &\geq 1 - C^t k \\ &\geq \frac{99}{100} \end{aligned} \tag{31}$$

Where in the last line, we took $C \leq 1/10$ and used the fact that $t = \log(k)$. Then if $\mathcal{E} = \cap_{i=1}^k \mathcal{E}_i$, we have $\Pr[\mathcal{E}] > 99/100$, which we condition on now, along with any fixing of the L_i 's that satisfies \mathcal{E} . Conditioned on this, it follows that $U_i \in \mathcal{U}_{md}$ for each $i \in [k]$. Using that $B \cong_{\mathcal{U}_d^t, S_n} D$, we can and apply Lemma 5.7 to obtain $\mathbf{B} \otimes \mathbf{Z} \cong_{\mathcal{U}_{dm}^t, \Gamma_{d,m}} \mathbf{D} \otimes \mathbf{Z}$. Thus, for each $i \in [k]$ we can obtain a bijection function $\psi_i : \Gamma_{d,m} \rightarrow \Gamma_{d,m}$ such that $(\mathbf{P}_{\sigma_i} \widetilde{\mathbf{B}} \mathbf{P}_{\sigma_i}^T)_{U_i} = (\mathbf{P}_{\psi_i(\sigma_i)} \widetilde{\mathbf{D}} \mathbf{P}_{\psi_i(\sigma_i)}^T)_{U_i}$ for each $\sigma_i \in \Gamma_{d,m}$. Thus we can create a coupling of draws from \mathcal{D}_1 with those from \mathcal{D}_2 conditioned on \mathcal{E} , so for any possible draw from the remaining randomness of \mathcal{D}_1 , which consists only of drawing some $(\sigma_1, \dots, \sigma_k) \in S_d^k$ generating a matrix \mathbf{A}_1 , we have a unique corresponding draw $(\psi_1(\sigma_1), \dots, \psi_k(\sigma_k)) \in S_d^k$ of the randomness in \mathcal{D}_2 which generates a matrix \mathbf{A}_2 , such that $(\mathbf{A}_1)_U = (\mathbf{A}_2)_U$. Thus conditioned on \mathcal{E} , any algorithm is correct on $\frac{\mathcal{D}_1 + \mathcal{D}_2}{2}$ with probability exactly $1/2$. Since \mathcal{E} occurred with probability $99/100$, it follows that the algorithm is correct with probability $51/100 < 3/5$, which is a contradiction. Thus we must have $\ell \geq Ck^2 = \Omega(k^2)$ as needed. ■

We are now ready to introduce our construction of the matrices as required in the prior lemma. Recall that $k \geq 3$, let C_k denote the cycle graph on k vertices.

Fact 5.9. Fix any $n \geq 3$. We have $\lambda_{\min}(C_{2n+1}) = -2 + \Theta(1/n^2)$ and $\lambda_{\min}(C_n \oplus C_{n+1}) = -2$.

Proof. The eigenvalues of the cycle C_ℓ are given by $2 \cos(\frac{2\pi t}{\ell})$ [Chu96] for $t = 0, \dots, \ell - 1$, which yields the result using the fact that $\cos(\pi(1 + \epsilon)) = 1 + \Theta(\epsilon^2)$ for small ϵ ■

Proposition 5.10. Fix any $n = n_1 + n_2$. For any $t \leq \min\{n_1, n_2\}/4$, we have $C_n \cong_{\mathcal{U}_n^t, S_n} C_{n_1} \oplus C_{n_2}$

Proof. We begin by fixing any set $U_i \in \mathcal{U}_n^t$. For any signed graph Σ and graph G on n vertices such that the maximum set of vertex disjoint edges in Σ is $t \leq \min\{n_1, n_2\}/4$, let $\mathcal{H}_\Sigma(G) = \{\sigma \in S_n \mid \mathbf{P}_\sigma \mathbf{A}_\Sigma \mathbf{P}_\sigma^T = \mathbf{A}_H, H \subset G\}$ and let $\mathcal{H}_\Sigma^{-1}(G) = \{\sigma \in S_n \mid \mathbf{A}_\Sigma = \mathbf{P}_\sigma \mathbf{A}_H \mathbf{P}_\sigma^T, H \subset G\}$. By Corollary 5.15, we have $|\mathcal{H}_\Sigma(\overline{C_n})| = |\mathcal{H}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})|$ whenever $|\Sigma|$ has no set of at least $\min\{n_1, n_2\}/4$ vertex disjoint edges. Since S_n is a group and has unique inverses, we also have $|\mathcal{H}_\Sigma^{-1}(\overline{C_n})| = |\mathcal{H}_\Sigma^{-1}(\overline{C_{n_1} \oplus C_{n_2}})| = |\mathcal{H}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})| = |\mathcal{H}_\Sigma^{-1}(\overline{C_{n_1} \oplus C_{n_2}})|$.

We now define a function $\psi_i : S_n \rightarrow S_n$ such that $(\mathbf{P}_\sigma \mathbf{A}_{\bar{C}_n} \mathbf{P}_\sigma^T)_{U_i} = (\mathbf{P}_{\psi_i(\sigma)} \mathbf{A}_{\bar{C}_{n_1} \oplus \bar{C}_{n_2}} \mathbf{P}_{\psi_i(\sigma)}^T)_{U_i}$ for every $\sigma \in S_n$. Now fix any signed graph Σ such that $\Sigma = (\mathbf{P}_\sigma \mathbf{A}_{\bar{C}_n} \mathbf{P}_\sigma^T)_{U_i}$ for some $\sigma \in S_n$. Note that the set of $\pi \in S_n$ such that $\Sigma = (\mathbf{P}_\pi \mathbf{A}_{\bar{C}_n} \mathbf{P}_\pi^T)_{U_i}$ is precisely $\mathcal{H}_\Sigma^{-1}(\bar{C}_n)$. Similarly, the set $\pi \in S_n$ such that $\Sigma = (\mathbf{P}_\pi \mathbf{A}_{\bar{C}_{n_1} \oplus \bar{C}_{n_2}} \mathbf{P}_\pi^T)_{U_i}$ is precisely $\mathcal{H}_\Sigma^{-1}(\bar{C}_{n_1} \oplus \bar{C}_{n_2})$. Also, by construction of \mathcal{U}_n^t , we know that the maximum set of vertex disjoint edges in U_i , and therefore in Σ is $t \leq \min\{n_1, n_2\}/4$. So by the above, we know there is a bijection $\psi_i^\Sigma : \mathcal{H}_\Sigma^{-1}(\bar{C}_n) \rightarrow \mathcal{H}_\Sigma^{-1}(\bar{C}_{n_1} \oplus \bar{C}_{n_2})$ for every such realizable matrix Σ . Taking $\psi_i(\sigma) = \psi_i^{\mathbf{P}_\sigma \mathbf{A}_{\bar{C}_n} \mathbf{P}_\sigma^T}(\sigma)$ satisfies the desired properties for $\bar{C}_n \cong_{\mathcal{U}_{n, S_n}^t} \bar{C}_{n_1} \oplus \bar{C}_{n_2}$. Notice that this implies that $C_n \cong_{\mathcal{U}_{n, S_n}^t} C_{n_1} \oplus C_{n_2}$, since C_n and $C_{n_1} \oplus C_{n_2}$ are both obtained from obtained \bar{C}_n and $\bar{C}_{n_1} \oplus \bar{C}_{n_2}$ by changing every entry with the value -1 to 0 .

■

Proposition 5.11. *Fix any $t > 1$, and set either $d_0 = 4t$, and $d = 2d_0 + 1$. Set $\mathbf{B} = 1/2(\mathbf{A}_{C_d} + \lambda \mathbf{I}_d)$ and $\mathbf{D} = 1/2(\mathbf{A}_{C_{d_0} \oplus C_{d_0+1}} + \lambda \mathbf{I}_d)$, where $\lambda = -2 \cos(\frac{2\pi d_0}{2d_0+1})$. Then we have that \mathbf{B} is PSD, $\lambda_{\min}(\mathbf{D}) < -\delta$ where $\delta = \Theta(1/d^2)$, $\|\mathbf{B}\|_\infty, \|\mathbf{D}\|_\infty \leq 1$, and $\mathbf{B} \cong_{\mathcal{U}_{d, S_d}^t} \mathbf{D}$.*

Proof. By Proposition 5.10, we know $C_{2d_0+1} \cong_{\mathcal{U}_{2d_0+1, S_{2d_0+1}}^t} C_{d_0} \oplus C_{d_0+1}$, so to show subgraph equivalence suffices to show that adding $\lambda \mathbf{I}_{2d_0+1}$ to both C_{2d_0+1} and $C_{d_0} \oplus C_{d_0+1}$ does not effect the fact that they are $\mathcal{U}_{d_0}^t, S_{d_0}$ subgraph-equivalent. But note that this fact is clear, since we have only changed the diagonal which is still equal to λ everywhere for both \mathbf{B}, \mathbf{D} . Namely, for any $\sigma, \pi \in S_{2d_0+1}$ and $i \in [2d_0+1]$ we have $(\mathbf{P}_\sigma \mathbf{B} \mathbf{P}_\sigma^T)_{(i,i)} = (\mathbf{P}_\pi \mathbf{D} \mathbf{P}_\pi^T)_{(i,i)} = \lambda$, thus the subgraph equivalence between C_{2d_0+1} and $C_{d_0} \oplus C_{d_0+1}$ still holds using the same functions ψ_i as required for $C_{2d_0+1} \cong_{\mathcal{U}_{2d_0+1, S_{2d_0+1}}^t} C_{d_0} \oplus C_{d_0+1}$. Note that the L_∞ bound on the entries follows from the fact that adjacency matrices are bounded by 1 and zero on the diagonal, $\lambda \leq 2$, and we scale each matrix down by $1/2$. Next, by Fact 5.9, we know that \mathbf{B} is PSD and $\lambda_{\min}(\mathbf{D}) = -\Theta(\frac{1}{d^2})$, which holds still after scaling by $1/2$, and completes the proof.

