Proceedings of Machine Learning Research vol 178:1–42, 2022 35th Annual Conference on Learning Theory

Private Matrix Approximation and Geometry of Unitary Orbits

Oren Mangoubi
Worcester Polytechnic Institute

Yikai Wu
Yale University

Satyen Kale
Google Research

Abhradeep Guha Thakurta
Google Research

Nisheeth K. Vishnoi
Yale University

Editors: Po-Ling Loh and Maxim Raginsky

Abstract

Consider the following optimization problem: Given $n \times n$ matrices A and Λ, maximize $\langle A, U\Lambda U^* \rangle$ where U varies over the unitary group $\mathrm{U}(n)$. This problem seeks to approximate A by a matrix whose spectrum is the same as Λ and, by setting Λ to be appropriate diagonal matrices, one can recover matrix approximation problems such as PCA and rank-k approximation. We study the problem of designing differentially private algorithms for this optimization problem in settings where the matrix A is constructed using users' private data. We give efficient and private algorithms that come with upper and lower bounds on the approximation error. Our results unify and improve upon several prior works on private matrix approximation problems. They rely on extensions of packing/covering number bounds for Grassmannians to unitary orbits which should be of independent interest.

Keywords: Unitary orbits, differential privacy, packing number, PCA, rank-k approximation

1. Introduction

In machine learning and statistical data analysis, a widely used technique is to represent data as a matrix and perform computations on the covariance matrix to extract statistical information from data. For instance, consider the setting with n users and where one represents the features of each user by a vector $x_i \in \mathbb{R}^d$, giving rise to the $d \times d$ covariance matrix $M = \sum_{i=1}^n x_i x_i^\top$. In many applications, approximations to such matrices are sought to reduce the space/time required to perform computations or, to replace them by matrices with a specified spectrum (Sarwar et al., 2000; Paterek, 2007; Koren et al., 2009; Beutel et al., 2015). An example of the first kind is the rank-k approximation problem where one is given a positive integer k and the goal is to find a rank-k matrix H which is "close" to M. An example of the latter class of problems is the rank-k PCA problem where one is given a positive integer k and the goal is to output the matrix H corresponding to the projection onto the subspace spanned by the top k eigenvectors of M. Closeness is usually measured using the spectral or Frobenius norm of $M - H$. All of these problems are extensively studied and algorithms for these problems have been well studied and deployed; see Blum et al. (2020).

Since such matrix approximation problems are often applied to matrices arising from user data (i.e. each user contributes one vector x_i to the sum above), an important concern is to protect the privacy of the users. Even without fixing a specific notion of privacy, traditional algorithms for these problems can leak information of users. For instance, suppose we know that a user vector x is part of exactly one of the two covariance matrices M and M', but we cannot access the data matrices directly and can only obtain the information of the data matrices using PCA. If we apply a traditional algorithm for rank-k PCA onto M and M', we obtain two projection matrices H and H' spanned by the top k eigenvectors of M and M', respectively. Then, if x is in the subspace of H but not in the subspace of H', we know for sure that x is part of M but not of M' – leading to a breach of privacy of the data vector x. Important examples of real-world privacy breaches in settings of this nature include the Netflix prize problem and (Bennett and Lanning, 2007) and recommendation systems of Amazon and Hunch (Calandrino et al., 2011). It is thus important to design private algorithms for fundamental matrix-approximation problems.

The notion of differential privacy has arisen as an important formalization of what it means to protect privacy of individuals in a dataset (Dwork, 2006). We say that two Hermitian PSD matrices M and M' are *neighbors* if each matrix is obtained from the other by replacing one user's vector by another user's vector. In other words, M and M' are neighbors if and only if there exists $x, y \in \mathbb{C}^d$ such that $\|x\|_2, \|y\|_2 \leqslant 1$ and $M' = M - xx^* + yy^*$. We can now define differentially private computations on matrices.

Definition 1.1 (Differential Privacy) *For a given $\varepsilon \geqslant 0$ and $\delta \geqslant 0$, a randomized mechanism $\mathcal{M}$ is said to be (ε, δ)-differentially private if for any two neighboring matrices M and M' and any measurable set of possible output S, it holds that*

$$\mathbb{P}[\mathcal{M}(M) \in S] \leqslant \exp(\varepsilon) \cdot \mathbb{P}[\mathcal{M}(M') \in S] + \delta.$$

When $\delta = 0$, the mechanism is said to be ε-differentially private.

There have been multiple works that give differentially private algorithms for matrix approximation problems, including rank-k approximation (Kapralov and Talwar, 2013; Upadhyay, 2018; Amin et al., 2019) and rank-k PCA (Chaudhuri et al., 2013; Dwork et al., 2014; Leake et al., 2021). Roughly, these algorithms can be divided into two categories: Those satisfying pure differential privacy (ε-differential privacy) (Chaudhuri et al., 2013; Kapralov and Talwar, 2013; Amin et al., 2019) and those satisfying (ε, δ)-differential privacy with a $\delta > 0$ (Chaudhuri et al., 2013; Dwork et al., 2014; Upadhyay, 2018). Pure differential privacy provides better privacy protection and we focus on pure differential privacy in this paper.

All the algorithms mentioned above that come with pure differential privacy guarantees utilize the exponential mechanism (McSherry and Talwar, 2007) (see Theorem B.3). This mechanism involves sampling from an exponential distribution which, in turn, depends on the utility function chosen. Among these algorithms, one of the algorithms by Chaudhuri et al. (2013) (PPCA) provides a near-optimal algorithm for PCA under pure differential privacy. However, the error upper and lower bounds are only proved for the first principal component (the top eigenvector). In addition, their algorithm satisfying pure differential privacy (PPCA) is implemented with a Gibbs sampler which is not shown to run in polynomial time.

Kapralov and Talwar (2013) provide two different algorithms under pure differential privacy for rank-1 approximation and rank-k approximation. They also provides error upper and lower bounds for both problems. Their rank-1 approximation relies on an efficient way to sample from a unit vector using the exponential mechanism. The algorithm outputs the sampled vector as the estimation of the first eigenvector of the input matrix. The rank-k approximation samples top k eigenvectors iteratively. The error bound is worse compared to the rank-1 case and there is a significant gap from the lower bound proved in the paper.

Amin et al. (2019) provide a differentially private algorithm for the version of rank-k approximation problem when $k = d$ (*covariance matrix estimation problem*). This problem is trivial without a privacy requirement: one can set the output H as the input covariance matrix M. In the differentially private case, Amin et al. (2019) give an algorithm that samples eigenvectors iteratively using an exponential mechanism. It uses a different error measure compared to Kapralov and Talwar (2013) and, hence, the error bounds cannot be compared directly. However, the algorithm of Amin et al. (2019) only applies to the covariance matrix estimation problem which is a special case of the rank-k approximation problem.

2. Our Work

2.1. Unitary Orbit Optimization

We first present a generalized problem that captures the matrix approximation problems mentioned above. The problem is a linear optimization problem over an orbit of the unitary group. Recall that a matrix $U \in \mathbb{C}^{d\times d}$ is said to be unitary if $UU^* = I$. The set of unitary matrices forms a group under matrix multiplication and is denoted by $\mathrm{U}(d)$. $\mathrm{U}(d)$ is also a non-convex manifold. For a given $d \times d$ Hermitian matrix H, $\mathrm{U}(d)$ acts on it by conjugation as follows: $H \mapsto UHU^*$ for a unitary matrix U. Note that H has the same eigenvalues as UHU^* for any unitary matrix U. Thus, the set of matrices obtainable from H under this action have the same set of eigenvalues. Given a diagonal matrix $\Lambda := \mathrm{diag}\,(\lambda_1, \ldots, \lambda_d)$, we denote its unitary orbit: $\mathcal{O}_\Lambda := \{U\Lambda U^* : U \in \mathrm{U}(d)\}$.

Problem 2.1 (Unitary orbit optimization) *Given a Hermitian matrix $M \in \mathbb{C}^{d\times d}$ with eigenvalues $\gamma_1 \geqslant \cdots \geqslant \gamma_d$ and a list of eigenvalues $\lambda_1 \geqslant \cdots \geqslant \lambda_d$, the goal is to find a Hermitian matrix $H \in \mathcal{O}_\Lambda$ with $\Lambda := \mathrm{diag}\,(\lambda_1, \ldots, \lambda_d) \in \mathbb{C}^{d\times d}$ that maximizes $\langle M, H\rangle := \mathrm{Tr}(M^*H)$.*

Since $U\Lambda U^*$ has the same eigenvalues as Λ, this problem asks to find the "closest" matrix to a given matrix M, with eigenvalues identical to those of Λ. This is a well-studied problem and the Schur-Horn Theorem implies that the optimal solution to this problem is the matrix $H = U\Lambda U^*$ where U is a unitary matrix whose columns are the eigenvectors of M, attaining the optimal value $\sum_{i=1}^d \lambda_i\gamma_i$ (Schur, 1923; Horn, 1954).

Rank-k PCA, rank-k approximation, and covariance matrix estimation of a given covariance matrix M can be reduced to Problem 2.1 by a careful choice of λ_i's. The rank-k PCA problem is obtained by setting $\lambda_1 = \cdots = \lambda_k = 1$ and $\lambda_{k+1} = \cdots = \lambda_d = 0$. The rank-$k$ approximation problem is obtained by setting λ_i equal to the i-th largest eigenvalue of M for $1 \leqslant i \leqslant k$, and to 0 for $i > k$. Finally, the covariance matrix estimation problem is obtained by setting λ_i equal to the i-th largest eigenvalue of M for all $i = 1, 2, \ldots, d$. In the rank-k approximation and covariance estimation problems, we consider both the setting where the eigenvalues $\lambda_1, \ldots, \lambda_k$ are given as prior "non-private" information, as well as the more challenging setting when $\lambda_1, \ldots, \lambda_k$ are private.

2.2. Upper Bound Results

Our first result is an ε-differentially private mechanism for Problem 2.1 when the matrix Λ is non-private (as in the case of rank-k PCA). Our algorithm (Algorithm 1) utilizes the exponential mechanism (McSherry and Talwar, 2007) and samples H from $\mathcal{O}_\Lambda$ from a density that is close in infinity distance to $\exp(\frac{\varepsilon}{\lambda_1}\langle M, H\rangle)$. However, to do this, we need a unitarily invariant measure μ_Λ on $\mathcal{O}_\Lambda$. Such a measure can be derived from the Haar measure on $\mathrm{U}(d)$; see Leake and Vishnoi (2021). Note that while Problem 2.1 makes sense for general Hermitian M, in our results we consider the case when M is Hermitian and positive semidefinite (PSD) as in the case of a covariance matrix.

Theorem 2.2 (Differentially private unitary orbit optimization) *For any $\varepsilon \in (0,1)$, there is a randomized ε-differentially private algorithm (Algorithm 1) such that given a $d \times d$ PSD Hermitian matrix $M \in \mathbb{C}^{d\times d}$ with eigenvalues $\gamma_1 \geqslant \gamma_2 \geqslant \cdots \geqslant \gamma_d \geqslant 0$, the maximum rank of the output matrix $k \in [d]$, and a list of top k nonnegative eigenvalues of the output matrix $\lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_k \geqslant 0$, outputs a $d\times d$ PSD Herimitian matrix $H \in \mathbb{C}^{d\times d}$ with eigenvalues $\lambda_1, \lambda_2, \ldots, \lambda_k, 0, \ldots, 0$, where there are $d - k$ 0's. Moreover, for any $\beta \in (0,1)$, with probability at least $1-\beta$, we have $\langle M, H\rangle \geqslant \sum_{i=1}^k \gamma_i\lambda_i - \tilde{O}\left(\frac{dk\lambda_1}{\varepsilon}\right)$, where $\tilde{O}$ hides logarithmic factors of $\frac{1}{\beta}$ and $\mathrm{Tr}(M)$. The number of arithmetic operations required by this algorithm is polynomial in $\log\frac{1}{\varepsilon}$, λ_1, $\gamma_1-\gamma_d$, and the number of bits representing $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_k)$ and $\gamma = (\gamma_1, \gamma_2, \ldots, \gamma_d)$.*

This theorem is a generalization of the result in Leake et al. (2021) which proved such a theorem for the special case of rank-k PCA (when the orbit eigenvalues are $\lambda_1 = \cdots = \lambda_k = 1$ and $\lambda_{k+1} = \cdots = \lambda_d = 0$). Our algorithm leverages efficient algorithms to sample approximately from such exponential densities on unitary orbits by Leake et al. (2021); Mangoubi and Vishnoi (2021) that provide guarantees on the closeness of the target distribution and the actual distribution in infinity distance.[1] Note that the error bound in Leake et al. (2021) improves on Kapralov and Talwar (2013) but is weaker than in Theorem 2.2 since Theorem 2.2 holds with high probability while Leake et al. (2021) only holds in expectation. The proof of Theorem 2.2 appears in Appendix C and uses a covering number bound for the orbit $\mathcal{O}_\Lambda$ (Lemma C.3) that generalizes the upper covering bounds for the Grassmannian (Szarek, 1982).

Our next result considers the setting of Problem 2.1 when Λ is the spectrum of M: $\lambda_i = \gamma_i$ for $1 \leqslant i \leqslant k$ and $\lambda_i = 0$ for $i > k$ (as in the rank-k covariance matrix approximation problem). In this case, Λ is also private and Algorithm 1 does not apply as such. However, we show that adding Laplace noise to λ_is, sorting them, and then using Algorithm 1 suffices; see Algorithm 2.

Theorem 2.3 (Differentially private rank-k approximation) *Given a PSD Hermitian input matrix $M \in \mathcal{H}_+^d$, a $k \in [d]$, and an $\varepsilon > 0$. Let the eigenvalues of M be $\lambda_1 \geqslant \cdots \geqslant \lambda_d \geqslant 0$. There exists a randomized ε-differentially private algorithm (Algorithm 2), which outputs a rank-k matrix $H \in \mathcal{H}_+^d$ and a list of estimated eigenvalues $\tilde{\lambda}_1, \ldots, \tilde{\lambda}_k$. For any $\beta \in (0,1)$, with probability at least $1-\beta$, for all $i \in [k]$, we have $|\tilde{\lambda}_i - \lambda_i| \leqslant O\left(\frac{1}{\varepsilon}\log\frac{1}{\beta}\right)$, and*

$$\|M - H\|_F^2 \leqslant \sum_{\ell=k+1}^{d} \lambda_\ell^2 + \tilde{O}\left(\frac{k}{\varepsilon^2} + \frac{dk}{\varepsilon}\left(\lambda_1 + \frac{1}{\varepsilon}\right)\right),$$

where $\tilde{O}$ hides logarithmic factors of $\frac{1}{\beta}$ and $\sum_{\ell=1}^k \lambda_\ell$. The number of arithmetic operations required by this algorithm is polynomial in $\log\frac{1}{\varepsilon}$, λ_1, and the number of bits representing $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_d)$.

1. For two densities ν and π, the infinity distance is $\mathrm{d}_\infty(\nu, \pi) := \sup_\theta |\log\frac{\nu(\theta)}{\pi(\theta)}|$.

The proof of Theorem 2.3 is an extension of the proof of Theorem 2.2 and appears in Appendix D. Since the (full-rank) covariance matrix estimation problem is a special case of the rank-k approximation problem (when we set $k = d$), the above result immediately applies in this case. Theorem 2.3 improves upon the bound in (Amin et al., 2019): Roughly, when the covariance matrix has its largest eigenvalue within a constant factor of its middle eigenvalue $\lambda_1 = O(\lambda_{\frac{k}{2}})$, our bound of $\tilde{O}(\frac{dk}{\varepsilon}(\lambda_1 + \frac{1}{\varepsilon}))$ is $O(d)$ better than the bound $\tilde{\Omega}(\frac{kd^2}{\varepsilon})$ of (Amin et al., 2019). This includes the setting when the input matrix M is a random sample covariance matrix from the Wishart distribution Wishart (1928) (that is $M = \frac{1}{m}X^\top X$, where X is a $d \times m$ matrix with i.i.d. standard Gaussian entries), as such a matrix has, with high probability, $\lambda_1 = O(\lambda_{\frac{k}{2}})$ for any m, d, where $k = \min(m, d)$. We discuss these examples in detail in Appendix A.

Note that in Theorem 2.3 we do not lose utility due to privatization of the eigenvalues whenever $\lambda_1 \geqslant \Omega(\frac{1}{\varepsilon})$, which is often the case in practice. In this case the utility bound $\tilde{O}(\frac{dk\lambda_1}{\varepsilon})$ in Theorem 2.2 (where we assume the eigenvalues are "public", and do not have an eigenvalue privatization step) is the same as the utility bound $\tilde{O}(\frac{dk}{\varepsilon}(\lambda_1 + \frac{1}{\varepsilon}))$ in Theorem 2.3 (where we *do* privatize eigenvalues).

2.3. Lower Bound Results

We give a lower bound for an ε-differentially private algorithm in the case where the eigenvalues $\gamma_1, \ldots, \gamma_d$ of the input matrix are equal to the eigenvalues $\lambda_1, \ldots, \lambda_d$ of the output matrix. Note that this lower bound holds even when the eigenvalues of the input matrix are given to the algorithm as prior non-private information.