■

We now state our main theorem, which is direct result of instantiating the general lower bound of Lemma 5.8 with the matrices as described above in Proposition 5.11.

Theorem 5.12. *Any non-adaptive sampling algorithm which solves with probability at least $2/3$ the PSD testing problem with ϵ - ℓ_2^2 gap must query at least $\tilde{\Omega}(\frac{1}{\epsilon^2})$ entries of the input matrix.*

Proof. Set $k = C \frac{1}{\epsilon \log^6(1/\epsilon)}$ for a small enough constant $C > 0$. Also set $t = \log k$, $d_0 = 4t$, $d = 2d_0 + 1$, and as before set $m = n/(dk)$. We first apply Lemma 5.8 with $\mathbf{Z} = \mathbf{1}^{m \times m}$, and the matrices $\mathbf{B} = 1/2(\mathbf{A}_{C_d} + \lambda \mathbf{I}_d)$ and $\mathbf{D} = 1/2(\mathbf{A}_{C_{d_0} \oplus C_{d_0+1}} + \lambda \mathbf{I}_d)$ from Proposition 5.11, where $\lambda = -2 \cos(\frac{2\pi d_0}{2d_0+1})$. Then by Lemma 5.8, using that $\mathbf{B} \cong_{\mathcal{U}_{d, S_d}^t} \mathbf{D}$ via Proposition 5.11, it follows that any non-adaptive sampling algorithm that distinguishes \mathcal{D}_1 from \mathcal{D}_2 requires $\Omega(k^2)$ samples.

We now demonstrate every instance of \mathcal{D}_1 and \mathcal{D}_2 satisfy the desired ℓ_2^2 -gap as defined in Problem 1.2. First, since the eigenvalues of the Kronecker product $\mathbf{Y} \otimes \mathbf{Z}$ of any matrices \mathbf{Y}, \mathbf{Z} are all pairwise eigenvalues of the matrices \mathbf{Y}, \mathbf{Z} , it follows that $\tilde{\mathbf{B}}$ is PSD as \mathbf{B} is PSD by Proposition 5.11 and $\mathbf{1}^{m \times m} = \mathbf{1}^m (\mathbf{1}^m)^T$ is PSD. By the same fact and Proposition 5.11, since $\lambda_1(\mathbf{1}^{m \times m}) = m$, we have that $\lambda_{\min}(\tilde{\mathbf{D}}) = -\Theta(m/(d^2)) = -\Theta(\frac{n}{d^3 k})$. Now note that if $\mathbf{A}_1 \sim \mathcal{D}_1$, then \mathbf{A}_1 is a block-diagonal matrix where each block is PSD, thus \mathbf{A}_1 is PSD. Note also that if $\mathbf{A}_2 \sim \mathcal{D}_2$, then \mathbf{A}_2 is a block-diagonal matrix where each block is has an eigenvalue smaller than $-C' \frac{n}{d^3 k}$ for some constant $C' > 0$. Since

the eigenvalues of a block diagonal matrix are the union of the eigenvalues of the blocks, it follows that

$$\begin{aligned}
\sum_{i: \lambda_i(\mathbf{A}_2) < 0} (\lambda_i(\mathbf{A}_2))^2 &= \sum_{i=1}^k \lambda_{\min}(\tilde{\mathbf{D}})^2 \\
&\geq k(C' \frac{n}{d^3 k})^2 \\
&= \left(\frac{(C')^2 \log^6(1/\epsilon)}{C(8 \log(\frac{C}{\epsilon \log^6(1/\epsilon)}) + 1)^6} \right) \cdot \epsilon n^2 \\
&\geq \epsilon n^2
\end{aligned} \tag{32}$$

Where the last inequality follows from setting the constant $C = \frac{(C')^2}{100^6}$ so that

$$\begin{aligned}
\frac{(C')^2 \log^6(1/\epsilon)}{C \left(8 \log(\frac{C}{\epsilon \log^6(1/\epsilon)}) + 1 \right)^6} &= \frac{100^6 \log^6(1/\epsilon)}{\left(8 \log(\frac{(C')^2}{100^6 \epsilon \log^6(1/\epsilon)}) + 1 \right)^6} \\
&\geq \frac{100^6 \log^6(1/\epsilon)}{(16 \log(\frac{1}{\epsilon}))^6} \\
&> 1
\end{aligned} \tag{33}$$

and using that the first inequality above holds whenever $(\frac{1}{\epsilon})^{16} \leq \left(\frac{(C')^2}{3 \cdot 100^6 \epsilon \log^6(1/\epsilon)} \right)^8$, which is true so long as $\epsilon < C_0$ for some constant C_0 . Note that if $\epsilon > C_0$, then a lower bound of $\Omega(1) = \Omega(1/\epsilon^2)$ follows from the one heavy eigenvalue ℓ_∞ gap lower bound. Thus $\mathbf{A}_1, \mathbf{A}_2$ satisfies the ϵ - L_2 gap property as needed, which completes the proof. ■

5.2.1 $C_{n_1+n_2}$ is Subgraph Equivalent to $C_{n_1} \oplus C_{n_2}$

In this section, we demonstrate the subgraph equivalence of the the cycle $C_{n_1+n_2}$ and union of cycles $C_{n_1} \oplus C_{n_2}$. In order to refer to edges which are not in the cycles $C_{n_1+n_2}$ and $C_{n_1} \oplus C_{n_2}$, it will actually be convenient to show that $\overline{C_{n_1+n_2}}$ is subgraph equivalent to $\overline{C_{n_1} \oplus C_{n_2}}$, where recall that \overline{G} for a simple graph G is the result of adding negative edges to G for each edge $e = (u, v) \notin E(G)$. Equivalently, the adjacency matrix of \overline{G} is the result of replacing the 0's on the off-diagonal of A_G with -1 's. Notice that, by the definition of subgraph equivalence, it does not matter whether these values are set to 0 or to -1 .

Overview of the bijection. We now intuitively describe the bijection of Lemma 5.13, which demonstrates that for any signed graph Σ such that any set of pairwise vertex disjoint edges $\{e_1, \dots, e_k\}$ (i.e. any matching) in Σ has size at most $k \leq \min\{n_1, n_2\}/4$, the number of subgraphs of $\overline{C_{n_1+n_2}}$ isomorphic to Σ is the same as the number of subgraphs of $\overline{C_{n_1} \oplus C_{n_2}}$ isomorphic to Σ . So let H be any subgraph of $\overline{C_{n_1+n_2}}$ that is isomorphic to Σ . For simplicity, let $n_1 = n_2$, and suppose H




Figure 2: An illustration of the bijection in Lemma 5.13, when H only contains positive edges. The three colored paths represent the graph H , which must be mapped from C_{2n} to $C_n \oplus C_n$. Since the paths intersect the edges $(2n, 1)$ and $(n, n+1)$ to be cut, we must first swap the last four vertices $\{n-4, \dots, n\}$ and $\{2n-4, \dots, 2n\}$ of C_{2n} before the two splitting points $n, 2n$, and then cut the cycle. Note that four is the smallest number of vertices which can be swapped, without swapping in the middle of a path of H .

contains only positive edges, so that H is actually a subgraph of the unsigned cycle C_{2n} . Since Σ has at most $n/4$ edges, $\Sigma \cong H$ must be a collection of disjoint paths. So the problem can be described as an arrangement problem: for each arrangement H of Σ in C_{2n} , map it to a unique arrangement H' of Σ in $C_n \oplus C_n$.

We would like to construct such a mapping by “splitting” the big cycle C_{2n} into two smaller cycles, see Figure 2 for an example. Specifically, we could split the cycle C_{2n} down the middle, cutting the edges $(n, n+1)$, and $(n, 1)$, and instead connecting the first vertex to the n -th and the $n+1$ -st to the $2n$ -th. Now if H does not contain either of the cut edges, then the resulting collection of paths will be an isomorphic copy of H living inside of $C_n \oplus C_n$. However, if H does contain such an edge, we cannot cut the cycle here, as the resulting paths inside of $C_n \oplus C_n$ would not be isomorphic. For example see Figure 2, where if we just cut the edge between $(n, n+1)$ and rerouted it to $(n, 1)$, then the red cycle with 4 vertices would be disconnected into a cycle of length three, and an isolated vertex. To handle this, before cutting and rerouting the edges $(n, n+1)$ and $(2n, 1)$, we first *swap* the last i vertices before the cutting points, for some i . Namely, we swap the vertices $(n-i, n-i+1, \dots, n)$ with $(2n-i, 2n-i+1, \dots, 2n)$ and *then* split the graph at the edges $(n, n+1)$ and $(2n, 1)$. For the resulting graphs to be isomorphic, we cannot swap in the middle of a path, thus the value i is chosen as the *smallest* $i \geq 0$ such that the edges $(n-i-1, n-i)$ and $(2n-i-1, 2n-i)$ do not exist in any path of H . Moreover, such an i must exist, so long as H has fewer than $\min\{n_1, n_2\}$ edges (the stronger bound of $\min\{n_1, n_2\}/4$ is only needed for the more general case, where negative edges are included).

One can show that this mapping is actually an involution; namely, given the collection of paths H' in $C_n \oplus C_n$ which are obtained from applying the function on H , one can similarly find the smallest $i \geq 0$ such that the edges $(n-i-1, n-i)$ and $(2n-i-1, 2n-i)$ are not in H' , which must in fact be the same value of i used when mapping H ! Then, by swapping the last i vertices before n

and $2n$, and then reconnecting $C_n \oplus C_n$ into a single cycle, one obtains the original graph H . From this, demonstrating bijectivity becomes relatively straightforward. Extending this to the case where H is allowed to contain negative edges of $\overline{C_{2n}}$ follows similar steps, albeit with a stronger condition on the choice of i . The full proof is now presented below.