Theorem 2.4 (Error lower bound) *Suppose that $\lambda_1 \geqslant \cdots \geqslant \lambda_d \geqslant 0$ and $\varepsilon > 0$. Then for any ε-differentially private algorithm $\mathcal{A}$ which takes as input a Hermitian matrix and outputs a rank-k Hermitian matrix with eigenvalues $\lambda_1, \ldots, \lambda_k$, there exists a $d \times d$ PSD Hermitian matrix M with eigenvalues $\gamma_i = \lambda_i$, $i \in [d]$, such that, with probability at least $\frac{1}{2}$, the output $H := \mathcal{A}(M)$ of the algorithm satisfies*

$$\|M - H\|_F^2 \geqslant \Omega\left(\sum_{\ell=k+1}^{d} \lambda_\ell^2 + \frac{d}{\max(\lambda_1\sqrt{\varepsilon}, \sqrt{d})^2} \max_{1 \leqslant i \leqslant \frac{d}{2}} i \times (\lambda_i - \lambda_{d-i+1})^2\right). \tag{1}$$

Note that the r.h.s. of our lower bound is never larger than $\Omega(\sum_{\ell=1}^{d} \lambda_\ell^2)$; this is true for any error lower bound since the diameter $D_{\|\cdot\|_F}(\mathcal{O}_\Lambda)$ of the unitary orbit is $D_{\|\cdot\|_F}(\mathcal{O}_\Lambda) := \sup_{M,H \in \mathcal{O}_\Lambda} \|M - H\|_F = O(\sqrt{\sum_{\ell=1}^{d} \lambda_\ell^2})$. The proof of Theorem 2.4 is given in Appendix F. The proof of Theorem 2.4 relies on a novel packing number lower bound for the unitary orbit $\mathcal{O}_\Lambda$ (Theorem 2.7). As a first attempt we show a packing number bound for the entire unitary orbit (Inequality 3). Unfortunately the resulting utility error lower bound (Inequality 11 in the proof overview) is (roughly) proportional to $e^{-\frac{1}{d} D_{\|\cdot\|_F}(\mathcal{O}_\Lambda)^2}$, which is exponentially small in the eigenvalues $\lambda_1, \ldots, \lambda_d$. To achieve an error bound polynomial in the λ's, we instead show a packing bound on a ball of radius ω inside the orbit, where ω is carefully chosen to ensure that the error bound is polynomial in $\lambda_1, \ldots, \lambda_d$.

Next, we give a corollary of Theorem 2.4, which provides a lower bound for the rank-k approximation problem (which includes the covariance matrix estimation problem as a special case).

Corollary 2.5 (Lower bound for covariance estimation) *Suppose that $\lambda_1 \geqslant \cdots \geqslant \lambda_d \geqslant 0$ and $\varepsilon > 0$. Then for any ε-differentially private algorithm $\mathcal{A}$ which takes as input a Hermitian matrix*

M and outputs a rank-k Hermitian matrix $H = \mathcal{A}(M)$, there exists a $d \times d$ PSD Hermitian matrix M with eigenvalues λ_i, $i \in [d]$, such that, with probability at least $\frac{1}{2}$, the output $H := \mathcal{A}(M)$ of the algorithm satisfies

$$\|M - H\|_F^2 \geqslant \Omega\left(\sum_{\ell=k+1}^{d} \lambda_\ell^2 + \frac{d}{\max(\lambda_1\sqrt{\varepsilon}, \sqrt{d})^2} \max_{1 \leqslant i \leqslant \frac{d}{2}} i \times (\lambda_i - \lambda_{d-i+1})^2\right).$$

The proof of Corollary 2.5 is given in Appendix F. Note that, unlike in Theorem 2.4, the output matrix in Corollary 2.5 is allowed to be any matrix, and need not have the same eigenvalues as the input matrix. To verify that the lower bounds in Theorem 2.4 and Corollary 2.5 are indeed lower than the upper bounds in Theorems 2.2 and 2.3, we observe that when the input matrix is of any rank $1 \leqslant k \leqslant d$, $\max_{1 \leqslant i \leqslant d} i \times (\lambda_i - \lambda_{d-i+1})^2 \leqslant k \times \lambda_1^2$. Thus, the r.h.s. of the lower bound in Theorem 2.4 and Corollary 2.5 is at most $\frac{dk}{\varepsilon}$, up to a constant factor. On the other hand the upper bounds in Theorems 2.2 and 2.3 are each at least as large as $\frac{dk}{\varepsilon}(\lambda_1 + \frac{1}{\varepsilon})$, which is greater than $\frac{dk}{\varepsilon}$.

When the input matrix is rank-k, and $\lambda_{\frac{k}{4}} - \lambda_{\frac{3k}{4}} = \Omega(\lambda_1)$, Corollary 2.5 implies that, with probability at least $\frac{1}{2}$, $\|M - \mathcal{A}(M)\|_F^2 \geqslant \frac{dk}{\varepsilon}$ if $\lambda_1 \geqslant \Omega(\frac{\sqrt{d}}{\varepsilon})$ and $\|M - \mathcal{A}(M)\|_F^2 \geqslant k\lambda_1^2$ if $\lambda_1 \leqslant \Omega(\frac{\sqrt{d}}{\varepsilon})$. Thus, our lower bound matches our upper bound from Theorem 2.3 up to a factor of $\frac{\lambda_1}{\varepsilon}$ if $\lambda_1 \geqslant \Omega(\frac{\sqrt{d}}{\varepsilon})$ and a factor of $\frac{d}{\lambda_1 \varepsilon}$ otherwise. This includes the setting when the input matrix M is a random sample covariance matrix from the Wishart distribution Wishart (1928), as such a matrix has, with high probability, $\lambda_1 = O(\lambda_{\frac{k}{2}})$ and $\lambda_{\frac{k}{4}} - \lambda_{\frac{3k}{4}} = \Omega(\lambda_1)$ for any m, d, where $k = \min(m, d)$.

The only previous lower bound we are aware of for the problem of (pure) differentially private rank-k covariance matrix estimation is from Kapralov and Talwar (2013). Roughly, their result says that if, for any $\omega > 0$ and $\lambda_1 > 0$, we have $\lambda_1 > \frac{1}{\varepsilon} k(d-k)\log(\frac{1}{\omega})$, then, for any ε-differentially private algorithm $\mathcal{A}$, there exists a matrix M with top eigenvalue λ_1, such that the error (measured in the spectral norm) has a lower bound of $\|M - \mathcal{A}(M)\|_2 \geqslant \lambda_{k+1} + \delta\lambda_1$ with positive probability, where λ_{k+1} is the $k+1$'st eigenvalue of the matrix M guaranteed by their result. Since only a condition on the top eigenvalue λ_1 is specified in their result, to show their result it is sufficient to produce an input matrix M satisfying their lower bound with $\lambda_{k+1} = 0$, that is, an input matrix of rank k, and this is what they show in their proof. Solving for the value of ω which maximizes their lower bound, one gets that their lower bound implies $\|M - \mathcal{A}(M)\|_2 - \lambda_{k+1} \geqslant \Omega(e^{-\frac{\lambda_1 \varepsilon}{k(d-k)}}\lambda_1)$.

While our lower bound is stated in terms of Frobenius norm, to see what our results give for the spectral norm error, we can use the fact that the Frobenius norm distance between two rank-k projection matrices is at most $O(\sqrt{k})$ times the spectral norm distance to obtain a spectral norm bound. In the case where the input matrix is rank-k, our result implies a error bound of $\|M - H\|_2 \geqslant \Omega(\frac{d}{\varepsilon})$ if, e.g., $\lambda_1 \varepsilon > \Omega(\sqrt{d})$. Thus, our lower bound is larger by a factor of roughly $\frac{d}{\varepsilon \lambda_1} e^{\frac{\lambda_1 \varepsilon}{k(d-k)}}$.

For the general unitary orbit approximation problem, Theorem 2.4 implies the following utility lower bound on the Frobenius norm utility:

Corollary 2.6 (Lower bound for general γ and λ) *Suppose that $\gamma_1 \geqslant \cdots \geqslant \gamma_d \geqslant 0$, and $\lambda_1 \geqslant \cdots \geqslant \lambda_d \geqslant 0$, and $\varepsilon > 0$. Then for any ε-differentially private algorithm $\mathcal{A}$ which takes as input a Hermitian matrix and outputs a rank-k Hermitian matrix with eigenvalues $\lambda_1, \ldots, \lambda_k$, there exists a $d \times d$ PSD Hermitian matrix M with eigenvalues γ_i, $i \in [d]$, such that, with probability at least*

$\frac{1}{2}$*, the output* $H := \mathcal{A}(M)$ *of the algorithm satisfies*

$$\|M - H\|_F^2 \geqslant \Omega\left(\sum_{\ell=k+1}^{d} \gamma_\ell^2 + \frac{d}{\max(\gamma_1\sqrt{\varepsilon}, \sqrt{d})^2} \max_{1 \leqslant i \leqslant \frac{d}{2}} i \times (\gamma_i - \gamma_{d-i+1})^2\right).$$

The proof of Corollary 2.6 is given in Appendix F. Corollary 2.6 says that, given any $\gamma_1 \geqslant \cdots \geqslant \gamma_d \geqslant 0$ and any $\lambda_1 \geqslant \cdots \geqslant \lambda_d \geqslant 0$, the same lower utility bound given in Theorem 2.4 holds (with γ_i taking the place of λ_i on the r.h.s. of the inequality) even if the eigenvalues λ_i of the output matrix are not equal to γ_i.

2.4. Packing Number Bounds for Unitary Orbits

As our main technical tool for proving the lower bounds on the error in Theorem 2.4 and Corollary 2.5, we will show packing number bounds for the unitary orbit. For any set S in a normed vector space with norm $|||\cdot|||$ and any $\zeta > 0$, we define a ζ-packing of the set S with respect to $|||\cdot|||$ to be any collection of points $\{z_1, \ldots, z_J\} \subseteq S$, where $J \in \mathbb{N}$, such that $|||z_s - z_t||| \geqslant \zeta$ for any $s, t \in [J]$. We define the packing number $P(S, |||\cdot|||, \zeta)$ to be the supremum of the number of points in any ζ-packing of S. We also denote by $B(X, r) := \{Z \in \mathbb{C}^{d \times d} : |||Z - X||| \leqslant r\}$ a ball of radius r with center X with respect to the norm $|||\cdot|||$. We show the following lower bound on the packing number of any unitary orbit $\mathcal{O}_\Lambda$ with respect to the Frobenius norm $\|\cdot\|_F$, and also provide a bound on the packing number of any ball $B \cap \mathcal{O}_\Lambda$ which is a subset of the unitary orbit. Since the ζ-packing and ζ-covering numbers of any set are equal up to a factor of 2 in ζ (see equation 4), our packing number lower bound also implies a lower bound on the covering number of the unitary orbit.

Theorem 2.7 (Packing number lower bound for unitary orbits) *There exist universal constants* $C > c > 0$ *such that, for any* $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_d)$*, and any* $\omega, \zeta > 0$*, and any* $X \in \mathcal{O}_\Lambda$*,*

$$\log P(B(X, \omega) \cap \mathcal{O}_\Lambda, \|\cdot\|_F, \zeta) \geqslant \max_{1 \leqslant i < j \leqslant d} 2i \times (d - j + 1) \times \log\left(\frac{\min(\omega, \lambda_1\sqrt{i}, \lambda_1\sqrt{d-j+1}) \times (\lambda_i - \lambda_j)}{2C\lambda_1\zeta}\right). \quad (2)$$

Moreover, we get the following bound for the packing number of the entire unitary orbit $\mathcal{O}_\Lambda$*:*

$$\log P(\mathcal{O}_\Lambda, \|\cdot\|_F, \zeta) \geqslant \max_{1 \leqslant i < j \leqslant d} 2i \times (d - j + 1) \times \log\left(\frac{c\min(\sqrt{i}, \sqrt{d-j+1}) \times (\lambda_i - \lambda_j)}{\zeta}\right). \quad (3)$$

The proof of Theorem 2.7 is given in Appendix E. The bound in Theorem 2.7 depends on the gaps $\lambda_i - \lambda_j$ between the eigenvalues of Λ, and is largest when there is a large gap between eigenvalues $\lambda_i - \lambda_j$ such that both i and $d - j$ are large. A special case of the unitary orbits is the Grasmannian manifold $\mathcal{G}_{d,k}$ for any $k \leqslant d$, which is the set of k-dimensional subspaces in a d-dimensional vector space. Identifying each subspace $\mathcal{V} \in \mathcal{G}_{d,k}$ with its projection matrix, the Grassmanian $\mathcal{G}_{d,k}$ has a one-to-one correspondence with the unitary orbit $\lambda_1 = \cdots = \lambda_k = 1$ and $\lambda_{k+1} = \cdots = \lambda_d = 0$, and any norm on the unitary orbit induces a norm on the Grassmannian. Theorem 2.7 generalizes the covering/packing number *lower* bounds for the (complex) Grassmannian of Szarek (1982, 1998) (restated as Lemma E.1 in the Appendix; see also e.g. Pajor (1998) and Kapralov and Talwar (2013) for different proofs of the same result), to a lower bound on the covering/packing number of any unitary orbit $\mathcal{O}_\Lambda$. Namely, in the special case where $\lambda_1 = \cdots = \lambda_k = 1$ and $\lambda_{k+1} = \cdots = \lambda_d = 0$, the r.h.s. of Theorem 2.7 is just $2d \times (d-k) \log(\frac{cD_{\|\cdot\|_F}(\mathcal{G}_{d,\,k})}{\zeta})$, since the diameter of the Grassmannian is $D_{\|\cdot\|_F}(\mathcal{G}_{d,\,k}) = c' \min(\sqrt{k}, \sqrt{d-k+1})$, for universal constant c'.

3. Proof Techniques

3.1. Upper Bounds: Theorem 2.2 (and Theorem 2.3)

Given $M = \sum_{i=1}^n x_i x_i^*$ for a dataset $\{x_1, \ldots, x_n\} \subseteq \mathbb{C}^n$, where $\|x_i\| \leqslant 1$ for each i, and a diagonal matrix Λ, the goal of our algorithm is to output a matrix $H \in \mathcal{O}_\Lambda$ which maximizes the utility $\langle M, H\rangle$ under the constraint that the output is ε-differentially private. Moreover, we would like our algorithm to run in time polynomial in the number of bits needed to represent M and Λ.

Privacy guarantee. Given data sets $\{x_i\}_{i=1}^n$ and $\{x_i'\}_{i=1}^n$, we say that two matrices $M = \sum_{i=1}^n x_i x_i^*$ and $M' = \sum_{i=1}^n x_i' x_i'^*$ are neighbors if $x_i = x_i'$ for all but one pair of points i. And we say that the output of any algorithm $\mathcal{A}$ is ε-differentially private if for any M, M' which are neighbors, and any set S in the output space of the algorithm, we have $\mathbb{P}(\mathcal{A}(M) \in S) \leqslant e^\varepsilon\, \mathbb{P}(\mathcal{A}(M') \in S)$. Our algorithm ensures that its output is ε-differentially private by applying the exponential mechanism of McSherry and Talwar (2007) to sample a matrix $H = U\Lambda U^*$, where U is a unitary matrix, from the unitary orbit $\mathcal{O}_\Lambda$. For any choice of query function $q(M, H)$ and $\Delta > 0$, a sample from the exponential mechanism with probability distribution proportional to $\exp\left(\frac{\varepsilon q(D,r)}{2\Delta}\right)$, is guaranteed to be ε-differentially private as long as Δ is no greater than the sensitivity

$$\sup_{\substack{M,M' \\ M,M' \text{ are neighbors}}} |q(M,H) - q(M',H)|$$

of the query function for all H. To ensure that matrices H with a larger utility $\langle M, H\rangle$ are sampled with a higher probability, we apply the exponential mechanism with the query function $q(M, H) = \langle M, H\rangle$, and sample H from the distribution $\exp(\frac{\varepsilon}{\lambda_1}\langle M, H\rangle)\mathrm{d}\mu_\Lambda$, where $\mathrm{d}\mu_\Lambda$ is a unitarily invariant measure on $\mathcal{O}_\Lambda$ obtained from the Haar measure on the unitary group. Since we show that whenever M and M' differ by only one point x_i $|\langle M, H\rangle - \langle M', H\rangle| = |x_i H x_i^* - x_i' H x_i'^*| \leqslant \lambda_1$ (Lemma C.1), the sensitivity is $\Delta \leqslant \lambda_1$. Thus, Algorithm 1 is ε-differentially private.

Running time. To generate the sample from the distribution $\nu(H) \propto \exp(\frac{\varepsilon}{\lambda_1}\langle M, H\rangle)\mathrm{d}\mu_\Lambda$, we use the Markov chain sampling algorithm from (Leake et al., 2021) (improved in Mangoubi and Vishnoi (2021)), which generates a sample from the log-linear distributions on unitary orbits. The distribution π of the output of this algorithm is guaranteed have sampling error at most $O(\varepsilon)$ in the infinity-distance metric, $\sup_H |\log \frac{\nu(H)}{\pi(H)}| < \varepsilon$. Thus, the output of the Markov chain sampling algorithm is $O(\varepsilon)$-differentially private as well. Its running time bound is polynomial in λ_1, $\gamma_1 - \gamma_d$ and the number of bits needed to represent $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_k)$ and $\gamma = (\gamma_1, \gamma_2, \ldots, \gamma_d)$.

Upper bound on error. Our upper bound on error is based on a covering number argument. For any set S and any $\zeta > 0$, we define a ζ-covering of the set S with respect to a norm $|||\cdot|||$ on this set to be any collection of balls $\{B_1, \ldots, B_J\}$ of radius ζ with centers in S, where $J \in \mathbb{N}$ such that $S \subseteq \bigcup_{i=1}^J B_i$. We define the covering number $N(S, |||\cdot|||, \zeta)$ to be the smallest number J of Balls in any ζ-covering of S. The packing and covering numbers are equal up to a factor of 2 in the radius ζ (see e.g. chapter 3.5 of Mohri et al. (2018)):

$$P(S, \|\cdot\|, 2\zeta) \leqslant N(S, \|\cdot\|, \zeta) \leqslant P(S, \|\cdot\|, \zeta) \qquad \forall \zeta > 0. \tag{4}$$

From a standard result about the exponential mechanism (McSherry and Talwar (2007)), we have that the utility of the exponential mechanism satisfies

$$\mathbb{P}(M \notin S_t) \leqslant \frac{\exp(-\frac{\varepsilon}{2\Delta}t)}{\mu_\Lambda(S_{\frac{t}{2}})}, \tag{5}$$

where S_t is the set of all matrices M with utility $\langle M, H\rangle > \mathrm{OPT} - t$ and $\mathrm{OPT} = \sum_{i=1}^d \lambda_i \gamma_i$ is the optimal value that $\langle M, H\rangle$ can take. The key ingredient we need to bound the utility is an upper bound on the volume $\mu_\Lambda(S_{\frac{t}{2}})$ in the denominator of equation 5. We bound this quantity via a covering number argument. First, we show that $S_{\frac{t}{2}}$ is contained in a spectral norm ball B of radius $\frac{t}{2\Gamma}$, where $\Gamma := \mathrm{tr}(M)$, with center at the optimal point H_0, since, whenever $\|H - H_0\| \leqslant \frac{\tau}{2\sum_i \gamma_i}$,

$$\langle M, H\rangle = \langle M, H_0\rangle - \langle M, H_0\rangle \geqslant \sum_{i=1}^d \lambda_i\gamma_i - \|H_0 - H\|_2 \mathrm{tr}(M) \geqslant \sum_{i=1}^d \lambda_i\gamma_i - \frac{t}{2}.$$

To obtain a bound on the volume of $\mu_\Lambda(B)$, we use the fact that the spectral norm $\|\cdot\|_2$ and the measure $\mu_\Lambda(B)$ are both unitarily invariant. We say a norm $\|\cdot\|$ is *unitarily invariant* if $\|UXV\| = \|X\|$ for any $X \in \mathbb{C}^{d\times d}$ and any unitary matrices $U, V \in \mathrm{U}(d)$; in particular $\|\cdot\|_2$ and $\|\cdot\|_F$ are unitarily invariant norms. And we say a measure μ is unitarily invariant if $\mu(USV) = \mu(S)$ for each subset S and each $U, V \in \mathrm{U}(d)$. Since μ_Λ and $\|\cdot\|_2$ are both unitarily invariant, every $\|\cdot\|_2$-norm ball of radius $\frac{t}{2\Gamma}$ in $\mathcal{O}_\Lambda$ has the same volume with respect to the measure μ_Λ. Thus, if we can find a covering of $\mathcal{O}_\Lambda$ of some size N consisting only of balls of radius $\frac{t}{2\Gamma}$, we would have $\mu_\Lambda(B) \geqslant \frac{1}{N}$.

Thus, in terms of the covering number, we can rewrite the utility bound equation 5 as

$$\mathbb{P}\left(\sum_i \gamma_i\lambda_i - \langle M, H\rangle \leqslant t\right) \geqslant N\left(\mathcal{O}_\Lambda, \|\cdot\|_2, \frac{t}{2\Gamma}\right) \exp\left(-\frac{\varepsilon}{2\Delta}t\right). \tag{6}$$

To bound the utility with equation 6, we will show that the covering number of $\mathcal{O}_\Lambda$ satisfies $N(\mathcal{O}_\Lambda, \|\cdot\|_2, t) \leqslant (1 + \frac{4\lambda_1}{\zeta})^{2dk}$ (Lemma C.3). Plugging our covering number bound, and the sensitivity bound $\Delta \leqslant \lambda_1$ into equation 6 we get that

$$\mathbb{P}\left(\sum_i \gamma_i\lambda_i - \langle M, H\rangle \leqslant t\right) \geqslant \left(1 + 8\lambda_1\Gamma t^{-1}\right)^{2dk} \exp\left(-\frac{\varepsilon}{2\lambda_1}t\right), \qquad \forall t > 0.$$

Plugging $t = \Theta\left(\frac{\lambda_1}{\varepsilon} dk \log(\frac{\Gamma}{\beta})\right)$, we get that $\sum_i \gamma_i\lambda_i - \langle M, H\rangle \leqslant \tilde{O}(\frac{\lambda_1}{\varepsilon}dk)$ w.p. at least $1 - \beta$.

In the rank-k covariance matrix estimation problem, the algorithm is not handed the eigenvalues $\lambda_1, \ldots, \lambda_d$ as private information. The Algorithm 2 in Theorem 2.3 perturbs the eigenvalues by adding random Laplace noise. The proof of Theorem 2.3, in addition to the proof of Theorem 2.2, requires us to carefully bound the distance between the eigenvalues λ_i of the covariance matrix and the perturbed eigenvalues $\tilde{\lambda}_i$; see Section D.

Bounding the covering number of $\mathcal{O}_\Lambda$. To bound the covering number of $\mathcal{O}_\Lambda$, we will first show a covering bound for the set S_k of $d \times k$ matrices with orthonormal columns, and then construct a map from S_k to the unitary orbit $\mathcal{O}_\Lambda$. Towards this end, we observe that the matrices in $\mathcal{O}_\Lambda$ are of the form $H = U\Lambda U^*$ where U is a unitary matrix, and, since Λ has only k nonzero eigenvalues, H only depends on the first k eigenvectors of U, which we denote by U_1. To bound the ζ-covering number of the space S_k of $d \times k$ rectangular matrices U_1 with orthonormal columns, observe that each $U_1 \in S_k$ has spectral norm $\|U_1\|_2$ at most 1. Thus, the set S_k of $d \times k$ complex matrices is the unit sphere in a $2dk$-dimensional (real) normed space. To bound the covering number of S_k, we apply a well-known result (see e.g., Lemma 6.27 in Mohri et al. (2018)) which says that a minimal ζ'-covering $B_1, \ldots B_t$ of the unit ball in any $2dk$-dimensional normed space has cardinality at most $(1 + \frac{2}{\zeta'})^{2dk}$. To obtain a covering with balls with centers on the unit sphere, we take any point x in $B_i \cap S_k$ (if such a point exists), and note that the ball centered at x of radius 2ζ contains B_i.

To obtain a covering of $\mathcal{O}_\Lambda$, we consider the map ϕ which maps each $U_1 \in S_k$ to a matrix $\phi(U_1) = U_1^* \Lambda U_1$. Since Λ has rank-k, $\phi : S_k \to \mathcal{O}_\Lambda$ is surjective, and thus $\phi(\hat{B}_1), \ldots, \phi(\hat{B}_t)$ is a covering of the unitary orbit; however we still need to bound the radius of the balls $\phi(\hat{B}_1)$ to show that it is a ζ-covering. Towards this end, we note that for any $U_1, U_1' \in S_k$ we have that

$$\|\phi(U_1) - \phi(U_1')\|_2 = \|U_1^* \Lambda U_1 - U_1'^* \Lambda U_1'\|_2 \leqslant 2\|U_1^* \Lambda (U_1 - U_1')\|_2 \leqslant \frac{2}{\lambda_1}\|U_1 - U_1'\|_2.$$

Thus, if we set $\zeta' = \frac{\zeta}{2\lambda_1}$, we obtain a ζ-covering of $\mathcal{O}_\Lambda$, and this covering has cardinality $(1 + \frac{2}{\zeta'})^{2dk} = (1 + \frac{4\lambda_1}{\zeta})^{2dk}$, which gives an upper bound on the covering number $N(\mathcal{O}_\Lambda, \|\cdot\|_2, \zeta)$ of $\mathcal{O}_\Lambda$.

3.2. Lower Bounds: Theorem 2.4

To prove a lower bound on the error in the covariance matrix approximation problem, it is sufficient to consider the setting where the eigenvalues of the output matrix are given as (non-private) prior information to the algorithm. This is because, any algorithm which works without this prior information can also be applied to this setting by simply ignoring the information about the eigenvalues of the input matrix. Thus, any lower bound for the setting where the eigenvalues are given as a prior will also imply a lower bound for the covariance matrix estimation problem.

Towards this end, we first show a bound for a special case of the unitary orbit minimization problem (Theorem 2.4), where the output matrix is in the orbit $\mathcal{O}_\Lambda$ with eigenvalues $(\lambda_1, \ldots, \lambda_d) = \mathrm{diag}(\Lambda)$ that are equal to the (non-private) eigenvalues of the input matrix (for simplicity, in this proof overview we assume that $\lambda_1 \geqslant \Omega(\sqrt{d})$). We then show that, roughly speaking, since the matrix H which minimizes the Frobenius norm distance $\|M - H\|_F$ is the matrix $H = M$ and is therefore in the orbit $\mathcal{O}_\Lambda$ of the input matrix M, our lower bound for the unitary orbit minimization problem also implies the same lower bound for the covariance matrix estimation problem (Corollary 2.5; see the end of this section for an overview of the proof of this corollary).

Our lower bound relies on a "packing number" lower bound for the orbit $\mathcal{O}_\Lambda$. As a first attempt, we consider a maximal ζ-packing of the orbit $\mathcal{O}_\Lambda$, $\{U_i \Lambda U_i^*\}_{i=1}^{\mathfrak{p}}$, where $\mathfrak{p} = P(\mathcal{O}_\Lambda, \|\cdot\|_F, \zeta)$ is the packing number of $\mathcal{O}_\Lambda$. We show, using a contradiction argument, that $\|M - \mathcal{A}(M)\|_F^2 \geqslant \zeta^2$ (with probability at least $\frac{1}{2}$) for any input matrix M and any ζ small enough such that

$$4\varepsilon D^2 \leqslant \log P(\mathcal{O}_\Lambda, \|\cdot\|_F, \zeta), \tag{7}$$

where D is the diameter of $\mathcal{O}_\Lambda$. Suppose, on the contrary, that for every $i \in [\mathfrak{p}]$, we have that

$$\|M_i - \mathcal{A}(M_i)\|_F^2 < \zeta^2 \tag{8}$$

with probability at least $\frac{1}{2}$. Next, observe that one can always find $m \leqslant \|M_i - M_j\|_F^2 + d$ vectors $x_1, \ldots, x_m$ with norm $\|x_i\| \leqslant 1$ such that $M_i - M_j = \sum_{s=1}^m x_s x_s^* + d$. Thus, if M_i and M_j are data matrices with unit-norm data points, one can transform M_i into M_j by modifying at most $2m$ points in the dataset. From equation 8, we have that for each i, the output $\mathcal{A}(M_i)$ is in the ball $B(M_i, \zeta)$ with probability at least $\frac{1}{2}$. Thus, since $\mathcal{A}$ is ε-differentially private, for every i, j we have

$$e^{-2\varepsilon D^2} \leqslant e^{-\varepsilon(\|M_i - M_j\|_F^2 + d)} \leqslant \frac{\mathbb{P}(\mathcal{A}(M_i) \in B(M_j, \zeta))}{\mathbb{P}(\mathcal{A}(M_i) \in B(M_i, \zeta))} \leqslant 2\mathbb{P}(\mathcal{A}(M_i) \in B(M_j, \zeta)) \tag{9}$$

since $M_i, M_j \in \mathcal{O}_\Lambda$, and D is the diameter of $\mathcal{O}_\Lambda$. Thus, since equation 9 holds for every $j \in \mathfrak{p}$,

$$1 \geqslant \sum_{j=1}^{\mathfrak{p}} \mathbb{P}(\mathcal{A}(M_i) \in B(M_j, \zeta)) \geqslant 2\mathfrak{p} \times e^{-2\varepsilon D^2}. \tag{10}$$

Rearranging equation 10, we get that $\varepsilon D^2 \leqslant \log \mathfrak{p} = \log P(\mathcal{O}_\Lambda, \|\cdot\|_F, \zeta)$, which contradicts our assumption in equation 7. Thus, by contradiction, we have that $\|M - \mathcal{A}(M)\|_F^2 \geqslant \zeta^2$ with probability $\Omega(1)$ for any $\zeta > 0$ satisfying equation 7.

Plugging in our packing number bound for $\mathcal{O}_\Lambda$ (equation 3 in Theorem 2.7), and solving for the largest value of ζ satisfying equation 7, gives the lower bound of

$$\|M - \mathcal{A}(M)\|_F^2 \geqslant c\hat{D}^2 \times (\lambda_i - \lambda_j)^2 \exp\left(\frac{-2\varepsilon D^2}{i \times (d - j + 1)}\right) \tag{11}$$

for every $1 \leqslant i, j \leqslant d$, where $\hat{D} := D_{\|\cdot\|_F}(\mathcal{G}_{d-j+i+1,\, i})$ is the diameter of the Grassmanian. Unfortunately, since the diameter of $\mathcal{O}_\Lambda$ is $D \geqslant \lambda_1 - \lambda_d$ this lower bound is *exponential* in $\lambda_1 - \lambda_d$. The D term in the exponent comes from the fact that, since we have used a packing for the entire unitary orbit $\mathcal{O}_\Lambda$, the distance between any two balls in our packing is upper bounded by D. To achieve a bound that is polynomial in $\lambda_1 - \lambda_d$, we would instead like to use a packing for a smaller *subset* of the orbit $\mathcal{O}_\Lambda$, of some radius (roughly) $\omega \leqslant \sqrt{\frac{i \times (d-j+1)}{-2\varepsilon}}$. However, restricting our packing to a ball of radius ω– rather than the entire orbit– requires us to prove a packing number for a subset of the orbit, (equation 2). This leads to additional challenges in the proof of our packing bound which we describe in the next subsection.

Replacing the D term in equation 3 with $\omega = \Theta(\frac{\lambda_1}{\lambda_i - \lambda_j}\zeta)$ and plugging in our bound for the ζ-packing number of a ball of radius ω inside the orbit, and solving for the largest value of ζ satisfying equation 7, gives the improved lower bound of

$$\|M - \mathcal{A}(M)\|_F^2 \geqslant \Omega\left(\frac{i \times (d - j + 1)}{\lambda_1^2 \varepsilon} \times (\lambda_i - \lambda_j)^2\right) \tag{12}$$

for every $1 \leqslant i < j \leqslant d$. Unlike the bound in equation 11 which is exponential in the $\lambda' s$, this bound is *polynomial* in the λ's and in d, $\frac{1}{\varepsilon}$. If we plug in $j = d$ in equation 12 and take the maximum over all $i \in [d]$, and then plug in $i = 1$ and take the maximum over all $j \in [d]$, and finally take the larger of these two maximum values, we recover the error lower bound of Theorem 2.4.

3.3. Packing Number Lower Bounds: Theorem 2.7

In this section we first explain how we bound the packing number of the entire unitary orbit $\mathcal{O}_\Lambda$ (equation 3), and we then explain how we extend the proof to obtain a bound on the packing number of any ball inside $\mathcal{O}_\Lambda$ (equation 2).

The general strategy for proving our packing bounds for the unitary orbit (Theorem 2.7), is to first construct a map $\phi : \Omega \to \mathcal{O}_\Lambda$ from some space Ω with previously known packing number bounds to the unitary orbit. And, once we have a map ϕ and a packing $X_1, \ldots, X_t \in \Omega$, we show that the map preserves (a lower bound for) distances between points in the packing: $\|\phi(X_i) - \phi(X_j)\|_F \geqslant \beta \|X_i - X_j\|_F$ for some $\beta > 0$, implying $\phi(X_1), \ldots, \phi(X_t)$ is a $\zeta\beta$-packing of $\mathcal{O}_\Delta$.

As a first attempt, we consider the space of unitary matrices $\mathrm{U}(d)$ for our choice of Ω, and the map $\phi : U \to U\Lambda U^*$. Unfortunately, there may be $U, U' \in \mathrm{U}(d)$ such that $\|\phi(U) - \phi(U')\|_F = 0$ even though $\|U - U'\|_F > 0$ (For instance, if $\mathrm{diag}(\Lambda) = (1, 1, 0)$ and $U = I$ and U' is the matrix $[e_2, e_1, e_3]^\top$ where e_i is the vector with a 1 in the i'th entry and zero everywhere else, we have $\phi(U) - \phi(U') = 0$ and yet $\|U - U'\|_F = 2$.).

To get around this problem we instead consider a map ϕ from the (complex) Grassmannian manifold $\mathcal{G}_{d,i}$, the collection of subspaces of dimension i in d-dimensional space, to $\mathcal{O}_\Lambda$. Identifying each subspace with its associated rank-i projection matrix, we construct a maximal ζ-packing for $P_1, \ldots, P_\mathfrak{p} \in \mathcal{G}_{d,i}$, where $\mathfrak{p}$ is the packing number of $\mathcal{G}_{d,i}$. To bound the size of this packing, we use the covering/packing number bound from Szarek (1982) for the Grassmannian $\mathcal{G}_{d,i}$, which says that

$$\mathfrak{p} =: P(\mathcal{G}_{d,i}, \|\cdot\|_F, \zeta) \geqslant \left(\zeta^{-1} c D_{\|\cdot\|_F}(\mathcal{G}_{d,i}))\right)^{2di},$$

where $D_{\|\cdot\|_F}(\mathcal{G}_{d,i}))$ is the diameter of $\mathcal{G}_{d,i}$ and c is a universal constant.