Lemma 5.13. *Fix any $n = n_1 + n_2$. Fix any simple graph $|\Sigma|$, such that any set of vertex disjoint edges $\{e_1, \dots, e_k\}$ in $|\Sigma|$ has size at most $k \leq \min\{n_1, n_2\}/4$, and let $\Sigma = (|\Sigma|, \sigma)$ be any signing of $|\Sigma|$. Let $\mathcal{F}_\Sigma(\overline{C_n})$ denote the set of subgraphs of $\overline{C_n}$ isomorphic to $|\Sigma|$, and similarly define $\mathcal{F}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})$. Then we have*

$$|\mathcal{F}_\Sigma(\overline{C_n})| = |\mathcal{F}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})|$$

Proof. Order the vertices of the cycle $C_n = \{1, 2, \dots, n\}$, which we will describe as the same vertex set for $C_{n_1} \oplus C_{n_2}$, where $\{1, \dots, n_1\}$ are the vertices of the first cycle C_{n_1} and $\{n_1 + 1, \dots, n\}$ are the vertices of C_{n_2} . We derive a bijection $\varphi : \mathcal{F}_\Sigma(\overline{C_n}) \rightarrow \mathcal{F}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})$. We describe a point $\mathbf{X} \in \mathcal{F}_H(\overline{C_n}) \cup \mathcal{F}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})$ by its (signed) adjacency matrix $\mathbf{X} \in \{-1, 0, 1\}^{n \times n}$. Namely, $\mathbf{X} \in \{-1, 0, 1\}^{n \times n}$ is any matrix obtained by setting a subset of the entries of $\mathbf{A}_{\overline{C_{n_1+n_2}}}$ or $\mathbf{A}_{\overline{C_{n_1} \oplus C_{n_2}}}$ equal to 0, such that the signed graph represented by \mathbf{X} is isomorphic to Σ . In this following, we will always modularly interpret the vertex $v_{n+i} = v_i$ for $i \geq 1$.

Thus, we can now think of φ as being defined on the subset of the matrices $\{-1, 0, 1\}^{n \times n}$ given by the adjacency matrices of signed graphs in $\mathcal{F}_\Sigma(\overline{C_n})$. In fact, it will be useful to define φ on a larger domain. Let $\mathcal{D} \subset \{-1, 0, 1\}^{n \times n}$ be the set of all adjacency matrices for signed graphs G with the property that any set of vertex disjoint edges $\{e_1, \dots, e_k\}$ in G has size at most $k \leq \min\{n_1, n_2\}/4$. Notice that \mathcal{D} contains both $\mathcal{F}_\Sigma(\overline{C_n})$ and $\mathcal{F}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})$. For a given $\mathbf{X} \in \mathcal{D}$, we will define $\varphi(\mathbf{X}) = \mathbf{P}_{\sigma_X} \mathbf{X} \mathbf{P}_{\sigma_X}^T$ for some permutation σ_X . Since the graph of $\mathbf{P}_{\sigma_X} \mathbf{X} \mathbf{P}_{\sigma_X}^T$ is by definition isomorphic to \mathbf{X} , it follows that $\mathbf{P}_{\sigma_X} \mathbf{X} \mathbf{P}_{\sigma_X}^T \in \mathcal{D}$, thus φ maps \mathcal{D} into \mathcal{D} . So in order to define the mapping $\varphi(\mathbf{X})$, it suffices to define a function $\phi : \mathcal{D} \rightarrow S_n$ mapping into the symmetric group so that $\varphi(\mathbf{X}) = \mathbf{P}_{\phi(\mathbf{X})} \mathbf{X} \mathbf{P}_{\phi(\mathbf{X})}^T$.

For $i = 0, 1, 2, \dots, \min\{n_1, n_2\} - 1$, define the permutation $\sigma_i \in S_n$ as follows. For $j \in \{0, 1, \dots, n_1 - i\} \cup \{n_1 + 1, \dots, n - i\}$, we set $\sigma_i(j) = j$. If $i > 0$, then for each $0 \leq j < i$, we set $\sigma_i(n_1 - j) = n - j$ and $\sigma_i(n - j) = n_1 - j$. In other words, the function σ_i swaps the last $\max\{0, i - 1\}$ vertices before the splitting points n_1, n of the cycle. Notice that σ_i is an involution, so $\sigma_i(\sigma_i) = \text{id}$ and $\sigma_i = \sigma_i^{-1}$.

We now define our bijection φ . For $\mathbf{X} \in \mathcal{D}$, let $i(\mathbf{X})$ be the smallest value of $i \geq 0$ such that $\mathbf{X}_{n_1-i, n_1-i+1} = \mathbf{X}_{n-i, n-i+1} = \mathbf{X}_{n_1-i, n-i+1} = \mathbf{X}_{n_1-i, n-i+1} = 0$. Equivalently, $i(\mathbf{X})$ is the smallest value of $i \geq 0$ such that none of the four edges of the cycle $c_i = (v_{n_1-i}, v_{n_1-i+1}, v_{n-i}, v_{n-i+1})$ exist in \mathbf{X} . We then define $\varphi(\mathbf{X}) = \sigma_{i(\mathbf{X})} \mathbf{X}$, so that $\varphi(\mathbf{X}) = \mathbf{P}_{\sigma_{i(\mathbf{X})}} \mathbf{X} \mathbf{P}_{\sigma_{i(\mathbf{X})}}^T$. Note that if the maximum number of vertex disjoint edges in \mathbf{X} is at most $\min\{n_1, n_2\}/4$, then $i(\mathbf{X})$ must always exist and is at most $\min\{n_1, n_2\}/2 + 1$. This can be seen by the fact that for each i such that $i(\mathbf{X}) > i + 1$, there must be at least one edge with endpoints in the set $\{v_{n_1-i}, v_{n_1-i+1}, v_{n-i}, v_{n-i+1}\}$, thus for each $i \geq 0$ with $i < i(\mathbf{X})$ we can assign an edge e_i , such that $e_0, e_2, e_4, \dots, e_{i(\mathbf{X})-1}$ are vertex disjoint.

We must first argue that if $\mathbf{X} \in \mathcal{F}_\Sigma(\overline{C_n})$, then $\varphi(\mathbf{X}) \in \mathcal{F}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})$, namely that the function maps into the desired co-domain. To do this, we must show that for every (i, j) with $(\mathbf{P}_{\sigma_X} \mathbf{X} \mathbf{P}_{\sigma_X}^T)_{i,j} \neq 0$, we have $(\mathbf{P}_{\sigma_X} \mathbf{X} \mathbf{P}_{\sigma_X}^T)_{i,j} = (\mathbf{A}_{\overline{C_{n_1} \oplus C_{n_2}}})_{i,j}$. This is equivalent to showing that for any signed edge $e = (v_i, v_j) \in \mathbf{X}$, v_i, v_j are connected in C_n if and only if $v_{\sigma_X(i)}, v_{\sigma_X(j)}$ are connected in $C_{n_1} \oplus C_{n_2}$. In the proof of this fact, we will only use that $\mathbf{X} \in \mathcal{D}$.

So suppose v_i, v_j were connected in C_n , and wlog $j > i$. First suppose that $i \notin \{n, n_1\}$. Then we

have $j = i + 1$. Since $e = (v_i, v_j) \in \mathbf{X}$ is an edge of the subgraph, we know $i \notin \{n - i(\mathbf{X}), n_1 - i(\mathbf{X})\}$ by construction of $i(\mathbf{X})$. Thus $(v_{\sigma_X(i)}, v_{\sigma_X(j)}) = (v_{i'}, v_{i'+1})$ for some $i' \notin \{n_1, n\}$, which is always an edge of $C_{n_1} \oplus C_{n_2}$. If $i = n$, then $j = 1$, and we have $i(\mathbf{X}) > 0$, so $\sigma(i) = n_1$ and $\sigma(j) = 1$, and (v_{n_1}, v_1) is an edge of $C_{n_1} \oplus C_{n_2}$. Similarly, if $i = n_1$, then $j = n_1 + 1$, and since again necessarily $i(\mathbf{X}) > 0$ we have $\sigma(i) = n$, $\sigma(j) = n_1 + 1$, and (v_n, v_{n_1+1}) is an edge of $C_{n_1} \oplus C_{n_2}$. We now consider the case where $(v_i, v_j) \in \mathbf{X}$ is not an edge in C_n . Suppose for the sake of contradiction that $(v_{\sigma_X(i)}, v_{\sigma_X(j)})$ is an edge in $C_{n_1} \oplus C_{n_2}$. WLOG, i, j are in the first cycle C_{n_1} . We can write $\sigma_X(i) = i'$, $\sigma_X(j) = i' + 1$ for some $i' \in \{1, 2, \dots, n_1\}$, where $i' + 1$ is interpreted as 1 if $i' = n_1$. If $i' \leq i(\mathbf{X}) - 1$, then both $i' = i$ and $i' + 1 = i + 1 = j$, but (v_i, v_{i+1}) is also connected in C_n . If $i' \geq i(\mathbf{X}) + 1$, then $i' = i + n_2$ and $i' + 1 = i + n_2 + 1$ (where $i + n_2 + 1$ is interpreted modularly as 1 if $i = n_1$), and again v_{i+n_2} and v_{i+n_2+1} are connected in C_n . Finally, if $i' = i(\mathbf{X})$, then $i = i'$ and $j = i + n_2 + 1$, but then we cannot have $(v_i, v_j) \in \mathbf{X}$ by construction of $i(\mathbf{X})$, which completes the of the claim that φ maps $\mathcal{F}_\Sigma(\overline{C_n})$ into $\mathcal{F}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})$.