To define our map $\phi(P)$ for any rank-i projection matrix $P \in \mathcal{G}_{d,i}$, we find a $d \times i$ matrix U_1 whose columns form an orthonormal basis for the space spanned by the columns of P; thus, $U_1 U_1^* = P$ (for now, we choose the matrix U_1 in an arbitrary manner, although we will choose U_1 more carefully for our proof of equation 2). We also find a $d \times (d - i)$ matrix U_2 whose columns are orthogonal to the columns of U_1. Thus, $[U_1, U_2]$ is a unitary matrix. This allows us to define the map ϕ by $\phi(P) = U\Lambda U^*$, where $U = [U_1, U_2]$.

To show that ϕ preserves a lower bound on the Frobenius norm distance, use the $\sin$-Θ theorem of Davis and Kahan (1970) (Lemma E.2) which gives a bound on how much the eigenvectors of a Hermitian matrix can "rotate" when the matrix is perturbed. More specifically, the $\sin$-Θ theorem says that if A, A' are Hermitian matrices, with eigenvalues $\lambda_1, \ldots, \lambda_d$ and $\lambda'_1, \ldots, \lambda'_d$, and V_1 and V'_1 are the matrix whose columns are the first i eigenvectors of A and A' respectively, then $\|V_1 V_1^* - V'_1 V_1'^*\|_F \leqslant \frac{\|A - A'\|_F}{\lambda_i - \lambda'_{i+1}}$. Applying the $\sin$-Θ Theorem, for any $P, P' \in \mathcal{G}_{d,i}$ we have

$$\|\phi(P) - \phi(P')\|_F = \|U\Lambda U^* - U'\Lambda U'^*\|_F \geqslant (\lambda_i - \lambda_{i+1}) \times \|P - P'\|_F, \tag{13}$$

for some unitary matrices $U = [U_1, U_2]$ and $U' = [U'_1, U'_2]$ such that $P = U_1 U_1^*$ and $P' = U'_1 U_1'^*$.

Inequality 13 implies that since $P_1, \ldots P_\mathfrak{p}$ is a ζ-packing of $\mathcal{G}_{d,i}$, $\phi(P_1), \ldots \phi(P_\mathfrak{p})$ must be a $\zeta \times (\lambda_i - \lambda_{i+1})$-packing of $\mathcal{O}_\Lambda$. Thus equation 13, together with the bound on the packing number of $\mathcal{G}_{d,i}$, gives the following bound on the packing number of $\mathcal{O}_\Lambda$

$$P(\mathcal{O}_\Lambda, \|\cdot\|_F, \zeta) \geqslant \left(\zeta^{-1} c D_{\|\cdot\|_F}(\mathcal{G}_{d,i}) \times (\lambda_i - \lambda_{i+1})\right)^{2di} \qquad \forall i \in [d]. \tag{14}$$

Improving the packing lower bound. While equation 14 gives a bound for the ζ-packing number of $\mathcal{O}_\Lambda$, the eigenvalue gap term $\lambda_i - \lambda_{i+1}$ may be much smaller than the eigenvalue gap term $\lambda_i - \lambda_j$ which appears in the packing number bounds we ultimately show in Theorem 2.4.

To get around this problem, we replace the map $\phi : \mathcal{G}_{d,i} \to \mathcal{O}_\Lambda$ and instead consider a more general map $\phi : \mathcal{G}_{d-j+i+1,i} \to \mathcal{O}_\Lambda$ for any $i, j \in [d]$. Namely, for any $(d-j+i+1) \times (d-j+i+1)$ rank-i projection matrix $P \in \mathcal{G}_{d-j+i+1,i}$, we choose a matrix U_1 with orthonormal columns such that $U_1 U_1^\star = P$, and choose U_2 such that $[U_1, U_2]$ is a $(d-j+i+1) \times (d-j+i+1)$ unitary matrix. And, denoting by A[i:j] the rows $i, \ldots, j$ of a given matrix A, we set

$$U = \begin{pmatrix} U_1[1:i] & 0 & U_2[1:i] \\ 0 & I & 0 \\ U_1[i+1:d-j+1] & 0 & U_2[i+1:d-j+1] \end{pmatrix} \in \mathrm{U}(d),$$

and set $\phi(P) = U\Lambda U^*$. Then, denoting $\tilde{\Lambda} = \mathrm{diag}(\lambda_1, \ldots, \lambda_i, \lambda_j, \ldots, \lambda_d)$,we have by the sin-$\Theta$ theorem, for any $P, P' \in \mathcal{G}_{d-j+i+1,i}$, that

$$\|\phi(P) - \phi(P')\|_F = \|\hat{U}\tilde{\Lambda}\hat{U}^* - \hat{U}'\tilde{\Lambda}\hat{U}'^*\|_F \geqslant (\lambda_i - \lambda_j)\|P - P'\|_F, \tag{15}$$

for unitary matrices $\hat{U} = [U_1, U_2]$ and $\hat{U}' = [U_1', U_2']$ such that $P = \hat{U}\hat{U}^*$ and $P' = \hat{U}'\hat{U}'^*$. Combining equation 15 with the lower bound for the packing number of the Grassmannian $\mathcal{G}_{d-j+i+1,i}$, we obtain our bound on the packing number for the unitary orbit $\mathcal{O}_\Lambda$ (Inequality 3 in Theorem 2.4).

Packing number lower bounds for $B \cap \mathcal{O}_\Lambda$. To obtain a packing number bound for a subset of the unitary orbit $B(X, \omega) \cap \mathcal{O}_\Lambda$ where $B(X, \omega)$ is some ball of radius ω with center $X \in \mathcal{O}_\Lambda$, we need to ensure that our packing lies inside a ball of radius ω. Towards this end, we first extend the packing number lower bound of Szarek (1982) for the Grassmannian $\mathcal{G}_{d,i}$, to a packing number lower bound for a ball inside the Grassmannian via a simple covering argument (Lemma E.4). Since $\|\cdot\|_F$ is unitarily invariant, the packing number is the same regardless of the center of the ball; thus, for simplicity we set the center of the ball in $\mathcal{G}_{d,i}$ to be the rank-i projection matrix I_i consisting of the first i columns of the identity matrix. While we have already shown that the map ϕ preserves a lower bound on the Frobenius distance $\|\phi(P) - \phi(P')\|_F$ between points in the packing (Inequality 13) to obtain a packing inside $\mathcal{O}_\Lambda$, in order to ensure that the packing lies inside a ball of radius ω we will also need to show that the map ϕ preserves an *upper* bound on this distance.

Unfortunately, if we construct the map $\phi(P)$ by choosing the columns of U_1 to be an arbitrary orthonormal basis for the column space of P and then set $\phi(P) = U\Lambda U^*$ where U is an arbitrary unitary matrix whose first i columns are U_1, we may have that $\|\phi(P) - \phi(P')\|_F > 1$ even when the distance $\|P - P'\|_F$ is arbitrarily small (e.g., if $\mathrm{diag}(\Lambda) = (2, 1, 0)$, $U = I$, and $U' = [e_2, e_1, e_3]R_\theta$, where R_η is a rotation matrix for a small angle $\eta > 0$, and we choose U_1 to be the first 2 columns of U and U' respectively, we have $\phi(U_1U_1^*) - \phi(U_1'U_1'^*) > 1$ and yet $\|U_1U_1^* - U_1'U_1'^*\|_F = \eta$.). This is because there are many ways to choose the basis U_1 for the column space of P.

To show a lower bound on $\|\phi(P) - \phi(P')\|_F$, when constructing the map $\phi(P)$ we will choose the eigenvectors U_1 of P such that, roughly speaking, they correspond to the "principal vectors" between the subspaces spanned by the columns of P and the columns of the projection matrix $P_0 = I_i$ which ϕ maps to the center $X = \Lambda$ of the ball B. We define the principle vectors and principle angles $\theta_1, \ldots, \theta_i$ between any two i-dimensional subspaces $\mathcal{U}$ and $\mathcal{V}$ recursively starting with $\ell = 1$ as follows (see e.g. Björck and Golub (1973)):

$$\theta_\ell = \min\left\{\frac{\arccos|\langle u, v\rangle|}{\|u\|\|v\|} : u \in \mathcal{U}, v \in \mathcal{V}, u \perp u_s, v \perp v_s \forall s \in 1, \ldots, \ell - 1\right\}. \tag{16}$$

Letting $\mathcal{U}$ be the subspace spanned by the columns of any rank-i projection matrix P, and $\mathcal{V}$ the subspace spanned by the columns of $P_0 = I_i$, we set $V_1 = [v_1, \ldots, v_i]$ and $U_1 = [u_1, \ldots, u_i]$ to be the principle vectors between the two subspaces. Thus, roughly speaking, equation 16 implies that we have chosen matrices U_1 and V_1 with the smallest possible angles between the columns of U_1 and the corresponding columns of V_1 under the constraint that $U_1 U_1^* = I_i$ and $V_1 V_1^* = P$. We then define the map to be $\phi(P) = W^* \Lambda W^*$ where W is a unitary matrix whose first i columns are $V_1 U_1^*$, and the last $d - i$ columns are obtained using a similar "principle angle" construction as the first i columns. In particular, we have $\phi(P_0) = \Lambda$.

We then show $\|U_1 - V_1\|_F^2 = 2k - 2\sum_{\ell=1}^{i} \cos(\theta_\ell) \leqslant \|V_1 V_1^* - I_i\|_F^2$, and hence (Lemma E.3),

$$\|V_1 U_1^* \hat{I}_i - \hat{I}_i\|_F \leqslant \|V_1 V_1^* - I_i\|_F = \|P - I_i\|_F$$

where $\hat{I}_i$ is the first i columns of the identity matrix. This in turn implies the bound

$$\|\phi(P) - \Lambda\|_F \leqslant 2\lambda_1 \|W - I\|_F \leqslant 4\lambda_1 \|P - I_i\|_F.$$

We now have an upper bound on the distance $\|\phi(P) - \Lambda\|_F$ between any matrix $\phi(P)$ in our packing and the center Λ of the ball $B(\Lambda, \omega)$ we would like to pack. Combining this bound with our lower bound on $\|\phi(P) - \phi(P')\|_F$ of the previous subsection allows us to show our packing number lower bound for the ball $B(\Lambda, \omega) \cap \mathcal{O}_\Lambda$ (Inequality 2 of Theorem 2.7).

Acknowledgments

This research was conducted in part while NKV was a visiting researcher at Google. OM was supported in part by an NSF grant (CCF-2104528).

References

Kareem Amin, Travis Dick, Alex Kulesza, Andres Muñoz Medina, and Sergei Vassilvitskii. Differentially Private Covariance Estimation. In *NeurIPS*, pages 14190–14199, 2019.

James Bennett and Stan Lanning. The Netflix Prize. In *Proceedings of KDD cup and workshop*, volume 2007, page 35. New York, NY, USA., 2007.

Alex Beutel, Amr Ahmed, and Alexander J. Smola. ACCAMS: Additive Co-Clustering to Approximate Matrices Succinctly. In *WWW*, pages 119–129. ACM, 2015.

Ake Björck and Gene H Golub. Numerical methods for computing angles between linear subspaces. *Mathematics of computation*, 27(123):579–594, 1973.

Avrim Blum, John Hopcroft, and Ravindran Kannan. *Foundations of Data Science*. Cambridge University Press, 2020.

Joseph A. Calandrino, Ann Kilzer, Arvind Narayanan, Edward W. Felten, and Vitaly Shmatikov. "You Might Also Like:" Privacy Risks of Collaborative Filtering. In *IEEE Symposium on Security and Privacy*, pages 231–246. IEEE Computer Society, 2011.

Kamalika Chaudhuri, Anand D. Sarwate, and Kaushik Sinha. A Near-Optimal Algorithm for Differentially-Private Principal Components. *J. Mach. Learn. Res.*, 14(1):2905–2943, 2013.

Chandler Davis and William Morton Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.

Cynthia Dwork. Differential Privacy. In *ICALP (2)*, volume 4052 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2006.

Cynthia Dwork and Aaron Roth. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014.

Cynthia Dwork, Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Analyze Gauss: Optimal Bounds for Privacy-Preserving Principal Component Analysis. In *STOC*, pages 11–20. ACM, 2014.

Alfred Horn. Doubly stochastic matrices and the diagonal of a rotation matrix. *American Journal of Mathematics*, 76(3):620–630, 1954.

Michael Kapralov and Kunal Talwar. On differentially private low rank approximation. In *SODA*, pages 1395–1414. SIAM, 2013.

Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix Factorization Techniques for Recommender Systems. *Computer*, 42(8):30–37, 2009.

Jonathan Leake and Nisheeth K. Vishnoi. Optimization and Sampling Under Continuous Symmetry: Examples and Lie Theory. *CoRR*, abs/2109.01080, 2021.

Jonathan Leake, Colin S. McSwiggen, and Nisheeth K. Vishnoi. Sampling matrices from Harish-Chandra-Itzykson-Zuber densities with applications to quantum inference and differential privacy. In *STOC*, pages 1384–1397. ACM, 2021.

Oren Mangoubi and Nisheeth K. Vishnoi. Sampling from Log-Concave Distributions with Infinity-Distance Guarantees. *CoRR*, abs/2111.04089, 2021.

Frank McSherry and Kunal Talwar. Mechanism Design via Differential Privacy. In *FOCS*, pages 94–103. IEEE Computer Society, 2007.

Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.

Alain Pajor. Metric entropy of the Grassmann manifold. *Convex Geometric Analysis*, 34(181-188): 0942–46013, 1998.

Arkadiusz Paterek. Improving regularized singular value decomposition for collaborative filtering. In *Proceedings of KDD Cup and Workshop*, volume 2007, pages 5–8, 2007.

Badrul Munir Sarwar, George Karypis, Joseph A. Konstan, and John Riedl. Analysis of Recommendation Algorithms for E-Commerce. In *EC*, pages 158–167. ACM, 2000.

Issai Schur. Uber eine klasse von mittelbildungen mit anwendungen auf die determinantentheorie. *Sitzungsberichte der Berliner Mathematischen Gesellschaft*, 22(9-20):51, 1923.

Stanisław J Szarek. Nets of Grassmann manifold and orthogonal group. In *Proceedings of research workshop on Banach space theory (Iowa City, Iowa, 1981)*, volume 169, page 185, 1982.

Stanisław J Szarek. Metric entropy of homogeneous spaces. *Banach Center Publications*, 1(43): 395–410, 1998.

Jalaj Upadhyay. The Price of Privacy for Low-rank Factorization. In *NeurIPS*, pages 4180–4191, 2018.

John Wishart. The generalised product moment distribution in samples from a normal multivariate population. *Biometrika*, pages 32–52, 1928.

Appendix A. Comparison of Our Bounds in Different Examples

In this section we compare our upper bound and lower bound theorems to key prior works. In our notation, the main result of (Amin et al., 2019) can be written as follows.

Theorem A.1 (Amin et al. (2019)) *Given a PSD symmetric covariance matrix $M \in \mathbb{R}^{d\times d}$, a privacy budget $\varepsilon > 0$, and privacy parameters $\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_d$, where $\sum_{i=0}^{d} \varepsilon_i = \varepsilon$. Let the eigenvalues of M be $\lambda_1 \geqslant \cdots \geqslant \lambda_d$. There is a polynomial time algorithm that outputs a matrix $H \in \mathcal{S}_+^d$ such that for any $\beta \in (0,1)$, with probability at least $1-\beta$, $\|M-H\|_F^2 \leqslant \tilde{O}\left(\frac{d}{\varepsilon_0^2} + \sum_{i=1}^{d} \frac{d}{\varepsilon_i}\lambda_i\right)$, where $\tilde{O}$ hides the logarithmic factors of $\frac{1}{\beta}$, d, and λ_i's.*

In the following we provide comparisons to our results for problems where the output matrix has (nearly) the same eigenvalues as the input matrix. For simplicity, we denote by $\lambda_1, \ldots, \lambda_d$ the eigenvalues of the input matrix, and by $\tilde{\lambda}_i$ the (privatized) eigenvalues of the output matrix.

Projection matrices. We first consider the case when the input matrix M is a scalar multiple of a projection matrix of some rank $k > 0$. In this case, the first k eigenvalues of the input matrix all have the same value as the top eigenvalue λ_1, and the remaining $d-k$ eigenvalues are all 0.

Upper bound (Theorem 2.3): When the input matrix is a rank-k projection matrix, Theorem 2.3 gives a bound of $\|M-H\|_F^2 \leqslant \tilde{O}\left(\frac{dk}{\varepsilon^2}\right)$ with probability at least $1-\beta$, where $\tilde{O}$ hides logarithmic factors of $\frac{1}{\beta}$ and k. When the input matrix is a scalar multiple of a rank-k projection matrix with top eigenvalue λ_1, Theorem 2.3 gives a bound of $\|M-H\|_F^2 \leqslant \tilde{O}\left(\frac{dk}{\varepsilon}(\lambda_1 + \frac{1}{\varepsilon})\right)$ with probability at least $1-\beta$.

Upper bound in Amin et al. (2019) (Theorem A.1): When the input matrix is a rank-k projection matrix, the error bound in Theorem A.1 is just $\|M-H\|_F^2 \leqslant E'$, where $E' = \tilde{O}\left(\frac{d}{\varepsilon_0^2} + \sum_{i=1}^{k} \frac{d}{\varepsilon_i}\right)$. This bound is minimized (up to a constant factor) by setting the privacy budget to be $\varepsilon_0 = \frac{\varepsilon}{2}$ and $\varepsilon_i = \frac{\varepsilon}{2k}$ for each $i \geqslant 1$. Hence, $E' \geqslant \tilde{\Omega}(\frac{kd^2}{\varepsilon})$.

More generally, when the input matrix is a scalar multiple of a rank-k projection matrix with top eigenvalue λ_1, the upper bound E' in Theorem A.1 has $E' \geqslant \tilde{\Omega}(\frac{kd^2\lambda_1}{\varepsilon})$. Thus, our bound in Theorem 2.3 is smaller than the bound E' of Theorem A.1 by a factor of $\tilde{O}(d)$.

Lower bounds (Theorem 2.4 and Corollary 2.5): When the input matrix is a scalar multiple of a rank-k projection matrix, our error lower bound is $\|M-H\|_F^2 \geqslant \Omega(\min(\frac{k(d-k)}{\varepsilon}, \lambda_1^2 k))$. Note that the r.h.s. of the lower bound cannot be greater than $k\lambda_1^2$, since $\sup_{M,H\in\mathcal{O}_\Lambda} \|M-H\|_F^2 = O(k\lambda_1^2)$ if Λ is rank-k. In this case our lower bound matches our upper bound up to a factor of $\frac{\lambda_1}{\varepsilon}$ if $\lambda_1 \geqslant \Omega(\sqrt{d})$ and a factor of $\frac{d}{\lambda_1\varepsilon}$ if $\lambda_1 \leqslant O(\sqrt{d})$. While Corollary 2.5 is stated in terms of the Frobenius norm, we can also get a bound for error defined in the spectral norm by using the fact that $\|M - \mathcal{A}(M)\|_F \leqslant \sqrt{2\min(k, d-k)}\|M - \mathcal{A}(M)\|_2$ since M and $\mathcal{A}(M)$ are rank-k matrices. Thus our Corollary 2.5 also implies a lower bound of $\|M - \mathcal{A}(M)\|_2 \geqslant \Omega(\frac{d}{\varepsilon})$ with probability at least $\frac{1}{2}$ when the input matrix is a scalar multiple of a rank-k projection matrix. In comparison, the lower bound from Kapralov and Talwar (2013), which also considers the setting where the input and output of the algorithm are (scalar multiples of) rank-k projection matrices, is $\|M - \mathcal{A}(M)\|_2 \geqslant \Omega(e^{-\frac{\lambda_1\varepsilon}{k(d-k)}}\lambda_1)$. Thus, our lower bound is larger by a factor of $\frac{d}{\varepsilon\lambda_1}e^{\frac{\lambda_1\varepsilon}{k(d-k)}}$.

Matrices with condition number $O(1)$ and large eigenvalue gaps. We consider the case where the eigenvalues $\lambda_1, \ldots, \lambda_d$ of the input matrix M are such that the input matrix has rank k with condition number $\frac{\lambda_1}{\lambda_k} = O(1)$ (and more generally when we may only have $\frac{\lambda_1}{\lambda_{\frac{k}{2}}} = O(1)$) and also has a gap in the eigenvalues of $\lambda_{\frac{k}{4}} - \lambda_{\frac{3k}{4}} = \Omega(\lambda_1)$.

Upper bound (Theorem 2.3): Theorem 2.3 gives a bound of $\|M - H\|_F^2 \leqslant \tilde{O}\left(\frac{dk}{\varepsilon}\left(\lambda_1 + \frac{1}{\varepsilon}\right)\right)$ with probability at least $1-\beta$, where $\tilde{O}$ hides logarithmic factors of $\frac{1}{\beta}$, λ_1, and k.

Upper bound in Amin et al. (2019) (Theorem A.1): By letting $\varepsilon_0 = O(\varepsilon)$ in Theorem A.1, the term in the error due to eigenvalue approximation is the same as $\tilde{O}\left(\frac{d}{\varepsilon_0^2}\right)$ for both algorithms and can thus be ignored. The remaining term in the bound in Theorem 2.3 is $E := \tilde{O}\left(\frac{dk\lambda_1}{\varepsilon}\right)$ and that in Theorem A.1 is $E' := \tilde{O}\left(\sum_{i=1}^{d} \frac{d}{\varepsilon_i}\lambda_i\right)$ which, in turn, depends on how the total privacy budget ε is distributed among the ε_is. With probability $1-\beta$, $\tilde{\lambda}_1 \approx O(\lambda_1 + \frac{1}{\varepsilon}\log\frac{1}{\beta})$. Thus, when $\lambda_1 \geqslant \frac{1}{\varepsilon}\log\frac{1}{\beta}$, $\tilde{\lambda}_1 = \Theta(\lambda_1)$. In this case, $E = \tilde{O}(\frac{dk\lambda_1}{\varepsilon})$. When $\frac{\lambda_1}{\lambda_k} = O(1)$ (or even if we just have the weaker condition that $\frac{\lambda_1}{\lambda_{\frac{k}{2}}} = O(1)$), $E' \geqslant \tilde{\Omega}(d\sum_{i=1}^{\frac{d}{2}} \frac{1}{\varepsilon_i}\lambda_i) \geqslant \tilde{\Omega}(d\lambda_1 \sum_{i=1}^{\frac{d}{2}} \frac{1}{\varepsilon_i})$. Since $\sum_{i=1}^{d} \varepsilon_i = \varepsilon$, the quantity $\sum_{i=1}^{\frac{k}{2}} \frac{1}{\varepsilon_i}$ E' is minimized when $\varepsilon_i := O(\frac{\varepsilon}{k})$ for each $i \leqslant \frac{k}{2}$ and $\varepsilon_i = 0$ for $i > \frac{k}{2}$. Hence, $E' \geqslant \tilde{\Omega}(\frac{d^2k\lambda_1}{\varepsilon})$. Thus, in this case, the bound E from our Theorem 2.3 is $\tilde{O}(d)$ smaller than the bound E' from Theorem A.1.

Lower bounds (Theorem 2.4 and Corollary 2.5): If the input matrix is rank-k ($\lambda_i = 0$ for $i > k$), and $\lambda_{\frac{k}{4}} - \lambda_{\frac{3k}{4}} = \Omega(\lambda_1)$, then the bound in Corollary 2.5 implies that $\|M - \mathcal{A}(M)\|_F^2 \geqslant \Omega(\min(\frac{dk}{\varepsilon}, k\lambda_1^2))$ with probability at least $\frac{1}{2}$. Thus, our lower bound matches our upper bound from Theorem 2.3 up to a factor of $\frac{\lambda_1}{\varepsilon}$ if $\lambda_1 \geqslant \Omega(\sqrt{d})$ and a factor of $\frac{d}{\lambda_1\varepsilon}$ if $\lambda_1 \leqslant O(\sqrt{d})$.

Wishart random matrices. We consider the setting where the input matrix M is a random sample covariance matrix from the Wishart distribution Wishart (1928) (that is $M = \frac{1}{d}X^\top X$, where X is an $m \times d$ matrix with i.i.d. standard Gaussian entries). As in the previous examples, we denote by $\lambda_1, \ldots, \lambda_d$ the eigenvalues of the *input* matrix.

Upper bound (Theorem 2.3): Theorem 2.3 gives a bound of $\|M - H\|_F^2 \leqslant \tilde{O}\left(\frac{dk}{\varepsilon}\left(\lambda_1 + \frac{1}{\varepsilon}\right)\right)$ with probability at least $\frac{1}{2}$ where $\tilde{O}$ hides logarithmic factors of $\frac{1}{\beta}$, λ_1, and k.

Upper bound in Amin et al. (2019) (Theorem A.1): From concentration results for random matrices, we have, with high probability, that $\lambda_1 = O(\lambda_{\frac{k}{2}})$ for any m, d, where $k = \min(m, d)$ is the rank of M. From the discussion in the previous section we have that, whenever $\lambda_1 = O(\lambda_{\frac{k}{2}})$, the bound E' of Theorem A.1 on the error $\|M - H\|_F^2$ satisfies $E' \geqslant \tilde{\Omega}(\frac{d^3\lambda_1}{\varepsilon})$. Thus, if the input matrix is a Wishart random matrix, with high probability, the bound given in our Theorem 2.3 is $\tilde{O}(d)$ smaller than the bound E'.

Lower bound (Theorem 2.4 and Corollary 2.5): From concentration results for random matrices, we also have that, with high probability, there is a large eigenvalue gap $\lambda_{\frac{k}{4}} - \lambda_{\frac{3k}{4}} = \Omega(\lambda_1)$ for any m, d, where $k = \min(m, d)$ is the rank of M. Thus, from the discussion in the previous section, the bound in Corollary 2.5 implies that $\|M - \mathcal{A}(M)\|_F^2 \geqslant \Omega(\min(\frac{dk}{\varepsilon}, k\lambda_1^2))$ with probability at least $\frac{1}{2}$. Thus, our lower bound matches the upper bound from Theorem 2.3 up to a factor of $\frac{\lambda_1}{\varepsilon}$ if $\lambda_1 \geqslant \Omega(\sqrt{d})$ and a factor of $\frac{d}{\lambda_1\varepsilon}$ if $\lambda_1 \leqslant O(\sqrt{d})$.

Appendix B. Preliminaries

B.1. Notation

For any vector $v \in \mathbb{C}^d$, we denote by $\|v\|$ its Euclidean (ℓ_2-norm) and by $\|v\|_p$ its ℓ_p-norm. For any matrix $M \in \mathbb{C}^{m\times n}$, we denote by $\|M\|$ its spectral norm (ℓ_2-operator norm), by $\|M\|_p$ its ℓ_p-operator norm, and by $\|M\|_F$ its Frobenius norm. We use the standard definition in the Euclidean space for inner products. For two vectors $u, v \in \mathbb{C}^d$, we denote the inner product of them as $\langle u, v\rangle := u^*v$. For two matrices $M, N \in \mathbb{C}^{m\times n}$, we denote their Frobenius inner product by $\langle M, N\rangle := \mathrm{Tr}\,(M^*N)$. For any $d \in \mathbb{Z}_+$, we denote by $\mathcal{S}_+^d \subset \mathbb{R}^d$ the set of $d \times d$ positive semi-definite (PSD) real matrices. For any $d \in \mathbb{Z}_+$, we denote by $\mathcal{H}_+^d \subset \mathbb{C}^d$ the set of $d \times d$ PSD Hermitian matrices.

B.2. Preliminaries on Differential Privacy

The Laplace distribution with mean 0 and parameter b is defined over $\mathbb{R}$ as $\mathrm{Lap}(x) := \frac{1}{2b}e^{-|x|}$.

Definition B.1 (Sensitivity) *Given collection of datasets $\mathcal{D}$ with a notion of neighboring datasets, the sensitivity of a query function $q : \mathcal{D} \to \mathbb{R}^d$ is denoted by Δq and defined as*

$$\Delta q := \sup_{\substack{D,D'\in\mathcal{D}\\ D,D' \text{ are neighbors}}} \|q(D) - q(D')\|_1.$$

Theorem B.2 (Laplace mechanism and its differential privacy (Dwork, 2006)) *For a given collection of datasets $\mathcal{D}$ and a privacy budget $\varepsilon > 0$, given any function $f : \mathcal{D} \to \mathbb{R}^d$, define the Laplace mechanism $\mathcal{M} : \mathcal{D} \to \mathbb{R}^d$ as $\mathcal{M}(D) := f(D) + (Y_1, \ldots, Y_d)$, where Y_i's are i.i.d. random variables drawn from $\mathrm{Lap}(\Delta f/\varepsilon)$. Then, $\mathcal{M}$ is ε-differentially private.*

Theorem B.3 (Exponential mechanism (McSherry and Talwar, 2007)) *For a given collection of datasets $\mathcal{D}$ with a notion of neighboring datasets, a measurable set of all possible results $\mathcal{R}$, and a privacy budget $\varepsilon > 0$, given any query function $q : \mathcal{D} \times \mathcal{R} \to \mathbb{R}$, define the exponential mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{R}$ as follows: For any dataset $D \in \mathcal{D}$, $\mathcal{M}(D)$ outputs an $r \in \mathcal{R}$ sampled from a distribution with probability density proportional to*

$$\exp\left(\frac{\varepsilon q(D, r)}{2\Delta q}\right).$$

Then, $\mathcal{M}$ is ε-differentially private.

Theorem B.4 (Utility guarantee for exponential mechanism (McSherry and Talwar, 2007)) *As in the setting in Theorem B.3, given a dataset D, a query function q and privacy budget ε, let $S_t := \{r : q(D, r) > \mathrm{OPT} - t\}$, where $\mathrm{OPT} := \max_r q(D, r)$. Then, we have $\mathbb{P}[r \notin S_t] \leqslant \frac{\exp\left(-\frac{\varepsilon}{2\Delta q}t\right)}{\mu(S_{t/2})}$, where μ is the base measure of the $\mathcal{R}$, the set of all possible results.*

Neighboring datasets. In our setting, $\mathcal{U}$ is the universe of users. For each $u \in \mathcal{U}$, we have a vector $v_u \in \mathbb{C}^d$ such that $\|v_u\|_2 \leqslant 1$. Given a dataset $D \subseteq \mathcal{U}$, define $A := \sum_{u\in D} v_u v_u^*$. Two $d \times d$ Hermitian PSD matrices A and A' are said to be *neighbors* if and only if there exists $u, v \in \mathbb{C}^d$ such that $\|u\|, \|v\| \leqslant 1$ and $A' = A - uu^* + vv^*$.

Appendix C. Differentially Private Optimization on Orbits: Proof of Theorem 2.2

The proof of Theorem 2.2 consists of four parts: the algorithm, its privacy guarantee, its utility guarantee, and its running time.

C.1. Algorithm

We first present the algorithm in Theorem 2.2.

Algorithm 1: Differentially private unitary orbit approximation

Input : A matrix $M \in \mathcal{H}_+^d \subset \mathbb{C}^{d\times d}$ with eigenvalues $\gamma_1 \geqslant \cdots \geqslant \gamma_d \geqslant 0$, the output matrix's maximum rank $k \in [d]$, a list of top k eigenvalues of the output matrix $\lambda_1 \geqslant \cdots \geqslant \lambda_k \geqslant 0$, a privacy budget $\varepsilon > 0$
Output: A matrix $H \in \mathcal{H}_+^d \subset \mathbb{C}^{d\times d}$
Algorithm:

1. Define $\Lambda \leftarrow \mathrm{diag}(\lambda_1, \ldots, \lambda_k, 0, \ldots, 0) \in \mathbb{C}^{d\times d}$
2. Sample $H \in \mathcal{O}_\Lambda$ from a distribution that is $\frac{\varepsilon}{4}$-close in infinity divergence distance to the distribution $d\nu(H) \propto \exp\left(\frac{\varepsilon}{4\lambda_1}\langle M, H\rangle\right) d\mu_\Lambda(H)$
3. Output H

C.2. Privacy guarantee

To prove the privacy guarantee we first need to bound the sensitivity of the utility function $\langle M, H\rangle$.

Lemma C.1 (Sensitivity bound) *Given d and a list of eigenvalues $\lambda_1 \geqslant \cdots \geqslant \lambda_k \geqslant 0$ for some $k \in [d]$, let $\Lambda := \mathrm{diag}(\lambda_1, \ldots, \lambda_k, 0, \ldots, 0) \in \mathbb{C}^{d\times d}$. For any two neighboring $d \times d$ PSD Hermitian $A, A' \in \mathcal{H}_+^d$ such that $A' = A - uu^* + vv^*$ for some u, v such that $\|u\|_2, \|v\|_2 \leqslant 1$, and for any PSD Hermitian matrix $H \in \mathcal{O}_\Lambda$, we have*

$$|\langle A, H\rangle - \langle A', H\rangle| \leqslant \lambda_1.$$

Proof Since $\min(0, \lambda_k) \leqslant v^*Hv \leqslant \max(0, \lambda_1)$ for any v with $\|v\|_2 \leqslant 1$,

$$|\langle A, H\rangle - \langle A', H\rangle| = |u^*Hu - v^*Hv| \leqslant \lambda_1.$$

∎

With Lemma C.1, we can prove the privacy guarantee for Algorithm 1.