We now show that φ is injective. To do this, we show that $\varphi(\varphi(\mathbf{X})) = \mathbf{X}$ for any $\mathbf{X} \in \mathcal{D}$ – namely that φ is an involution on \mathcal{D} . This can be seen by showing that we always have $i(\mathbf{X}) = i(\varphi(\mathbf{X}))$. To see this, observe that $i(\mathbf{X})$ is defined as the first $i \geq 0$ such that none of the four edges of the cycle $c_i = (v_{n_1-i}, v_{n_1-i+1}, v_{n-i}, v_{n-i+1})$ exist in \mathbf{X} . Thus it suffices to show that for each $\min\{n_1, n_2\} - 1 > i \geq 0$, the number of edges in c_i is preserved after permuting the vertices by $\sigma_{i(\mathbf{X})}$. To see this, note that if $i(\mathbf{X}) > i$, then $(\sigma_{i(\mathbf{X})}(v_{n_1-i}), \sigma_{i(\mathbf{X})}(v_{n_1-i+1}), \sigma_{i(\mathbf{X})}(v_{n-i}), \sigma_{i(\mathbf{X})}(v_{n-i+1})) = (v_{n-i}, v_{n-i+1}, v_{n_1-i}, v_{n_1-i+1})$, which is the same cycle. If $i(\mathbf{X}) < i$, then $\sigma_{i(\mathbf{X})}$ does not move any of the vertices in c_i . Finally, if $i(\mathbf{X}) = i$, then $(\sigma_{i(\mathbf{X})}(v_{n_1-i}), \sigma_{i(\mathbf{X})}(v_{n_1-i+1}), \sigma_{i(\mathbf{X})}(v_{n-i}), \sigma_{i(\mathbf{X})}(v_{n-i+1})) = (v_{n_1-i}, v_{n_1-i+1}, v_{n-i}, v_{n_1-i+1})$, which again is the same cycle c_i (just with the ordering of the vertices reversed). So $\varphi(\varphi(\mathbf{X})) = \mathbf{X}$ for any $\mathbf{X} \in \mathcal{D}$, so in particular $\varphi : \mathcal{F}_\Sigma(\overline{C_n}) \rightarrow \mathcal{F}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})$ is injective.

To show surjectivity, it suffices to show that if $\mathbf{X} \in \mathcal{F}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})$ then $\varphi(\mathbf{X}) \in \mathcal{F}_\Sigma(\overline{C_n})$. Namely, that φ can also be defined as a valid function $\varphi : \mathcal{F}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}}) \rightarrow \mathcal{F}_\Sigma(\overline{C_n})$. Again, this is equivalent to showing that for any signed edge $e = (v_i, v_j) \in \mathbf{X}$, v_i, v_j are connected in $C_{n_1} \oplus C_{n_2}$ if and only if $v_{\sigma_X(i)}, v_{\sigma_X(j)}$ are connected in C_n . Since σ_X is an involution, this is the same as asking that for any signed edge $e = (v_i, v_j) \in \mathbf{X}$, $v_{\sigma_X(\sigma_X(i))}, v_{\sigma_X(\sigma_X(j))}$ are connected in $C_{n_1} \oplus C_{n_2}$ if and only if $v_{\sigma_X(i)}, v_{\sigma_X(j)}$ are connected in C_n . Setting $i' = \sigma_X(i)$, $j' = \sigma_X(j)$, this states that for all signed edges $(v_{i'}, v_{j'}) \in \mathbf{P}_{\sigma_X} \mathbf{X} \mathbf{P}_{\sigma_X}^T = \mathbf{Y} \in \mathcal{D}$, we have that $v_{i'}, v_{j'}$ are connected in C_n if and only if $v_{\sigma_X(i')}, v_{\sigma_X(j')}$ are connected in $C_{n_1} \oplus C_{n_2}$. But as shown above, we have that $i(\mathbf{X}) = i(\varphi(\mathbf{X}))$, so $\sigma_X = \sigma_Y$, and then this fact was already proven above for any $\mathbf{Y} \in \mathcal{D}$, which completes the proof. ■

Now for any signed graph Σ on n vertices, let \mathbf{A}_Σ be its adjacency matrix. Note that we can equivalently define via $\mathcal{F}_\Sigma(\overline{C_n}) = \{H \subseteq \overline{C_n} \mid \mathbf{P}_\sigma \mathbf{A}_\Sigma \mathbf{P}_\sigma^T = \mathbf{A}_H, \sigma \in S_n\}$. Here $H \subseteq \overline{C_n}$ means H is a subgraph of $\overline{C_n}$. On the other hand, we may be interested in the potentially much larger set of all possible permutations σ such that $\mathbf{P}_\sigma \mathbf{A}_\Sigma \mathbf{P}_\sigma^T = \mathbf{A}_H$ for some $H \subset \overline{C_n}$. So define $\mathcal{H}_\Sigma(\overline{C_n}) = \{\sigma \mid \mathbf{P}_\sigma \mathbf{A}_\Sigma \mathbf{P}_\sigma^T = \mathbf{A}_H, H \subset \overline{C_n}, \sigma \in S_n\}$. It is not difficult to show that $|\mathcal{H}_\Sigma(\overline{C_n})| = |\text{Aut}(\Sigma)| |\mathcal{F}_\Sigma(\overline{C_n})|$, where $\text{Aut}(\Sigma)$ is the set of (signed) graph automorphisms of Σ .

Fact 5.14. *We have $|\mathcal{H}_\Sigma(\overline{C_n})| = |\text{Aut}(\Sigma)| |\mathcal{F}_\Sigma(\overline{C_n})|$.*

Proof. Fix any $H \subset \overline{C_n}$ such that $\mathbf{P}_\sigma \mathbf{A}_\Sigma \mathbf{P}_\sigma^T = \mathbf{A}_H$ for some $\sigma \in S_n$. We show that there are exactly $|\text{Aut}(\Sigma)|$ elements $\sigma' \in S_n$ such that $\mathbf{P}_{\sigma'} \mathbf{A}_\Sigma \mathbf{P}_{\sigma'}^T = \mathbf{A}_H$. By definition, $\text{Aut}(\Sigma)$ is the set of permutations

$\pi \in S_n$ with $\mathbf{P}_\sigma \mathbf{A}_\Sigma \mathbf{P}_\sigma^T = \mathbf{A}_\Sigma$. For every $\pi \in \text{Aut}(\Sigma)$, we have $\mathbf{P}_\sigma \mathbf{P}_\pi \mathbf{A}_\Sigma \mathbf{P}_\pi \mathbf{P}_\sigma^T = \mathbf{P}_{\sigma \circ \pi} \mathbf{A}_\Sigma \mathbf{P}_{\sigma \circ \pi}^T = \mathbf{A}_H$, and moreover the set of elements $|\{\sigma \circ \pi \mid \pi \in \text{Aut}(\Sigma)\}| = |\text{Aut}(\Sigma)|$ since S_n is a group. Now suppose we have some $\lambda \in S_n$ such that $\mathbf{P}_\lambda \mathbf{A}_\Sigma \mathbf{P}_\lambda^T = \mathbf{A}_H$ and $\lambda \notin \{\sigma \circ \pi \mid \pi \in \text{Aut}(\Sigma)\}$. Then $\mathbf{P}_\sigma \mathbf{A}_\Sigma \mathbf{P}_\sigma^T = \mathbf{P}_\lambda \mathbf{A}_\Sigma \mathbf{P}_\lambda^T$, so $\mathbf{P}_{\sigma^{-1} \circ \lambda} \mathbf{A}_\Sigma \mathbf{P}_{\sigma^{-1} \circ \lambda}^T = \mathbf{A}_\Sigma$, which by definition implies that $\sigma^{-1} \circ \lambda = x$ for some $x \in \text{Aut}(\Sigma)$. Thus $\lambda = \sigma \circ x \in \{\sigma \circ \pi \mid \pi \in \text{Aut}(\Sigma)\}$, which is a contradiction. ■

Corollary 5.15. *Fix any $n = n_1 + n_2$. Fix any simple graph $|\Sigma|$, such that any set of vertex disjoint edges $\{e_1, \dots, e_k\}$ in $|\Sigma|$ has size at most $k \leq \min\{n_1, n_2\}/4$, and let $\Sigma = (|\Sigma|, \sigma)$ be any signing of $|\Sigma|$. Let $\mathcal{F}_\Sigma(\overline{C_n})$ denote the set of subgraphs of C_n isomorphic to $|\Sigma|$, and similarly define $\mathcal{F}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})$. Then we have*

$$|\mathcal{H}_\Sigma(\overline{C_n})| = |\mathcal{H}_\Sigma(\overline{C_{n_1} \oplus C_{n_2}})|$$

5.3 Lower Bounds for Schatten, Ky-Fan, and Tail Error Testing

In this section, we demonstrate how our construction of subgraph equivalent matrices with gaps in their spectrum result in lower bounds for a number of other spectral testing problems via Lemma 5.8. We begin by proving a lower bound for testing Schatten norms. To do this, we must first demonstrate that there is a gap in the Schatten 1 norm between a cycle and the union of two disjoint cycles.

Fact 5.16 (Theorem 1 of [Kna09]). *Fix any $a, b, n \in \mathbb{R}$ with $\sin(b/2) \neq 0$. Then we have*

$$\sum_{k=0}^{n-1} \cos(a + kb) = \frac{\sin(\frac{nb}{2})}{\sin(\frac{b}{2})} \cos\left(a + \frac{(n-1)b}{2}\right)$$

Proposition 5.17. *Fix any $d \geq 6$ be any integer divisible by 4. Then*

$$\|C_d\|_{S_1} = 4 \cdot \frac{\cos(\pi/d)}{\sin(\pi/d)}$$

Proof. By [Chu96], for any $d \geq 3$ the eigenvalues of C_d are given by $2 \cdot \cos(\frac{2\pi j}{d})$ for $j = 0, 1, \dots, d-1$. Let $a_1 = \lfloor d/4 \rfloor, a_2 = \lfloor 3d/4 \rfloor, a_3 = d - a_2 - 1$.

$$\begin{aligned} \|C_d\|_1 &= 2 \sum_{j=0}^{d-1} \left| \cos\left(\frac{2\pi j}{d}\right) \right| \\ &= 2 \left(\sum_{j=0}^{a_1} \cos\left(\frac{2\pi j}{d}\right) - \sum_{j=a_1+1}^{a_2} \cos\left(\frac{2\pi j}{d}\right) + \sum_{j=a_2+1}^{d-1} \cos\left(\frac{2\pi j}{d}\right) \right) \\ &= 2 \left(\sum_{j=-a_3}^{a_1} \cos\left(\frac{2\pi j}{d}\right) - \sum_{j=a_1+1}^{a_2} \cos\left(\frac{2\pi j}{d}\right) \right) \end{aligned} \tag{34}$$

We analyze each term in the above via Fact 5.16. Firstly:

$$\begin{aligned}
\sum_{j=-a_3}^{a_1} \cos\left(\frac{2\pi j}{d}\right) &= \sum_{j=0}^{a_1+a_3} \cos\left(\frac{2\pi j}{d} - \frac{2\pi a_3}{d}\right) \\
&= \frac{\sin((a_1 + a_3 + 1)\pi/d)}{\sin(\pi/d)} \cos\left(\frac{(a_1 + a_3)\pi}{d} - \frac{2\pi a_3}{d}\right)
\end{aligned} \tag{35}$$

Note that if d is divisible by 4, the above becomes $2 \cos(\pi/d)/\sin(\pi/d)$. Next, for the second term, we have

$$\begin{aligned}
\sum_{j=a_1+1}^{a_2} \cos\left(\frac{2\pi j}{d}\right) &= \sum_{j=0}^{a_2-a_1-1} \cos\left(\frac{2\pi j}{d} - \frac{2\pi(a_1+1)}{d}\right) \\
&= \frac{\sin((a_2 - a_1)\pi/d)}{\sin(\pi/d)} \cos\left(\frac{(a_2 - a_1 - 1)\pi}{d} - \frac{2\pi(a_1+1)}{d}\right)
\end{aligned} \tag{36}$$

Again, note that if d is divisible by 4, the above becomes $2 \cos(\pi/d)/\sin(\pi/d)$. Putting these two equations together, we have that

$$\|C_d\|_1 = 4 \cdot \frac{\cos(\pi/d)}{\sin(\pi/d)}$$

■

Proposition 5.18. *Fix any d larger than some constant. Then we have*

$$\left| \|C_{8d}\|_{S_1} - \|C_{4d} \oplus C_{4d}\|_{S_1} \right| \gtrsim \frac{1}{d^3}$$

Proof. By the prior Lemma, we have $\|C_d\|_{S_1} = 4 \cot(\pi/d)$ for any d divisible by 4. Thus using the Taylor expansion of cotangent, we have

$$\|C_{8d}\|_{S_1} = 4 \left(\frac{8d}{\pi} + \frac{\pi}{24d} + \frac{\pi^3}{45 \cdot 512 \cdot d^3} + O(1/d^5) \right) \tag{37}$$

and

$$\begin{aligned}
\|C_{4d} \oplus C_{4d}\|_{S_1} &= 2\|C_{4d}\|_{S_1} \\
&= 4 \left(\frac{8d}{\pi} + \frac{\pi}{24d} + \frac{\pi^3}{45 \cdot 128 \cdot d^3} + O(1/d^5) \right)
\end{aligned} \tag{38}$$

Thus

$$\left| \|C_{8d}\|_{S_1} - \|C_{4d} \oplus C_{4d}\|_{S_1} \right| \gtrsim \frac{1}{d^3} \tag{39}$$

■

Theorem 5.19. *Fix any $\frac{1}{\sqrt{n}} \leq \epsilon \leq 1$. Then given $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$, any non-adaptive sampling algorithm which distinguishes between the cases*

1. $\|\mathbf{A}\|_{S_1} > \epsilon_0 n^{1.5}$

$$2. \|\mathbf{A}\|_{S_1} < \epsilon_0 n^{1.5} - \epsilon n^{1.5}$$

with probability at least $3/4$, where $\epsilon_0 = \tilde{\Theta}(\epsilon)$, must query at least $\tilde{\Omega}(1/\epsilon^4)$ entries of \mathbf{A} .

Proof. We use the hard instance $\mathcal{D}_1, \mathcal{D}_2$ as earlier. Set $k = C \frac{1}{\epsilon^2 \log^9(1/\epsilon)}$, $t = \log k$, and $d = 4k$, and $m = n/(dk)$. We instantiate the matrices $(\mathbf{B}, \mathbf{D}, \mathbf{Z})$ in the hard instance via $\mathbf{B} = \mathbf{C}_{2d}$, $\mathbf{D} = \mathbf{C}_d \oplus \mathbf{C}_d$, and let $\mathbf{Z} = \delta_{i,j}$ for $i \leq j$, where $\delta_{i,j} \in \{-1, 1\}$ are i.i.d. Bernoulli random variables, so that $\mathbf{Z} \in \mathbb{R}^{m \times m}$ is a symmetric random Bernoulli matrix. Using the fact that $\|\mathbf{Z}\|_2 \leq O(\sqrt{n})$ with high probability [Ver10], along with the fact that $\|\mathbf{Z}\|_F^2 = n^2$ deterministically, we have that $\|\mathbf{Z}\|_{S_1} > C_1 m^{1.5}$ with non-zero probability for some constant $C_1 > 0$, as the former two facts imply that \mathbf{Z} has $\Omega(n)$ eigenvalues with magnitude $\Theta(\sqrt{n})$. Thus, we can deterministically fix \mathbf{Z} to be such a matrix with $\{1, -1\}$ entries such that $\|\mathbf{Z}\|_{S_1} \geq C_1 m^{1.5}$. Given this, we have $\|\tilde{\mathbf{B}}\|_{S_1} = \|\mathbf{B} \otimes \mathbf{Z}\|_{S_1} = \|\mathbf{B}\|_{S_1} \cdot \|\mathbf{Z}\|_{S_1}$, and so by Proposition 5.18, we have

$$\left| \|\tilde{\mathbf{B}}\|_{S_1} - \|\tilde{\mathbf{D}}\|_{S_1} \right| \geq C_0 \frac{m^{1.5}}{d^3}$$

for some absolute constant $C_0 \geq 0$. Note also that we have $\|\mathbf{B}\|_{S_1} > \Omega(d)$, where we use the fact that a constant fraction of the eigenvalues $2 \cdot \cos(\frac{2\pi j}{d})$ for $j = 0, 1, \dots, d-1$ of \mathbf{B} are $\Omega(1)$. Thus we have $\|\tilde{\mathbf{B}}\|_{S_1} = dm^{1.5}$.

Now by Proposition 5.10, we obtain that $\mathbf{B} \cong_{\mathcal{U}_{2d}^t, S_{2d}} \mathbf{D}$, and thus $\tilde{\mathbf{B}} = \mathbf{B} \otimes \mathbf{Z} \cong_{\mathcal{U}_{2d}^t, \Gamma_{2d, 2dm}} \mathbf{D} \otimes \mathbf{Z} = \tilde{\mathbf{D}}$ by Lemma 5.7. Thus by Lemma 5.8, we have that distinguishing \mathcal{D}_1 from \mathcal{D}_2 requires $\Omega(k^2) = \tilde{\Omega}(1/\epsilon^4)$ samples for any non-adaptive algorithm. It suffices then to show that if $\mathbf{A}_1 \sim \mathcal{D}_1$ and $\mathbf{A}_2 \sim \mathcal{D}_2$, then we have the desired gap in Schatten norms. We have

$$\begin{aligned} \left| \|\mathbf{A}_1\|_{S_1} - \|\mathbf{A}_2\|_{S_1} \right| &\geq \sum_{i=1}^k C_0 \frac{m^{1.5}}{d^3} \\ &\geq C_0 \frac{n^{1.5}}{d^{4.5} k^{1/2}} \\ &\geq \epsilon n^{1.5} \end{aligned} \tag{40}$$

Where the last inequality follows setting C large enough, and assuming that $1/\epsilon$ is larger than some constant as in Theorem 5.12. Again, if $1/\epsilon$ is not larger than some constant, a $\Omega(1)$ lower bound always applies, since an algorithm must read at least one entry of the matrix to have any advantage. Now note that we also have $\|\mathbf{A}_1\|_{S_1} = k \|\tilde{\mathbf{B}}\|_{S_1} = kdm^{1.5} = n^{1.5}/\sqrt{dk} = \tilde{\Theta}(\epsilon n^{1.5})$ as desired. To complete the proof, we can scale down all the entries of the input matrix by $1/2$, which results in the required bounded entry property, and only changes the gap by a constant factor. ■

We now present our lower bound for testing Ky-Fan norms. Recall that for a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $1 \leq s \leq n$, the Ky-Fan s norm is defined as $\|\mathbf{A}\|_{KF(s)} = \sum_{i=1}^s \sigma_i(\mathbf{A})$, where $\sigma_i(\mathbf{A})$ is the i -th singular value of \mathbf{A} .