Lemma C.2 (Privacy guarantee for Algorithm 1) *The randomized algorithm $\mathcal{M}$ as described in Algorithm 1 is ε-differentially private, for the given privacy budget $\varepsilon > 0$.*

Proof Given neighboring $d \times d$ PSD Hermitian matrices $A, A' \in \mathcal{H}_+^d$ such that $A' = A - uu^* + vv^*$ for some u, v with $\|u\|_2, \|v\|_2 \leqslant 1$, and any matrix $H \in \mathcal{O}_\Lambda$, we want to bound the ratio of probability density of $\mathcal{M}$ at H for A and A'. Let $\tilde{\nu}_A(H)$ be the output density of H for $\mathcal{M}(A)$ and $\nu_A(H)$ be the target density of H, which is given by $\exp(\frac{\varepsilon}{4\lambda_1}\langle A, H\rangle)$. We have $D_\infty(\tilde{\nu}_A \| \nu_A) \leqslant \frac{\varepsilon}{4}$.

$$
\begin{aligned}
\frac{\nu_A(H)}{\nu_{A'}(H)} = \frac{\frac{e^{\frac{\varepsilon}{4}\langle A,H\rangle}}{\int_{Q\in\mathcal{O}_\Lambda} e^{\frac{\varepsilon}{4}\langle A,Q\rangle} d\mu_\Lambda(Q)}}{\frac{e^{\frac{\varepsilon}{4}\langle A',H\rangle}}{\int_{Q\in\mathcal{O}_\Lambda} e^{\frac{\varepsilon}{4}\langle A',Q\rangle} d\mu_\Lambda(Q)}} &= e^{\frac{\varepsilon}{4}\langle A-A',H\rangle} \cdot \frac{\int_{Q\in\mathcal{O}_\Lambda} e^{\frac{\varepsilon}{4}\langle A',Q\rangle} d\mu_\Lambda(Q)}{\int_{Q\in\mathcal{O}_\Lambda} e^{\frac{\varepsilon}{4}\langle A,Q\rangle} d\mu_\Lambda(Q)} \\
&\leqslant e^{\frac{\varepsilon}{4}} \cdot \frac{\int_{Q\in\mathcal{O}_\Lambda} e^{\frac{\varepsilon}{4}\langle A,Q\rangle + \frac{\varepsilon}{2}\langle A'-A,Q\rangle} d\mu_\Lambda(Q)}{\int_{Q\in\mathcal{O}_\Lambda} e^{\frac{\varepsilon}{4}\langle A,Q\rangle} d\mu_\Lambda(Q)} \\
&\leqslant e^{\frac{\varepsilon}{4}} \cdot \max_{Q\in\mathcal{O}_\Lambda} e^{\frac{\varepsilon}{4}|\langle A'-A,Q\rangle|} \\
&\leqslant e^{\frac{\varepsilon}{2}}.
\end{aligned}
$$

Using the infinity divergence bounds between $\tilde{\nu}_A$ and ν_A, we then further have that

$$
\frac{\tilde{\nu}_A(H)}{\tilde{\nu}_{A'}(H)} = \frac{\tilde{\nu}_A(H)/\mu_A(H)}{\tilde{\nu}_{A'}(H)/\nu_{A'}(H)} \cdot \frac{\mu_A(H)}{\mu_{A'}(H)} \leqslant \frac{e^{\frac{\varepsilon}{4}}}{e^{-\frac{\varepsilon}{4}}} \cdot e^{\frac{\varepsilon}{2}} = e^{\varepsilon}.
$$

$$
\frac{\tilde{\nu}_A(H)}{\tilde{\nu}_{A'}(H)} \leqslant \exp(\varepsilon).
$$

■

C.3. Utility guarantee

In this section we prove a guarantee on the utility of Algorithm 1. Towards this end, we first prove a covering number lemma for the unitary orbit.

Lemma C.3 (Covering number for $\mathcal{O}_\Lambda$) *For any $\zeta > 0$, the covering number of $\mathcal{O}_\Lambda$ is at most $N(\mathcal{O}_\Lambda, \|\cdot\|_2, \zeta) \leqslant (1 + \frac{8\lambda_1}{\zeta})^{2dk}$, with Λ defined in Algorithm 1.*

Proof

First consider S_k, the set of $k \times d$ complex matrices with orthonormal rows. Fix any $M \in S_k$ and let U be a unitary matrix such that the first k rows of U are the rows of M. Letting $\|\cdot\|_2$ denote the $2 \to 2$ operator norm, we have that $\|M\|_2 = \|M^*\|_2 = 1$ since M^*M is a PSD projection. Hence, the set S_k can be considered a subset of the unit sphere in a dk-dimensional complex normed vector space. By a standard result (see e.g., Lemma 6.27 in Mohri et al. (2018)), we can cover the complex unit ball in such a space with respect to any norm by at most $(1 + \frac{2}{\zeta})^{2dk}$ balls of radius ζ for any $\zeta > 0$. By replacing each such ball B with a ball of radius 2ζ centered about any $M \in B \cap S_k$ (if such a point exists), we have that we can cover S_k by at most $(1 + \frac{4}{\zeta})^{2dk}$ balls centered in S_k of radius ζ for any $\zeta > 0$.

Consider the map $\phi : M \to M^* \operatorname{diag}(\lambda) M$, which maps S_k to $\mathcal{O}_\Lambda$. Given any M, M' with $\|M - M'\|_2 < \frac{\zeta}{2\lambda_1}$, we have

$$\begin{aligned}
\|\phi(M) - \phi(M')\|_2 &= \|M^* \operatorname{diag}(\lambda) M - M'^* \operatorname{diag}(\lambda) M'\|_2 \\
&\leqslant \|M^* \operatorname{diag}(\lambda)(M - M')\|_2 + \|(M - M')^* \operatorname{diag}(\lambda) M'\|_2 \\
&\leqslant \lambda_1 \frac{\zeta}{2\lambda_1} + \lambda_1 \frac{\zeta}{2\lambda_1} \\
&= \zeta.
\end{aligned}$$

Thus, for any ball B with radius $\frac{\zeta}{2\lambda_1}$ centered at some $M \in S_k$, we have $\phi(B \cap S_k)$ contained in an ζ-ball centered at $\phi(M) \in \mathcal{O}_\Lambda$. Since ϕ is surjective, $\mathcal{O}_\Lambda$ can be covered with at most $(1 + \frac{8\lambda_1}{\zeta})^{2dk}$ balls centered in $\mathcal{O}_\Lambda$ of radius ζ for any $\zeta > 0$. ■

To prove our utility bound, we need the utility bound on the exponential mechanism (Theorem B.4). We use the notation $\Gamma := \sum_{i=1}^d \gamma_i$. The following lemma assumes that we can sample exactly from the distribution proportional to $\exp(\frac{\varepsilon}{4\lambda_1} \langle M, H \rangle)$.

Lemma C.4 (Probability bound assuming exact sampling for the sampling step in Algorithm 1) *Let the input and output be as listed in Algorithm 1. Assume $H \in \mathcal{O}_\Lambda$ is sampled exactly from the distribution $\exp(\frac{\varepsilon}{4\lambda_1} \langle M, H \rangle)$, then we have*

$$\mathbb{P}\left[\langle M, H \rangle \leqslant \sum_{i=1}^k \lambda_i \gamma_i - \tau\right] \leqslant N\left(\mathcal{O}_\Lambda,\ \|\cdot\|_2,\ \frac{\tau}{2\sum_{i=1}^d \gamma_i}\right) \times \exp\left(-\frac{\varepsilon\tau}{4\lambda_1}\right). \tag{17}$$

Proof Using Theorem B.4 in this case, we have

$$S_t = \left\{ H : \langle M, H \rangle > \sum_{i=1}^k \lambda_i \gamma_i - t \right\}.$$

Let $H_0 = U \Lambda U^*$, where U is the unitary matrix obtained by diagonalizing $A = U \operatorname{diag}(\gamma) U^*$. H_0 is the optimal output and we have $\langle A, H_0 \rangle = \sum_{i=1}^k \lambda_i \gamma_i$. We fix any $H \in \mathcal{O}_\Lambda$ such that $\|H_0 - H\|_2 \leqslant \frac{\tau}{2\Gamma}$. We can then apply the Hölder's inequality to get

$$\begin{aligned}
\langle M, H \rangle &= \langle M, H_0 \rangle - \langle M, H_0 - H \rangle \\
&\geqslant \sum_{i=1}^k \lambda_i \gamma_i - \|H_0 - H\|_2 \operatorname{Tr}(M) && \text{(Using Hölder's inequality)} \\
&= \sum_{i=1}^k \lambda_i \gamma_i - \|H_0 - H\|_2 \Gamma && \text{(Substitute definitions)} \\
&\geqslant \sum_{i=1}^k \lambda_i \gamma_i - \frac{\tau}{2\Gamma} \Gamma && \text{(Substitute } \|H_0 - H\|_2\text{)} \\
&\geqslant \sum_{i=1}^k \lambda_i \gamma_i - \frac{\tau}{2}.
\end{aligned}$$

Thus, if $\|H_0 - H\|_2 \leqslant \frac{\tau}{2\Gamma}$, then $H \in S_{\frac{\tau}{2}}$. Thus, every H contained in the ball of radius $\frac{\tau}{2\Gamma}$ centered at M is also in $S_{\frac{\tau}{2}}$. Let μ be the unitarily invariant probability measure on $\mathcal{O}_\Lambda$. By the definition of covering number, the number of balls centered in $\mathcal{O}_\Lambda$ of radius $\frac{\tau}{2\Gamma}$ required to cover the set $\mathcal{O}_\Lambda$ is at most $N(\mathcal{O}_\Lambda, \|\cdot\|_2, \frac{\tau}{2\Gamma})$.Thus, there exists some ball $B_{\frac{\tau}{2\Gamma}}(H')$ centered at H' with $\mu(B_{\frac{\tau}{2\Gamma}}(H')) \geqslant N(\mathcal{O}_\Lambda, \|\cdot\|_2, \frac{\tau}{2\Gamma})^{-1}$. Thus, since μ is unitarily invariant,

$$\begin{aligned}\mu\left(S_{\frac{\tau}{2}}\right) &\geqslant \mu(B_{\frac{\tau}{2\Gamma}(M)})\\ &= \mu(B_{\frac{\tau}{2\Gamma}(H')})\\ &\geqslant N(\mathcal{O}_\Lambda, \|\cdot\|_2, \frac{\tau}{2\Gamma})^{-1}.\end{aligned}$$

Using Theorem B.4, with query function $q(M,H) = \langle M, H\rangle$, we have

$$\begin{aligned}\mathbb{P}\left[\langle M, H\rangle \leqslant \sum_{i=1}^k \lambda_i\gamma_i - \tau\right] &= \mathbb{P}\left[H \notin S_\tau\right]\\ &\leqslant \frac{\exp\left(-\frac{\varepsilon}{4\lambda_1}\tau\right)}{\mu\left(S_{\frac{\tau}{2}}\right)}\\ &\leqslant N(\mathcal{O}_\Lambda, \|\cdot\|_2, \frac{\tau}{2\Gamma}) \times \exp\left(-\frac{\varepsilon\tau}{4\lambda_1}\right)\end{aligned}$$

■

This gives the following alternative form of the utility bound.

Lemma C.5 (Utility bound for Algorithm 1) *Let the input and output be as listed in Algorithm 1. For any $\beta \in (0,1)$, with probability at least $1-\beta$, the randomized algorithm $\mathcal{M}$ in Algorithm 1 outputs a matrix $H \in \mathcal{O}_\Lambda$ satisfying*

$$\sum_{i=1}^k \gamma_i\lambda_i - \langle M, H\rangle \leqslant O\left(\frac{\lambda_1}{\varepsilon}\left(dk\log\sum_{i=1}^d \gamma_i + \log\frac{1}{\beta}\right)\right).$$

Proof We can choose a suitable τ to give a utility bound on $\langle M, H\rangle$. By letting

$$\tau = \frac{2\lambda_1}{\varepsilon}\log\left(e + \frac{(2+8\Gamma)^{4dk}}{\beta}\right) + \lambda_1,$$

since $\varepsilon \in (0,1)$, we have

$$\tau > \frac{2\lambda_1}{\varepsilon} > \lambda_1.$$

Thus,

$$\frac{8\lambda_1\Gamma}{\tau} \leqslant 8\Gamma. \tag{18}$$

Since H is sampled from a distribution which is $\frac{\varepsilon}{4}$-close to the distribution $\exp(\frac{\varepsilon}{4\lambda_1}\langle M, H\rangle)$ in infinity divergence, by plugging in this choice of τ together with the covering number bound of Lemma C.3 into equation 17, we have

$$\begin{aligned}
\mathbb{P}\left[\langle M, H\rangle \leqslant \sum_{i=1}^{k} \lambda_i \gamma_i - \tau\right] &\leqslant \exp\left(\frac{\varepsilon}{4}\right) \times \exp\left(-\frac{\varepsilon\tau}{4\lambda_1}\right) \times N\left(\mathcal{O}_\Lambda,\ \|\cdot\|_2,\ \frac{\tau}{2\sum_{i=1}^{d}\gamma_i}\right) && \text{(From equation 17)}\\
&\leqslant \exp\left(-\frac{\varepsilon(\tau-\lambda_1)}{4\lambda_1}\right)\left(1+\frac{16\lambda_1\Gamma}{\tau}\right)^{2dk} && \text{(From Lemma C.3)}
\end{aligned}$$

We then substitute τ,

$$\begin{aligned}
&\mathbb{P}\left[\langle M, H\rangle \leqslant \sum_{i=1}^{k} \lambda_i \gamma_i - O\left(\frac{\lambda_1}{\varepsilon}\left(dk\log\Gamma + \log\frac{1}{\beta}\right)\right)\right] && \text{(Substitute } \tau\text{)}\\
&\leqslant \mathbb{P}\left[\langle M, H\rangle \leqslant \sum_{i=1}^{k} \lambda_i \gamma_i - \tau\right]\\
&\leqslant \exp\left(-\frac{\varepsilon(\tau-\lambda_1)}{4\lambda_1}\right)\left(1+\frac{16\lambda_1\Gamma}{\tau}\right)^{2dk}\\
&\leqslant \exp\left(-\frac{\varepsilon(\tau-\lambda_1)}{4\lambda_1}\right)(1+16\Gamma)^{2dk} && \text{(From equation 18)}\\
&= \exp\left(-\log\left(e+\frac{(1+16\Gamma)^{4dk}}{\beta}\right)\right)(1+16\Gamma)^{2dk} && \text{(Substitute } \tau\text{)}\\
&= \left(e+\frac{(1+16\Gamma)^{4dk}}{\beta}\right)^{-1}(1+16\Gamma)^{2dk}\\
&\leqslant \frac{\beta}{(1+16\Gamma)^{4dk}}(1+16\Gamma)^{2dk}\\
&= \beta.
\end{aligned}$$

Thus, with probability at least $1-\beta$, we have

$$\langle M, H\rangle \geqslant \sum_{i=1}^{k} \lambda_i \gamma_i - O\left(\frac{\lambda_1}{\varepsilon}\left(dk\log\Gamma + \log\frac{1}{\beta}\right)\right).$$

∎

C.4. Running time

Lemma C.6 (Running time for Algorithm 1) *The number of arithmetic operations required by the algorithm Algorithm 1 is polynomial in* $\log\frac{1}{\varepsilon}$, λ_1, $\gamma_1 - \gamma_d$, *and the number of bits representing* $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_k)$ *and* $\gamma = (\gamma_1, \gamma_2, \ldots, \gamma_d)$.

Proof This follows directly by using the algorithm from Corollary 2.7 of (Mangoubi and Vishnoi, 2021) in Algorithm 1. ∎

C.5. Completing the proof of Theorem 2.2

Proof (of Theorem 2.2) The privacy Guarantee for Algorithm 1 is provided in Lemma C.2. The utility guarantee is Lemma C.5. The running time bound for Algorithm 1 is from Lemma C.6.

■

Appendix D. Differentially Private Rank-k Approximation: Proof of Theorem 2.3

Our algorithm in the proof of Theorem 2.3 has two parts. The first part approximates the eigenvalues of M and the second part is just Algorithm 1.

D.1. Algorithm

Algorithm 2: Differentially private rank-k approximation

Input : A data matrix $M \in \mathcal{H}_+^d \subset \mathbb{C}^{d\times d}$, the rank of output matrix $k \in [d]$, a privacy budget $\varepsilon > 0$

Output: A matrix $H \in \mathcal{H}_+^d \subset \mathbb{C}^{d\times d}$

Algorithm:

1. Compute the eigenvalues of M and let them be $\lambda_1 \geqslant \cdots \geqslant \lambda_d$.
2. Compute $\tilde{\lambda}_i \leftarrow \lambda_i + \mathrm{Lap}\left(\frac{4}{\varepsilon}\right)$, for all $i \in [k]$
3. Sort $\tilde{\lambda}_i$s so that $\tilde{\lambda}_1 \geqslant \cdots \geqslant \tilde{\lambda}_k$
4. Define $\Lambda \leftarrow \mathrm{diag}(\tilde{\lambda}_1, \ldots, \tilde{\lambda}_k, 0, \ldots, 0) \in \mathbb{C}^{d\times d}$
5. Sample $H \in \mathcal{O}_\Lambda$ from a distribution that is $\frac{\varepsilon}{8}$-close in infinity divergence distance to the distribution $d\nu(H) \propto \exp\left(\frac{\varepsilon}{8\lambda_1}\langle M, H\rangle\right) d\mu(H)$
6. Output H and the list of estimated eigenvalues $\tilde{\lambda}_1, \ldots, \tilde{\lambda}_k$

This algorithm has two parts. The first part (Step 1 to 2) approximates eigenvalues and is shown to be $\frac{\varepsilon}{2}$-differentially private in Theorem D.1. The second part (Step 3 to 6) is Algorithm 1 with privacy budget $\frac{\varepsilon}{2}$.

D.2. First part: Differentially private eigenvalue approximation

Theorem D.1 (Differentially private approximation of eigenvalues) *Given a positive semidefinite (PSD) Hermitian input matrix $M \in \mathcal{H}_+^d$ and a privacy budget $\varepsilon > 0$. Let the eigenvalues of M be $\lambda_1, \ldots, \lambda_d \in \mathbb{R}$. Outputting $\tilde{\lambda}_1, \ldots, \tilde{\lambda}_d$, where $\tilde{\lambda}_i = \lambda_i + \mathrm{Lap}\left(\frac{2}{\varepsilon}\right)$ is an ε-differentially private algorithm for approximating the eigenvalues of M. In addition, for any $i \in [d]$, $\mathbb{E}[\tilde{\lambda}_i] = \lambda_i$. With probability at least $1-\beta$, $|\tilde{\lambda}_i - \lambda_i| = O\left(\frac{1}{\varepsilon}\log\frac{1}{\beta}\right)$ for all i.*

Since we need to deal with the eigenvalues, we use the following notation: For any matrix $M \in \mathbb{C}^{d\times d}$, we denote $\lambda(M) := (\lambda_1(M), \dots, \lambda_d(M))$ as its eigenvalues with $\lambda_1(M) \geqslant \cdots \geqslant \lambda_d(M)$. To prove Theorem D.1, we need the following lemma.