Theorem 5.20. *Fix any $1 \leq s \leq n/(\text{poly log } n)$. Then there exists a fixed constant $c > 0$ such that given $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$, any non-adaptive sampling algorithm which distinguishes between the cases*

$$1. \|\mathbf{A}\|_{KF(s)} > \frac{c}{\log(s)} n$$

$$2. \|\mathbf{A}\|_{KF(s)} < (1 - \epsilon_0) \frac{c}{\log(s)} n$$

with probability at least $3/4$, where $\epsilon_0 = \Theta(1/\log^2(s))$, must query at least $\tilde{\Omega}(s^2)$ entries of \mathbf{A} .¹²

Proof. The proof is nearly the same as the usage of the hard instance in Theorem 5.12. Set $k = s$, and let $d_0 = \Theta(\log s)$ and $d = 2d_0 + 1$. We apply Lemma 5.8 with the hard instance as instantiated with $\mathbf{Z} = \mathbf{1}^{m \times m}$, and the matrices $\mathbf{B} = 1/4(\mathbf{A}_{C_d} - 2\mathbf{I}_d)$ and $\mathbf{D} = 1/4(\mathbf{A}_{C_{d_0} \oplus C_{d_0+1}} - 2\mathbf{I}_d)$. Notice that since the eigenvalues of C_d are given by $2 \cdot \cos(\frac{2\pi j}{d})$ for $j = 0, 1, \dots, d-1$ [Chu96], we have $\lambda_{\min}(\mathbf{A}_{C_d}) = -2 \cos(\frac{2\pi d_0}{2d_0+1}) = -2 + \Theta(1/\log^2(1/\epsilon))$, $\lambda_{\min}(\mathbf{A}_{C_{d_0} \oplus C_{d_0+1}}) = -2$, and $\lambda_{\max}(\mathbf{A}_{C_d}) = \lambda_{\max}(\mathbf{A}_{C_{d_0} \oplus C_{d_0+1}}) = 2$. Thus $\|\mathbf{D}\|_2 = 4$ and $\|\mathbf{B}\|_2 = 4 - \Theta(1/d^2)$, and moreover $\|\mathbf{D} \otimes \mathbf{Z}\|_2 = 4m$ $\|\mathbf{B} \otimes \mathbf{Z}\|_2 = 4m(1 - \Theta(1/\log^2(1/\epsilon)))$. Thus if $\mathbf{A}_1 \sim \mathcal{D}_1$, we have $\|\mathbf{A}_1\|_{KF(s)} > \sum_{i=1}^k 4m = 4km$, and $\|\mathbf{A}_2\|_{KF(s)} < 4km(1 - \Theta(1/\log^2(1/\epsilon)))$. The proof then follows from the $\Omega(k^2)$ lower bound for this hard instance via Lemma 5.8. ■

We now present our lower bound for testing the magnitude of the s -tail $\|\mathbf{A} - \mathbf{A}_s\|_F^2$, where $\mathbf{A}_s = \mathbf{U}_{\Sigma_s} \mathbf{V}^T$ is the truncated SVD (the best rank- s approximation to \mathbf{A}). Note that $\|\mathbf{A} - \mathbf{A}_s\|_F^2 = \sum_{j>s} \sigma_j^2(\mathbf{A})$.

Theorem 5.21. Fix any $1 \leq s \leq n/(\text{poly } \log n)$. Then there exists a fixed constant $c > 0$ (independent of ϵ), such that given $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$, any non-adaptive sampling algorithm which distinguishes between the cases

1. $\|\mathbf{A} - \mathbf{A}_s\|_F^2 > \frac{c}{\log(s)} \cdot \frac{n^2}{s}$
2. $\|\mathbf{A} - \mathbf{A}_s\|_F^2 < (1 - \epsilon_0) \cdot \frac{c}{\log(s)} \cdot \frac{n^2}{s}$

with probability at least $3/4$, where $\epsilon_0 = \tilde{\Theta}(1)$, must query at least $\tilde{\Omega}(s^2)$ entries of \mathbf{A} .

Proof. We set $s = k$, and use the same hard instance as in Theorem 5.20 above. Note that if $\mathcal{D}_1, \mathcal{D}_2$ are defined as in Theorem 5.20, if $\mathbf{A}_1 \sim \mathcal{D}_1, \mathbf{A}_s \sim \mathcal{D}_s$, we have $\sum_{i=1}^s \lambda_i(\mathbf{A}_1) = s(4m)^2 = 16n^2/(sd^2)$ and $\sum_{i=1}^s \lambda_i(\mathbf{A}_2) = 16n^2/(sd^2)(1 - \Theta(1/\log^2 s))$. Now note that $\|\mathbf{A}_1\|_F^2 = \|\mathbf{A}_2\|_F^2 = kdm^2 = n^2/(dk) = n^2/(ds)$, using that each of the single cycle and union of two smaller cycles has d edges, so the Frobenius norm of each block is dm^2 in both cases. Using that $d = \Theta(\log s)$, we have that if $\|(\mathbf{A}_1) - (\mathbf{A}_1)s\|_F^2 > n^2/(ds) - 16n^2/(sd^2) = c \frac{n^2}{s \log(s)}$ for some constant $c > 0$, and $\|(\mathbf{A}_2) - (\mathbf{A}_2)s\|_F^2 > n^2/(ds) - 16n^2/(sd^2)(1 - \Theta(1/\log^2 s)) = c \frac{n^2}{s \log(s)} + \tilde{\Theta}(\frac{n^2}{s})$, which completes the proof after applying Lemma 5.8. ■

5.4 Lower Bound For Estimating Ky-Fan of $\mathbf{A}\mathbf{A}^T$ via Submatrices

In this section, we demonstrate a $\Omega(1/\epsilon^4)$ query lower bound for algorithms which estimate the quantity $\sum_{i=1}^k \sigma_i^2(\mathbf{A}) = \|\mathbf{A}\mathbf{A}^T\|_{KF(k)}$ for any $k \geq 1$ by querying a sub-matrix. The following lemma as a special case states that for $\epsilon = \Theta(1/\sqrt{n})$, additive ϵn^2 approximation of $\|\mathbf{A}\mathbf{A}^T\|_{KF(k)}$ requires one to read the entire matrix \mathbf{A} .

¹² $\tilde{\Omega}$ hides $\log(s)$ factors here.

Lemma 5.22. Fix any $1 \leq k \leq n$, and fix any $\frac{100}{\sqrt{n}} \leq \epsilon \leq 1/4$. Any algorithm that queries a submatrix $\mathbf{A}_{S \times T}$ of $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$ and distinguishes with probability at least $4/5$ between the case that either:

- $\sum_{i=1}^k \sigma_i^2(\mathbf{A}) > n^2/2 + \epsilon n^2$.
- $\sum_{i=1}^k \sigma_i^2(\mathbf{A}) \leq n^2/2$

must make $|S| \cdot |T| = \Omega(1/\epsilon^4)$ queries to the matrix \mathbf{A} .

Proof. We design two distributions $\mathcal{D}_1, \mathcal{D}_2$. If $\mathbf{A}_1 \sim \mathcal{D}_1$, we independently set each row of \mathbf{A}_1 equal to the all 1's vector with probability $p_1 = 1/2 + 2\epsilon$, and then return either $\mathbf{A} = \mathbf{A}_1$ or $\mathbf{A} = \mathbf{A}_1^T$ with equal probability. If $\mathbf{A}_2 \sim \mathcal{D}_2$, we independently set each row of \mathbf{A}_2 equal to the all 1's vector with probability $p_2 = 1/2 - 2\epsilon$, and then return either $\mathbf{A} = \mathbf{A}_2$ or $\mathbf{A} = \mathbf{A}_2^T$ with equal probability. Our hard instance then draws $\mathbf{A} \sim \frac{\mathcal{D}_1 + \mathcal{D}_2}{2}$ from the mixture. First note that in both cases, we have $\|\mathbf{A}\|_2^2 = \|\mathbf{A}\|_F^2 = \sum_{i=1}^k \sigma_i^2(\mathbf{A})$, since the matrix is rank 1. Since $\frac{100}{\sqrt{n}} \geq \epsilon$, by Chernoff bounds, we have that if $\mathbf{A}_1 \sim \mathcal{D}_1$ then $\sum_{i=1}^k \sigma_i^2(\mathbf{A}) > n^2/2 + \epsilon n^2$ with probability at least $99/100$. Similarly, we have that if $\mathbf{A}_2 \sim \mathcal{D}_2$ then $\sum_{i=1}^k \sigma_i^2(\mathbf{A}) \leq n^2/2$ with probability at least $99/100$.

Now suppose that such an algorithm sampling $|S| \cdot |T| < \frac{c^2}{\epsilon^4}$ entries exists, for some constant $c > 0$. Then by Yao's min-max principle, there is a fixed submatrix $S, T \subset [n]$ such that, with probability $9/10$ over the distribution $\frac{\mathcal{D}_1 + \mathcal{D}_2}{2}$, the algorithm correctly distinguishes \mathcal{D}_1 from \mathcal{D}_2 given only $\mathbf{A}_{S \times T}$. Suppose WLOG that $|S| \leq \frac{c}{\epsilon^2}$. Then consider the case only when \mathbf{A}_1 or \mathbf{A}_2 is returned by either of the distributions, and not their transpose, which occurs with probability at least $1/2$. Then $\mathbf{A}_{S \times T}$ is just a set of $|S|$ rows, each of which are either all 0's or all 1's. Moreover, each row is set to being the all 1's row independently with probability p_1 in the case of \mathcal{D}_1 , and p_2 in the case of \mathcal{D}_2 . Thus, by Independence across rows, the behavior of the algorithm can be assumed to depend only on the number of rows which are set to 1. Thus, in the case of \mathcal{D}_1 the algorithm receives $X_1 \sim \text{Bin}(|S|, p_1)$ and in \mathcal{D}_2 the algorithm receives $X_2 \sim \text{Bin}(|S|, p_2)$. Then if $d_{TV}(X_1, X_2)$ is the total variational distance between X_1, X_2 , then by Equation 2.15 of [AJ06], assuming that $\epsilon\sqrt{|S|}$ is smaller than some constant (which can be obtained by setting c small enough), we have

$$d_{TV}(X_1, X_2) \leq O(\epsilon\sqrt{|S|})$$

Which is at most $1/100$ for c a small enough constant. Thus any algorithm can correctly distinguish these two distributions with advantage at most $1/100$. Since we restricted our attention to the event when rows were set and not columns, and since we conditioned on the gap between the norms which held with probability $99/100$, it follows that the algorithm distinguishes \mathcal{D}_1 from \mathcal{D}_2 with probability at most $1/2 + 1/4 + (2/100) < 4/5$, which completes the proof. ■

6 Conclusion

In this work, we gave an optimal (up to $\log(1/\epsilon)$ factors) algorithm for testing if a matrix was PSD, or was far in spectral norm distance from the PSD cone. In addition, we gave a query efficient algorithm for testing if a matrix was PSD, or was ϵn^2 far from the PSD-cone in ℓ_2^2 distance. Furthermore, we established a new technique for proving lower bounds based on designing “subgraph-equivalent” matrices. We believe that this technique is quite general, as shown by its immediate application to

lower bounds for the Schatten-1 norm, Ky-Fan norm, and tail error testing. Our construction could also likely be useful for proving lower bounds against testing of *graph properties*, which is a well studied area [Gol10]. We pose the open problem to design (or demonstrate the non-existence of) additional subgraph-equivalent matrices beyond the cycle graph construction utilized in this work, which have gaps in their spectral or graph-theoretic properties.

Additionally, we pose the open problem of determining the exact non-adaptive query complexity of PSD testing with ℓ_2^2 gap. As discussed in Section 1.1, there appear to be several key barriers to improving the complexity beyond $O(1/\epsilon^4)$. Indeed, it seems that perhaps the main tool that is lacking is a concentration inequality for the eigenvalues of random principal submatrices. Since most such decay results apply only to norms [Tro08, RV07], progress in this direction would likely result in important insights into eigenvalues of random matrices.

Finally, we note that the complexity of the testing problems for several matrix norms, specifically the Schatten p and Ky-Fan norms, are still open in the bounded entry model. In particular, for the Schatten 1 norm, to the best of our knowledge no non-trivial algorithms exist even for estimation with additive error $\Theta(n^{1.5})$, thus any improvements would be quite interesting.

Acknowledgement

We thank Erik Waingarten for many useful suggestions, being closely involved in the early stages of this project, and for feedback on early drafts of this manuscript. We also thank Roie Levin, Ryan O'Donnell, Pedro Paredes, Nicolas Resch, and Goran Zuzic for illuminating discussions related to this project.

References

- [ACK⁺16] Alexandr Andoni, Jiecao Chen, Robert Krauthgamer, Bo Qin, David P Woodruff, and Qin Zhang. On sketching quadratic forms. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 311–319. ACM, 2016. 8
- [AHK05] Sanjeev Arora, Elad Hazan, and Satyen Kale. Fast algorithms for approximate semidefinite programming using the multiplicative weights update method. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 339–348. IEEE, 2005. 1
- [AJ06] José A Adell and Pedro Jodrá. Exact kolmogorov and total variation distances between some familiar discrete distributions. *Journal of Inequalities and Applications*, 2006(1):64307, 2006. 53
- [AMS96] Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 20–29. ACM, 1996. 8
- [AN13] Alexandr Andoni and Huy L. Nguyen. Eigenvalues of a matrix in the streaming model. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 1729–1737. Society for Industrial and Applied Mathematics, 2013. 8

- [Aro98] Sanjeev Arora. Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems. *Journal of the ACM (JACM)*, 45(5):753–782, 1998. [7](#)
- [ARV09] Sanjeev Arora, Satish Rao, and Umesh Vazirani. Expander flows, geometric embeddings and graph partitioning. *Journal of the ACM (JACM)*, 56(2):1–37, 2009. [1](#)
- [ÄzÖÖ93] M Necati Ä-zisik, M Necati Özısıık, and M Necati Özısıık. *Heat conduction*. John Wiley & Sons, 1993. [1](#)
- [BBG18] Siddharth Barman, Arnab Bhattacharyya, and Suprovat Ghoshal. Testing sparsity over known and unknown bases. In *International Conference on Machine Learning*, pages 491–500, 2018. [7](#)
- [BCK⁺18] Vladimir Braverman, Stephen Chestnut, Robert Krauthgamer, Yi Li, David Woodruff, and Lin Yang. Matrix norms in data streams: Faster, multi-pass and row-order. In *International Conference on Machine Learning*, pages 649–658, 2018. [8](#)
- [BCW19] Ainesh Bakshi, Nadiia Chepurko, and David P Woodruff. Robust and sample optimal algorithms for psd low-rank approximation. *arXiv preprint arXiv:1912.04177*, 2019. [1](#), [7](#)
- [BH11] Maria-Florina Balcan and Nicholas JA Harvey. Learning submodular functions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 793–802, 2011. [7](#)
- [BJWY20] Omri Ben-Eliezer, Rajesh Jayaram, David P Woodruff, and Eylon Yogev. A framework for adversarially robust streaming algorithms. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 63–80, 2020. [8](#)
- [BKKS19] Vladimir Braverman, Robert Krauthgamer, Aditya Krishnan, and Roi Sinoff. Schatten norms in matrix streams: Hello sparsity, goodbye dimension. *arXiv preprint arXiv:1907.05457*, 2019. [8](#)
- [BLWZ19] Maria-Florina Balcan, Yi Li, David P Woodruff, and Hongyang Zhang. Testing matrix rank, optimally. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 727–746. SIAM, 2019. [1](#), [4](#), [5](#), [7](#)
- [BSST13] Joshua Batson, Daniel A Spielman, Nikhil Srivastava, and Shang-Hua Teng. Spectral sparsification of graphs: theory and algorithms. *Communications of the ACM*, 56(8):87–94, 2013. [8](#)
- [BVKS19] Jess Banks, Jorge Garza Vargas, Archit Kulkarni, and Nikhil Srivastava. Pseudospectral shattering, the sign function, and diagonalization in nearly matrix multiplication time. *arXiv preprint arXiv:1912.08805*, 2019. [31](#), [38](#)
- [BW18] Ainesh Bakshi and David Woodruff. Sublinear time low-rank approximation of distance matrices. In *Advances in Neural Information Processing Systems*, pages 3782–3792, 2018. [7](#), [8](#)

- [BZ16] Maria-Florina F Balcan and Hongyang Zhang. Noise-tolerant life-long matrix completion via adaptive sampling. In *Advances in Neural Information Processing Systems*, pages 2955–2963, 2016. [7](#)
- [Chu96] Fan RK Chung. Lectures on spectral graph theory. *Lecture Notes*, 1996. [43](#), [49](#), [52](#)
- [CRT06] Emmanuel J Candes, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 59(8):1207–1223, 2006. [8](#)
- [CW17] Kenneth L Clarkson and David P Woodruff. Low-rank approximation and regression in input sparsity time. *Journal of the ACM (JACM)*, 63(6):54, 2017. [8](#)
- [Dat10] Jon Dattorro. *Convex optimization & Euclidean distance geometry*. Lulu. com, 2010. [7](#)
- [DDHK07] James Demmel, Ioana Dumitriu, Olga Holtz, and Robert Kleinberg. Fast matrix multiplication is stable. *Numerische Mathematik*, 106(2):199–224, 2007. [31](#), [38](#)
- [DJS⁺19] Huaian Diao, Rajesh Jayaram, Zhao Song, Wen Sun, and David Woodruff. Optimal sketching for kronecker product regression and low rank approximation. In *Advances in Neural Information Processing Systems*, pages 4737–4748, 2019. [8](#)
- [DK19] Ilias Diakonikolas and Daniel M Kane. Recent advances in algorithmic high-dimensional robust statistics. *arXiv preprint arXiv:1911.05911*, 2019. [1](#)
- [DL09] Michel Marie Deza and Monique Laurent. *Geometry of cuts and metrics*, volume 15. Springer, 2009. [1](#)
- [EK12] Yonina C Eldar and Gitta Kutyniok. *Compressed sensing: theory and applications*. Cambridge university press, 2012. [8](#)
- [GGR98] Oded Goldreich, Shari Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *Journal of the ACM (JACM)*, 45(4):653–750, 1998. [1](#)
- [GI10] Anna Gilbert and Piotr Indyk. Sparse recovery using sparse matrices. *Proceedings of the IEEE*, 98(6):937–947, 2010. [8](#)
- [Gle94] Paul Glendinning. *Stability, instability and chaos: an introduction to the theory of nonlinear differential equations*, volume 11. Cambridge university press, 1994. [1](#)
- [GLSS18] Ankit Garg, Yin Tat Lee, Zhao Song, and Nikhil Srivastava. A matrix expander chernoff bound. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1102–1114, 2018. [8](#)
- [Gol10] Oded Goldreich. Introduction to testing graph properties. In *Property testing*, pages 105–141. Springer, 2010. [1](#), [54](#)
- [Gol17] Oded Goldreich. *Introduction to property testing*. Cambridge University Press, 2017. [1](#)