Lemma D.2 (Inequality for eigenvalues) *Given a positive semi-definite Hermitian matrix $M \in \mathcal{H}_+^d \subset \mathbb{C}^{d\times d}$ and a vector $v \in \mathbb{C}^d$. Let $A := M - vv^*$. For any $i \in [d]$, $\lambda_i(A) \leqslant \lambda_i(M)$. In addition, $\|\lambda(M) - \lambda(A)\|_1 = \|v\|_2^2$.*

Proof Let $\mathbb{S}^{d-1}$ be the sphere of unit vectors in $\mathbb{C}^d$. For any $u \in \mathbb{S}^{d-1}$, we have

$$u^*Au = u^*Mu - u^*vv^*u = u^*Mu - (v^*u)^2 \leqslant u^*Mu.$$

Thus, pick any $i \in [d]$ and any subspace $U \subseteq \mathbb{C}^d$ with dimension i, we have

$$\min_{u\in U\cap\mathbb{S}^{d-1}} u^*Au \leqslant \min_{u\in U\cap\mathbb{S}^{d-1}} u^*Mu. \tag{19}$$

Thus, using the min-max theorem (Courant–Fischer–Weyl min-max principle), we have

$$\begin{aligned}
\lambda_i(A) &= \max_{U\subseteq\mathbb{C}^{d\times d}:\ \dim(U)=i}\ \min_{u\in U\cap\mathbb{S}^{d-1}} u^*Au \\
&\leqslant \max_{U\subseteq\mathbb{C}^{d\times d}:\ \dim(U)=i}\ \min_{u\in U\cap\mathbb{S}^{d-1}} u^*Mu \qquad \text{(equation 19 holds for any } U \in \mathbb{C}^{d\times d}) \\
&= \lambda_i(M).
\end{aligned}$$

This leads to $\lambda_i(A) \leqslant \lambda_i(M)$ for any $i \in [d]$. In addition,

$$\begin{aligned}
\|\lambda(M) - \lambda(A)\|_1 &= \sum_{i=1}^d |\lambda_i(M) - \lambda_i(A)| \\
&= \sum_{i=1}^d (\lambda_i(M) - \lambda_i(A)) \qquad (\lambda_i(A) \leqslant \lambda_i(M)) \\
&= \mathrm{Tr}(M) - \mathrm{Tr}(A) \\
&= \mathrm{Tr}(vv^*) \\
&= \|v\|_2^2.
\end{aligned}$$

■

Using this lemma, we can then prove Theorem D.1.

Proof (of Theorem D.1) Given two neighboring PSD Hermitian matrix $M, M' \in \mathcal{H}_+^d \subset \mathbb{C}^{d\times d}$, so that there exist $u, v \in \mathbb{C}^d$ with $\|u\|_2, \|v\|_2 \leqslant 1$ such that $M' = M - uu^* + vv^*$. Let $A := M - uu^*$. Using Lemma D.2, we have

$$\|\lambda(M) - \lambda(A)\|_1 = \|u\|_2^2 \leqslant 1.$$

Similarly, we have

$$\|\lambda(M') - \lambda(A)\|_1 = \|v\|_2^2 \leqslant 1.$$

Thus,

$$\|\lambda(M) - \lambda(M')\|_1 \leqslant \|\lambda(M) - \lambda(A)\|_1 + \|\lambda(M') - \lambda(A)\|_1 \leqslant 2. \tag{20}$$

Thus, for any neighboring PSD Hermitian matrix $M, M' \in \mathcal{H}_+^d$, the ℓ_1 distance between their eigenvalue vector is at most 2. According to Defintion B.1, the sensitivity of the eigenvalue computation is 2. Thus, outputting $\tilde{\lambda}_i = \lambda_i + \mathrm{Lap}\left(\frac{2}{\varepsilon}\right)$ follows exactly the Laplace mechanism in Theorem B.2. Thus, the eigenvalue approximation satisfies ε-differential privacy.

In addition, for any $\beta \in (0, 1)$, with probability $1 - \beta$,

$$\left|\mathrm{Lap}\left(\frac{2}{\varepsilon}\right)\right| \leqslant \frac{2}{\varepsilon} \ln \frac{1}{\beta} = O\left(\frac{1}{\varepsilon} \log \frac{1}{\beta}\right).$$

Thus, with probability at least $1 - \beta$,

$$\left|\tilde{\lambda}_i - \lambda_i\right| = O\left(\frac{1}{\varepsilon} \log \frac{1}{\beta}\right),$$

for all $i \in [d]$. ∎

D.3. Second part: Completing the proof

Combining Theorem D.1 and Theorem 2.2, we can prove Theorem 2.3.

Proof (of Theorem 2.3)

Running time: From Theorem D.1, the number of arithmetic operations required by the first part of Theorem 2 (eigenvalue approximations) is $\tilde{O}(d)$. From Theorem 2.2, the number of arithmetic operations required by the second part of the algorithm is polynomial in $\log \frac{1}{\varepsilon}$, λ_1, and the number of bits representing $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_d)$.

Privacy guarantee: The first part (eigenvalue approximation) in Algorithm 2 is done by letting each $\tilde{\lambda}_i := \lambda_i + \mathrm{Lap}\left(\frac{4}{\varepsilon}\right)$. Theorem D.1 implies that this approximation is $\frac{\varepsilon}{2}$-differentially private. The second part of Algorithm 2 is just Algorithm 1 where we set the privacy budget to be $\frac{\varepsilon}{2}$. Thus, from Lemma C.2, the second part is $\frac{\varepsilon}{2}$-differentially private. From the composition theorem of differential privacy (Dwork and Roth, 2014), it follows that Algorithm 2 is ε-differentially private.

Utility bound: From Theorem D.1, it follows that for any $\beta \in (0, 1)$, with probability at least $1 - \beta$, for all $i \in [k]$, we have $|\tilde{\lambda}_i - \lambda_i| \leqslant O\left(\frac{1}{\varepsilon} \log \frac{1}{\beta}\right)$. Note that the values $\tilde{\lambda} = (\tilde{\lambda}_1, \ldots, \tilde{\lambda}_k)$ may not be sorted. Let $\hat{\lambda}$ be the vector generated by sorting the entries of $\tilde{\lambda}$ in non-increasing order. For the second part, we note that Theorem 2.2 (applied with $\hat{\lambda}_i$s for γ_js) implies that, for any $\beta \in (0, 1)$, with probability at least $1 - \beta$,

$$\langle M, H \rangle \geqslant \sum_{i=1}^{k} \lambda_i \hat{\lambda}_i - O\left(\frac{\hat{\lambda}_1}{\varepsilon}\left(dk \log \sum_{i=1}^{d} \lambda_i + \log \frac{1}{\beta}\right)\right) \tag{21}$$

Thus, with probability at least $1-\beta$,

$$\begin{aligned}
\|M-H\|_F^2 &= \|M\|_F^2 + \|H\|_F^2 - 2\langle M, H\rangle \\
&\leqslant \sum_{i=1}^{d} \lambda_i^2 + \sum_{i=1}^{k} \hat{\lambda}_i^2 - 2\sum_{i=1}^{k} \lambda_i \hat{\lambda}_i + 2O\left(\frac{\hat{\lambda}_1}{\varepsilon}\left(dk \log \sum_{i=1}^{d} \lambda_i + \log \frac{1}{\beta}\right)\right) \\
&= \sum_{i=k+1}^{d} \lambda_i^2 + \sum_{i=1}^{k}\left(\lambda_i^2 + \hat{\lambda}_i^2 - 2\sum_{i=1}^{k} \lambda_i \hat{\lambda}_i\right) + O\left(\frac{\hat{\lambda}_1}{\varepsilon}\left(dk \log \sum_{i=1}^{d} \lambda_i + \log \frac{1}{\beta}\right)\right) \\
&= \sum_{i=k+1}^{d} \lambda_i^2 + \sum_{i=1}^{k}\left(\lambda_i - \hat{\lambda}_i\right)^2 + O\left(\frac{\hat{\lambda}_1}{\varepsilon}\left(dk \log \sum_{i=1}^{d} \lambda_i + \log \frac{1}{\beta}\right)\right). \qquad (22)
\end{aligned}$$

We would like to replace $\hat{\lambda}_1$ in the last term by λ_1 and use Theorem D.1 to prove an upper bound on the second term. For the former, observe that: Since with probability at least $1-\beta$, $\left|\tilde{\lambda}_i - \lambda_i\right| = O\left(\frac{1}{\varepsilon}\log\frac{1}{\beta}\right)$ for all i, and $\hat{\lambda}_1 = \max_i \tilde{\lambda}_i$, it follows that with probability at least $1-\beta$, $\hat{\lambda}_1 = \lambda_1 + O\left(\frac{1}{\varepsilon}\log\frac{1}{\beta}\right)$. For the latter, notice that for any vector $v \in \mathbb{R}^k$, $\sum_{i=1}^{k} (\lambda_i - v_i)^2$ is minimized when the entries in v are sorted in non-increasing order. Using these, along with equation 22, we get that

$$\begin{aligned}
\|M-H\|_F^2 &\leqslant \sum_{i=k+1}^{d} \lambda_i^2 + \sum_{i=1}^{k}\left(\lambda_i - \tilde{\lambda}_i\right)^2 + O\left(\frac{\lambda_1 + \frac{1}{\varepsilon}\log\frac{1}{\beta}}{\varepsilon}\left(dk \log \sum_{i=1}^{d} \lambda_i + \log \frac{1}{\beta}\right)\right) \\
&\leqslant \sum_{i=k+1}^{d} \lambda_i^2 + kO\left(\frac{1}{\varepsilon^2}\log^2\frac{1}{\beta}\right) + O\left(\frac{\lambda_1 + \frac{1}{\varepsilon}\log\frac{1}{\beta}}{\varepsilon}\left(dk \log \sum_{i=1}^{d} \lambda_i + \log \frac{1}{\beta}\right)\right) \\
&= \sum_{i=k+1}^{d} \lambda_i^2 + O\left(\frac{k}{\varepsilon^2}\log^2\frac{1}{\beta} + \frac{\lambda_1 + \frac{1}{\varepsilon}\log\frac{1}{\beta}}{\varepsilon}\left(dk \log \sum_{i=1}^{d} \lambda_i + \log \frac{1}{\beta}\right)\right) \\
&= \sum_{i=k+1}^{d} \lambda_i^2 + \tilde{O}\left(\frac{k}{\varepsilon^2} + \frac{dk}{\varepsilon}\left(\lambda_1 + \frac{1}{\varepsilon}\right)\right),
\end{aligned}$$

where $\tilde{O}$ hides logarithmic factors of $\sum_{i=1}^{d} \lambda_i$ and $\frac{1}{\beta}$. The above equation uses the fact that with probability at least $1-\beta$, $\left|\tilde{\lambda}_i - \lambda_i\right| = O\left(\frac{1}{\varepsilon}\log\frac{1}{\beta}\right)$ for all i. Thus, we have proved the utility bounds for Algorithm 2.

Combining the running time, the privacy guarantee, and the utility bound, we have proved Theorem 2.3. ■

Appendix E. Packing Number Lower Bound: Proof of Theorem 2.7

We will use the following result from (Szarek, 1982) which bounds the covering number $N(\mathcal{G}_{d,k}, \zeta)$ for the (complex) Grassmannian $\mathcal{G}_{d,k}$ with respect to the metric induced by the operator norm on the projection matrices P_V for the subspaces $V \in \mathcal{G}_{d,k}$ (see Proposition 8 of (Szarek, 1998), and the note about about extension to complex spaces).

Lemma E.1 (Covering number of (complex) Grassmannian $G_{d,k}$ (Szarek, 1982, 1998)) *There exist universal constants $C > c > 0$ such that for every unitarily invariant norm $|||\cdot|||$ and every $0 < \zeta < D_{|||\cdot|||}$, the covering number $N(\mathcal{G}_{d,k}, |||\cdot|||, \zeta)$ of the (complex) Grassmannian $G_{d,k}$ satisfies*

$$\left(\frac{cD_{|||\cdot|||}(\mathcal{G}_{d,k})}{\zeta}\right)^{2k(d-k)} \leqslant N(\mathcal{G}_{d,k}, |||\cdot|||, \zeta) \leqslant \left(\frac{CD_{|||\cdot|||}(\mathcal{G}_{d,k})}{\zeta}\right)^{2k(d-k)},$$

where $D_{|||\cdot|||}(\mathcal{G}_{d,k}) := \sup_{\mathcal{U},\mathcal{V}\in\mathcal{G}_{d,k}} |||P_{\mathcal{U}} - P_{\mathcal{V}}|||$ is the diameter of $\mathcal{G}_{d,k}$ with respect to $|||\cdot|||$.

In the case of the Frobenius norm, we have $D_{\|\cdot\|_F} = \sup_{\mathcal{U},\mathcal{V}\in G_{d,k}} \|P_{\mathcal{U}} - P_{\mathcal{V}}\|_F \geqslant \min(\sqrt{k}, \sqrt{d-k})$, and in the case of the operator norm we have $D_{\|\cdot\|_2} = \sup_{\mathcal{U},\mathcal{V}\in G_{d,k}} \|P_{\mathcal{U}} - P_{\mathcal{V}}\|_2 \geqslant 1$
We will also make use of the following Sin-Θ theorem of (Davis and Kahan, 1970):

Let A and $\hat{A}$ be two Hermitian matrices with eigenvalue decompositions

$$A = U\Lambda U^* = (U_1, U_2)\begin{pmatrix} \Lambda_1 & \\ & \Lambda_2 \end{pmatrix}\begin{pmatrix} U_1^* \\ U_2^* \end{pmatrix} \tag{23}$$

$$\hat{A} = \hat{U}\hat{\Lambda}\hat{U}^* = (\hat{U}_1, \hat{U}_2)\begin{pmatrix} \hat{\Lambda}_1 & \\ & \hat{\Lambda}_2 \end{pmatrix}\begin{pmatrix} \hat{U}_1^* \\ \hat{U}_2^* \end{pmatrix}, \tag{24}$$

(although when we apply the Sin-Theta theorem we will only need the special case where $\hat{\Lambda} = \Lambda$).

Lemma E.2 (sin-Θ Theorem (Davis and Kahan, 1970)) *Let $A, \hat{A}$ be two Hermitian matrices with eigenvalue decompositions given in equation 23 and equation 24. Suppose that there are $\alpha > \beta > 0$ and $\Delta > 0$ such that the spectrum of Λ_1 is contained in the interval $[\alpha, \beta]$ and the spectrum of $\hat{\Lambda}_2$ lies entirely outside of the interval $(\alpha - \Delta, \beta + \Delta)$. Then*

$$|||U_1U_1^* - \hat{U}_1\hat{U}_1^*||| \leqslant \frac{|||\hat{A} - A|||}{\Delta},$$

where $|||\cdot|||$ denotes the operator norm or Frobenius norm (or, more generally, any unitarily invariant norm).

Lemma E.3 *Suppose that I_k is the $d \times d$ diagonal matrix with the first k diagonal entries 1 and the remaining $d-k$ diagonal entries 0, and let $\hat{I}_k$ be the first k columns of I_k. Let P be any Hermitian rank-k projection matrix. Then there exists a $d \times k$ matrix $\hat{W}$ with orthonormal columns such that $\hat{W}\hat{W}^* = P$ and $\|\hat{W} - \hat{I}_k\|_F \leqslant \|P - I_k\|_F$.*

Proof Denote by $\mathcal{I}_k$ the column space of I_k and $\mathcal{P}$ the column space of P. Let $\theta_1 \leqslant \cdots \leqslant \theta_k$ be the k principal angles between $\mathcal{I}_k$ and $\mathcal{P}$. Let $u_1, \ldots, u_k$ and $v_1, \ldots, v_k$, form an orthonormal basis for $\mathcal{I}_k$ and $\mathcal{P}$ respectively, and where the angles between corresponding vectors u_i and v_i in the two bases are equal to the i'th principle angle θ_i for every $i \in [k]$. The existence of such a basis is guaranteed by the the variational definition of principle angles between subspaces (see e.g. Björck and Golub (1973)).

Let $\mathcal{I}_k^{\perp}$ and $\mathcal{P}^{\perp}$ be the orthogonal complements of $\mathcal{I}_k$ and $\mathcal{P}$, respectively. Let $u_{k+1}, \ldots, u_d$ be a basis for $\mathcal{I}_k^{\perp}$, and let $v_{k+1}, \ldots, v_d$ be a basis for $\mathcal{P}^{\perp}$. Let $U_1 = [u_1, \ldots, u_k]$ and $V_1 = [v_1, \ldots, v_k]$. And let $U_2 = [u_{k+1}, \ldots, u_d]$ and $V_2 = [v_{k+1}, \ldots, v_d]$ Let $U = [U_1, U_2]$ and $V = [V_1, V_2]$.