- [GT11] Alex Gittens and Joel A Tropp. Tail bounds for all eigenvalues of a sum of random matrices. *arXiv preprint arXiv:1104.4513*, 2011. [5](#), [8](#), [11](#), [33](#)
- [GW95] Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995. [1](#)
- [HMAS17] Insu Han, Dmitry Malioutov, Haim Avron, and Jinwoo Shin. Approximating spectral sums of large-scale matrices using stochastic chebyshev approximations. *SIAM Journal on Scientific Computing*, 39(4):A1558–A1585, 2017. [8](#)
- [Ind06] Piotr Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *Journal of the ACM (JACM)*, 53(3):307–323, 2006. [8](#)
- [IVWW19] Piotr Indyk, Ali Vakilian, Tal Wagner, and David Woodruff. Sample-optimal low-rank approximation of distance matrices. *arXiv preprint arXiv:1906.00339*, 2019. [7](#), [8](#)
- [JST11] Hossein Jowhari, Mert Sağlam, and Gábor Tardos. Tight bounds for lp samplers, finding duplicates in streams, and related problems. In *Proceedings of the Thirtieth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS '11*, pages 49–58, New York, NY, USA, 2011. ACM. [8](#)
- [JSTW19] Rajesh Jayaram, Gokarna Sharma, Srikanta Tirthapura, and David P Woodruff. Weighted reservoir sampling from distributed streams. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 218–235, 2019. [8](#)
- [JW18] Rajesh Jayaram and David P Woodruff. Perfect lp sampling in a data stream. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 544–555. IEEE, 2018. [8](#)
- [JW19] Rajesh Jayaram and David P Woodruff. Towards optimal moment estimation in streaming and distributed models. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019. [8](#)
- [KBV09] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009. [1](#)
- [KIDP16] Ramakrishnan Kannan, Mariya Ishteva, Barry Drake, and Haesun Park. Bounded matrix low rank approximation. In *Non-negative Matrix Factorization Techniques*, pages 89–118. Springer, 2016. [1](#)
- [KLS20] Rasmus Kyng, Kyle Luh, and Zhao Song. Four deviations suffice for rank 1 matrices. *Advances in Mathematics*, 375:107366, 2020. [8](#)
- [Kna09] Michael P Knapp. Sines and cosines of angles in arithmetic progression. *Mathematics Magazine*, 82(5):371–372, 2009. [49](#)

- [KNW10] Daniel M Kane, Jelani Nelson, and David P Woodruff. On the exact space complexity of sketching and streaming small norms. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 1161–1178. SIAM, 2010. [8](#)
- [KOSZ13] Jonathan A Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. A simple, combinatorial algorithm for solving sdd systems in nearly-linear time. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 911–920, 2013. [1](#)
- [KS03] Robert Krauthgamer and Ori Sasson. Property testing of data dimensionality. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 18–27. Society for Industrial and Applied Mathematics, 2003. [7](#)
- [KS18] Rasmus Kyng and Zhao Song. A matrix chernoff bound for strongly rayleigh distributions and spectral sparsifiers from a few random spanning trees. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 373–384. IEEE, 2018. [8](#)
- [LNW14] Yi Li, Huy L Nguyen, and David P Woodruff. On sketching matrix norms and the top singular vector. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1562–1581. SIAM, 2014. [8](#)
- [LNW19] Yi Li, Huy L Nguyen, and David P Woodruff. On approximating matrix norms in data streams. *SIAM Journal on Computing*, 48(6):1643–1697, 2019. [8](#)
- [LPP91] Françoise Lust-Piquard and Gilles Pisier. Non commutative khintchine and paley inequalities. *Arkiv för matematik*, 29(1-2):241–260, 1991. [8](#)
- [LW16a] Yi Li and David P Woodruff. On approximating functions of the singular values in a stream. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 726–739, 2016. [8](#)
- [LW16b] Yi Li and David P Woodruff. Tight bounds for sketching the operator norm, schatten norms, and subspace embeddings. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016. [5](#), [8](#)
- [LW17] Yi Li and David P Woodruff. Embeddings of schatten norms with applications to data streams. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017. [8](#)
- [LWW14] Yi Li, Zhengyu Wang, and David P Woodruff. Improved testing of low rank matrices. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 691–700, 2014. [1](#), [7](#), [8](#)
- [MSS15] Adam W Marcus, Daniel A Spielman, and Nikhil Srivastava. Interlacing families ii: Mixed characteristic polynomials and the kadison–Tsinger problem. *Annals of Mathematics*, pages 327–350, 2015. [8](#)

- [MW10] Morteza Monemizadeh and David P Woodruff. 1-pass relative-error lp-sampling with applications. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 1143–1160. SIAM, 2010. [8](#)
- [MW17] Cameron Musco and David P Woodruff. Sublinear time low-rank approximation of positive semidefinite matrices. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 672–683. IEEE, 2017. [1](#), [7](#), [8](#)
- [Pis09] Gilles Pisier. Remarks on the non-commutative khintchine inequalities for $0 < p < 2$. *Journal of Functional Analysis*, 256(12):4128–4161, 2009. [8](#)
- [PR03] Michal Parnas and Dana Ron. Testing metric properties. *Information and Computation*, 187(2):155–195, 2003. [7](#)
- [PR17] Gilles Pisier and Éric Ricard. The non-commutative khintchine inequalities for $0 < p < 1$. *Journal of the Institute of Mathematics of Jussieu*, 16(5):1103–1123, 2017. [8](#)
- [RV07] Mark Rudelson and Roman Vershynin. Sampling from large matrices: An approach through geometric functional analysis. *Journal of the ACM (JACM)*, 54(4):21–es, 2007. [8](#), [12](#), [33](#), [54](#)
- [Sch35] Isaac J Schoenberg. Remarks to maurice frechet’s article “sur la definition axiomatique d’une classe d’espace distances vectoriellement applicable sur l’espace de hilbert. *Annals of Mathematics*, pages 724–732, 1935. [1](#)
- [SKZ14] Alaa Saade, Florent Krzakala, and Lenka Zdeborová. Spectral clustering of graphs with the bethe hessian. In *Advances in Neural Information Processing Systems*, pages 406–414, 2014. [1](#)
- [SL⁺91] Jean-Jacques E Slotine, Weiping Li, et al. *Applied nonlinear control*, volume 199. Prentice hall Englewood Cliffs, NJ, 1991. [1](#)
- [ST04] Daniel A Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 81–90, 2004. [1](#)
- [ST11] Daniel A Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011. [8](#)
- [Ste10] David Steurer. Fast sdp algorithms for constraint satisfaction problems. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 684–697. SIAM, 2010. [1](#)
- [SWYZ19] Xiaoming Sun, David P Woodruff, Guang Yang, and Jialin Zhang. Querying a matrix through matrix-vector products. *arXiv preprint arXiv:1906.05736*, 2019. [8](#)
- [SZ20] Zhao Song and Ruizhe Zhang. Hyperbolic polynomials i: Concentration and discrepancy. *arXiv preprint arXiv:2008.09593*, 2020. [8](#)

- [Tao11] Terence Tao. Topics in random matrix theory. *Lecture Notes*, 2011. [33](#)
- [Tho72] Robert C Thompson. Principal submatrices ix: Interlacing inequalities for singular values of submatrices. *Linear Algebra and its Applications*, 5(1):1–12, 1972. [38](#)
- [Tro08] Joel A Tropp. Norms of random submatrices and sparse approximation. *Comptes Rendus Mathématique*, 346(23-24):1271–1274, 2008. [8](#), [54](#)
- [Tro15] Joel Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015. [5](#), [8](#), [11](#)
- [Tur48] Alan M Turing. Rounding-off errors in matrix processes. *The Quarterly Journal of Mechanics and Applied Mathematics*, 1(1):287–308, 1948. [1](#)
- [VB96] Lieven Vandenberghe and Stephen Boyd. Semidefinite programming. *SIAM review*, 38(1):49–95, 1996. [1](#), [6](#)
- [Ver10] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010. [8](#), [51](#)
- [Wai19] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019. [1](#)
- [WSV12] Henry Wolkowicz, Romesh Saigal, and Lieven Vandenberghe. *Handbook of semidefinite programming: theory, algorithms, and applications*, volume 27. Springer Science & Business Media, 2012. [1](#)

A Proof of Eigenvalue Identity

Proposition A.1. Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be any real symmetric matrix. Then $\min_{\mathbf{B} \geq 0} \|\mathbf{A} - \mathbf{B}\|_F^2 = \sum_{i: \lambda_i(\mathbf{A}) < 0} \lambda_i^2(\mathbf{A})$.

Proof. Let g_i be the eigenvector associated with $\lambda_i = \lambda_i(\mathbf{A})$. First, setting $\mathbf{B} = \sum_{i: \lambda_i(\mathbf{A}) \geq 0} \lambda_i g_i g_i^\top$, which is a PSD matrix, we have $\|\mathbf{A} - \mathbf{B}\|_F^2 = \|\sum_{i: \lambda_i(\mathbf{A}) < 0} \lambda_i (g_i g_i^\top \mathbf{A})\|_2^2 = \sum_{i: \lambda_i(\mathbf{A}) < 0} \lambda_i^2(\mathbf{A})$, where the second equality follows from the Pythagorean Theorem, which proves that $\min_{\mathbf{B} \geq 0} \|\mathbf{A} - \mathbf{B}\|_F^2 \leq \sum_{i: \lambda_i(\mathbf{A}) < 0} \lambda_i^2(\mathbf{A})$. To see the other direction, fix any PSD matrix \mathbf{B} , and let $\mathbf{Z} = \mathbf{B} - \mathbf{A}$. Then $\mathbf{Z} + \mathbf{A} \geq 0$, where \geq is the Lowner ordering, thus $\mathbf{Z} \geq -\mathbf{A}$, which by definition implies that $x^\top \mathbf{Z} x \geq -x^\top \mathbf{A} x$ for all $x \in \mathbb{R}^n$. Then by the Courant-Fischer variational characterization of eigenvalues, we have that $\lambda_i(\mathbf{Z}) \geq -\lambda_i(\mathbf{A})$ for all i . In particular, $|\lambda_i(\mathbf{Z})| \geq |\lambda_i(\mathbf{A})|$ for all i such that $\lambda_i(\mathbf{A}) < 0$. Thus $\|\mathbf{Z}\|_F^2 = \sum_i \lambda_i^2(\mathbf{Z}) \geq \sum_{i: \lambda_i(\mathbf{A}) < 0} \lambda_i^2(\mathbf{Z}) \geq \sum_{i: \lambda_i(\mathbf{A}) < 0} \lambda_i^2(\mathbf{A})$, which completes the proof. ■