Therefore, we have that

$$\begin{aligned}\|U_1 - V_1\|_F^2 &= \mathrm{tr}((U_1 - V_1)^*(U_1 - V_1)) \\ &= \mathrm{tr}((U_1^*U_1 - U_1^*V_1 - V_1^*U_1 + V_1^*V_1) \\ &\mathrm{tr}(U_1^*U_1) - 2\mathrm{tr}(U_1^*V_1) + \mathrm{tr}(V_1^*V_1) \\ &= 2k - 2\mathrm{tr}(U_1^*V_1) \\ &= 2k - 2\sum_{i=1}^{k} u_i^* v_i \\ &= 2k - 2\sum_{i=1}^{k} \cos(\theta_i). \end{aligned} \tag{25}$$

But, by the variational definition of principal angles, we also have that the largest singular values of $I_k P^*$ are also $\cos(\theta_1) \geqslant \cos(\theta_2) \geqslant \cdots \geqslant \cos(\theta_k)$, with the remaining singular values equal to 0.

Therefore, we have that

$$\begin{aligned}\|V_1V_1^* - I_k\|_F^2 &= \|V_1V_1^* - U_1U_1^*\|_F^2 \\ &= \mathrm{tr}((V_1V_1^* - U_1U_1^*)^*(V_1V_1^* - U_1U_1^*)) \\ &= \mathrm{tr}(V_1V_1^*)^2) - \mathrm{tr}((U_1U_1^*)(V_1V_1^*)) - \mathrm{tr}((V_1V_1^*)(U_1U_1^*)) + \mathrm{tr}((U_1U_1^*)^2) \\ &= 2k - 2\mathrm{tr}((U_1U_1^*)(V_1V_1^*)) \\ &= 2k - 2\mathrm{tr}(I_kP^*) \\ &\geqslant 2k - 2\sum_{i=1}^{k} \cos(\theta_i), \end{aligned} \tag{26}$$

where the inequality holds since $\mathrm{tr}((U_1U_1^*)(V_1V_1^*))$ is the sum of the eigenvalues of $(U_1U_1^*)(V_1V_1^*)$, and the sum of the singular values of any matrix is at least as large as the sum of its eigenvalues. Therefore, combining equation 25 a equation 26 we have that

$$\|U_1 - V_1\|_F \leqslant \|V_1V_1^* - I_k\|_F. \tag{27}$$

But, since $V_1 = V_1U_1^*U_1$ we also have that

$$\begin{aligned}\|V_1 - U_1\|_F^2 &= \|V_1U_1^*U_1 - U_1U_1^*U_1\|_F \\ &= \|(V_1U_1^* - I_k)U_1\|_F^2 \\ &= \mathrm{tr}(((V_1U_1^* - I_k)U_1)^*(V_1U_1^* - I_k)U_1) \\ &= \mathrm{tr}(U_1^*(V_1U_1^* - I_k)^*(V_1U_1^* - I_k)U_1) \\ &= \mathrm{tr}((V_1U_1^* - I_k)^*(V_1U_1^* - I_k)U_1U_1^*) \\ &= \mathrm{tr}((V_1U_1^* - I_k)^*(V_1U_1^* - I_k)) \\ &= \|V_1U_1^* - I_k\|_F^2 \\ &= \|V_1U_1^*U_1U_1^* - I_k\|_F^2 \\ &= \|V_1U_1^*I_k - I_k\|_F^2 \\ &= \|V_1U_1^*\hat{I}_k - \hat{I}_k\|_F^2. \end{aligned} \tag{28}$$

Therefore, combining equation 27 and equation 28 we have that

$$\|V_1U_1^*\hat{I}_k - \hat{I}_k\|_F \leqslant \|V_1V_1^* - I_k\|_F \tag{29}$$

Now, since U is a unitary matrix, VU^* is also unitary. But, since $U_1U_1^* = I_k$, U_1 must have all zeros below the $k'th$ row, and U_2 must have all zeros above the $k+1$'st row. Therefore, the first k columns of VU^* are the same as the first k columns of $V_1U_1^*$. Thus, the first k columns of $V_1U_1^*$ must be orthonormal to each other. Since the columns of $V_1U_1^*\hat{I}_k$ are the same as the first k columns of $V_1U_1^*$, the columns of $V_1U_1^*\hat{I}_k$ must be orthonormal to each other as well. Therefore, setting $\hat{W} = V_1U_1^*\hat{I}_k$, and since $V_1V_1^* = P$, from equation 29 we have that

$$\|\hat{W} - \hat{I}_k\|_F \leqslant \|P - I_k\|_F$$

where $\hat{W} = V_1U_1^*\hat{I}_k$, is a matrix with orthonormal columns.

The remaining $d-k$ orthonormal columns of the unitary matrix W (with the first k columns being the columns of $\hat{W}$) can be found by diagonalizing the projection matrix for the subspace $\mathcal{P}^\perp$.

■

Lemma E.4 (Packing number of ball inside $\mathcal{O}_\Lambda$) *For any $M \in \mathcal{O}_\Lambda$, any unitarily invariant norm $|||\cdot|||$, and any $\delta > r > 0$, and denoting by $B(x, |||\cdot|||, \delta)$ a ball of radius δ centered at M with respect to the norm m $|||\cdot|||$, we have*

$$P(B(M, |||\cdot|||, \delta) \cap \mathcal{O}_\lambda, |||\cdot|||, r) \geqslant \frac{P(\mathcal{O}_\Lambda, |||\cdot|||, r)}{N(\mathcal{O}_\Lambda, |||\cdot|||, \delta)}.$$

Proof Let $M, M' \in \mathcal{O}_\Lambda$. Then there exists $U \in U(d)$ such that $M' = UMU^*$. Since $|||\cdot|||$ is unitarily invariant, $|||W_1 - W_2||| = |||U(W_1 - W_2)U^*||| = |||UW_1U^* - UW_2U^*|||$ for any $W_1, W_2 \in \mathcal{O}_\Lambda$. Thus, for any $n \in \mathbb{N}$, we have that a collection of matrices $M_1, \ldots, M_n$ is an r-packing of $B(M, |||\cdot|||, \delta)$ if and only if $UM_1U^*, \ldots, UM_nU^*$ is an r-packing of $B(UMU^*, |||\cdot|||, \delta)$. Thus, for every $M, M' \in \mathcal{O}_\Lambda$ we have that

$$P(B(M, |||\cdot|||, \delta), |||\cdot|||, r) = P(B(M', |||\cdot|||, \delta), |||\cdot|||, r). \tag{30}$$

Let $M_1, \ldots, M_\alpha$, where $\alpha \geqslant P(\mathcal{O}_\Lambda, |||\cdot|||, r)$, be an r-packing of $\mathcal{O}_\Lambda$. And let $C_1, \ldots, C_\beta$ where $C_1, \ldots, C_\beta$ are balls of radius δ centered in $\mathcal{O}_\Lambda$ and $\beta \leqslant N(\mathcal{O}_\Lambda, |||\cdot|||, \delta)$, be a δ-covering of $\mathcal{O}_\Lambda$. Then by the pigeonhole principle there exists $i \in [\beta]$ such that a subset of $M_1, \ldots, M_\alpha$ of size $\geqslant \frac{\alpha}{\beta}$ is an r-packing of C_i. Hence, $P(C_i, r) \geqslant \frac{\alpha}{\beta}$ and by equation 30 we have that

$$\begin{aligned} P(B(M, |||\cdot|||, \delta), |||\cdot|||, r) &= P(C_i, r) \\ &\geqslant \frac{\alpha}{\beta} \\ &\geqslant \frac{P(\mathcal{O}_\Lambda, |||\cdot|||, r)}{N(\mathcal{O}_\Lambda, |||\cdot|||, \delta)}. \end{aligned} \tag{31}$$

■

For any $\zeta > 0$ and norm $\|\cdot\|$, we define the packing number $P(S, \|\cdot\|, \zeta)$ of a subset $S \subset E$ of a normed space E with metric ρ induced by the norm $\|\cdot\|$ to be the largest subset $\{x_1, \ldots, x_n\}$, for any $n \in \mathbb{N}$, such that $\rho(x_i, x_j) > \zeta$.

We will now use the covering number for the Grassmannian and the Sin-Θ theorem to prove Theorem 2.7.

Proof (of Theorem 2.7)
Bounding the packing number of a ball in the Grassmannian: Denote by $I_k \in \mathcal{P}_{d,k}$, the matrix with its first k diagonal entries 1 and each all other entries 0. Plugging in the upper and lower bounds for the covering number in E.1 into Lemma E.4, for any $k > 0$, we get that the packing number of any ball of radius ζ inside any Grassmannian manifold $\mathcal{G}_{d,k}$ (which we represent by the set of rank-k projection matrices $\mathcal{P}_{d,k}$), with center $I_k \in \mathcal{P}_{d,k}$ [2], satisfies

$$\begin{aligned}
P(B(I_k, \|\cdot\|_F, \omega) \cap \mathcal{P}_{d,k}, \|\cdot\|_F, \zeta) &\overset{\text{Lemma E.4}}{\geqslant} \frac{P(\mathcal{P}_{d,k}, \|\cdot\|_F, \zeta)}{N(\mathcal{P}_{d,k}, \|\cdot\|_F, \omega)} \\
&= \frac{P(\mathcal{P}_{d,k}, \|\cdot\|_F, \zeta)}{N(\mathcal{P}_{d,k}, \|\cdot\|_F, \omega)} \\
&\geqslant \frac{N(\mathcal{G}_{d,k}, \|\cdot\|_F, 2\zeta)}{N(\mathcal{G}_{d,k}, \|\cdot\|_F, \omega)} \\
&\overset{\text{Lemma E.1}}{\geqslant} \left(\frac{\min(\omega, D_{\|\cdot\|_F}(\mathcal{G}_{d,k})c}{2\zeta C}\right)^{2k(d-k)} \\
&\geqslant \left(\frac{\min(\omega, \sqrt{k}, \sqrt{d-k})c}{2\zeta C}\right)^{2k(d-k)}. \qquad (32)
\end{aligned}$$

Constructing the map from the Grassmannian to the orbit: For any projection matrix $M \in \mathcal{P}_{d-j+i+1,\, i}$ define a map $\psi : \mathcal{P}_{d-j+i+1,\, i} \to \mathcal{U}(d-j+i+1)$, from $\mathcal{P}_{d-j+i+1,\, i}$ to the group of $(d-j+i+1) \times (d-j+i+1)$ unitary matrices $\mathcal{U}(d-j+i+1)$, as follows:

- $\psi(I_i) = I$, where I_i is the $(d-j+i+1) \times (d-j+i+1)$ diagonal matrix with the first i diagonal entries 1 and all other entries 0, and I is the $(d-j+i+1) \times (d-j+i+1)$ identity matrix.

- $\psi(M) = U$, where $U \in \mathcal{U}(d-j+i+1)$ is a unitary matrix such that its first i columns U_1 satisfy $U_1 U_1^* = M$, and $\|U - I\|_F \leqslant 2\|M - I_i\|_F$.

We still need to show that a matrix $\psi(M) = U$ satisfying the above conditions exists. We can construct the matrix $\psi(M) = U$ by applying Lemma E.3 twice. First, we apply Lemma E.3 which guarantees the existence of matrix U_1 with orthonormal columns such that $U_1 U_1^* = M$ and $\|U_1 - \hat{I}_i\|_F \leqslant \|M - I_i\|_F$. Next, we apply Lemma E.3 a second time to obtain a matrix U_2 with orthonormal columns such that $U_2 U_2^* = I - M$ is a projection matrix for the orthogonal complement of the space spanned by the columns of M, and $\|U_2 - (\hat{I} - \hat{I}_i)\|_F \leqslant \|(I - M) - (I - I_i)\|_F = \|M - I_i\|_F$.

2. Note that the choice of center here is arbitrary, and we would get the same bound regardless of choice of center since $\|\cdot\|_F$ is unitarily invariant.

Define the matrix $U := [U_1, U_2]$. Then we have that

$$\begin{aligned}\|U - I\|_F &\leqslant \|U_1 - \hat{I}_i\|_F + \|U_2 - (\hat{I} - \hat{I}_i)\|_F \\ &\leqslant 2\|M - I_i\|_F.\end{aligned}$$

Moreover, since $U_1 U_1^* = M$ and $U_2 U_2^* = I - M$, we have that the columns of U_1 and U_2 are orthogonal to each other and hence that the matrix U is a $(d - j + i + 1) \times (d - j + i + 1)$ unitary matrix.

Showing that the map preserves Frobenius norm distance (lower bound): For convenience, we denote the submatrix of any matrix H consisting of the entries in rows $k, \ldots, \ell$ by $H[k : \ell]$. For convenience, in the remainder of the proof, we denote the restriction of ψ to the first i columns of its output by $\psi_1(M) = U_1$. And we denote the last $d - j + 1$ columns of U by U_2, and the restriction of ψ to these columns by $\psi_2(M) = U_2$.

Next, consider the map $\Psi : \mathcal{P}_{i,d-j+i+1} \to \mathcal{U}(d)$ defined as follows:

$$\Psi(M) := \begin{pmatrix} \psi_1(M)[1 : i] & 0 & \psi_2(M)[1 : i] \\ 0 & I_{(j-i-1)\times(j-i-1)} & 0 \\ \psi_1(M)[i + 1 : d - j + 1] & 0 & \psi_2(M)[i + 1 : d - j + 1] \end{pmatrix},$$

where $I_{(j-i-1)\times(j-i-1)}$ denotes the $(j - i - 1) \times (j - i - 1)$ identity matrix. And define the map $\phi : \mathcal{P}_{i,d-j+i+1} \to \mathcal{O}_\Lambda$ as follows:

$$\phi(M) = \Psi(M)\, \Lambda\, \Psi(M)^* \tag{33}$$

Define $\tilde{\Lambda} := \operatorname{diag}(\lambda_1, \ldots, \lambda_i, \lambda_j, \ldots, \lambda_d)$. For any projection matrices $M, M' \in \mathcal{P}_{i,d-j+i+1}$, we have

$$\begin{aligned}\|\phi(M) - \phi(M')\|_F &= \|\Psi(M)\Lambda\Psi(M)^* - \Psi(M')\Lambda\Psi(M')^*\|_F \\ &= \|\psi(M)\tilde{\Lambda}\psi(M)^* - \psi(M')\tilde{\Lambda}\psi(M')^*\|_F \\ &\geqslant (\lambda_i - \lambda_j) \times \|\psi_1(M)\psi_1(M)^* - \psi_1(M')\psi_1(M')^*\|_F \\ &= (\lambda_i - \lambda_j) \times \|M - M'\|_F, \end{aligned} \tag{34}$$

where the inequality holds by the Sin-Θ Theorem of (Davis and Kahan, 1970) (restated above as Lemma E.2), since $\|\cdot\|_F$ is a unitarily invariant norm.

Showing that the map preserves Frobenius norm distance (upper bound): Moreover, we also have that

$$\begin{aligned}
\|\phi(M) - \Lambda\|_F &= \|\phi(M) - \phi(I_i)\|_F \\
&= \|\Psi(M)\Lambda\Psi(M)^* - \Psi(I_i)\Lambda\Psi(I_i)^*\|_F \\
&= \|\psi(M)\tilde{\Lambda}\psi(M)^* - \psi(I_i)\tilde{\Lambda}\psi(I_i)^*\|_F \\
&= \|\psi(M)\tilde{\Lambda}\psi(M)^* - I\tilde{\Lambda}I^*\|_F \\
&= \|(\psi(M) - I)\tilde{\Lambda}\psi(M)^* - I\tilde{\Lambda}(I^* - \psi(M))\|_F \\
&\leqslant \|(\psi(M) - I)\tilde{\Lambda}\psi(M)^*\|_F + \|I\tilde{\Lambda}(I^* - \psi(M))\|_F \\
&= 2\|(\psi(M) - I)\tilde{\Lambda}\psi(M)^*\|_F \\
&\leqslant 2\|(\psi(M) - I)\tilde{\Lambda}\|_F \times \|\psi(M)^*\|_2 \\
&\leqslant 2\|(\psi(M) - I)\tilde{\Lambda}\|_F \\
&\leqslant 2\|\psi(M) - I\|_F \times \|\tilde{\Lambda}\|_2 \\
&\leqslant 2\lambda_1\|\psi(M) - I\|_F, \qquad (35)
\end{aligned}$$

where the second and fourth inequalities hold because the Freobenius norm is sub-multiplicative with respect to the operator norm, and the third inequality holds because $\|\psi(M)^*\|_2 = 1$ since $\psi(M)^*$ is a unitary matrix.

Bounding the packing number (subset of orbit): From equation 32, we have that

$$P\left(B\left(I_k, \|\cdot\|_F, \frac{\omega}{2\lambda_1}\right) \cap \mathcal{P}_{i,d-j+i+1},\ \|\cdot\|_F,\ \frac{\zeta}{\lambda_i - \lambda_j}\right) \geqslant J,$$

where

$$J = \left(\frac{\min(\omega, \lambda_1\sqrt{i}, \lambda_1\sqrt{d-j+1})c \times (\lambda_i - \lambda_j)}{2\lambda_1\zeta C}\right)^{2i\times(d-j+1)}.$$

Therefore, we have that there exists a $\{M_1, \dots M_J\} \subseteq \mathcal{P}_{i,d-j+i+1}$ of size J such that,

$$\|M_s - M_t\|_F > \frac{\zeta}{\lambda_i - \lambda_j} \qquad \forall s,t \in [J] \qquad (36)$$

and

$$\|M_s - I_k\|_F < \frac{\omega}{2\lambda_1} \qquad \forall s \in [J]. \qquad (37)$$

Therefore, plugging by equation 34 into that equation 36 we have that,

$$\|\phi(M_s) - \phi(M_t)\|_F > \zeta \qquad \forall s,t \in [J] \qquad (38)$$

and, moreover, plugging equation 35 into equation 36 we have that

$$\|\phi(M_s) - \Lambda\|_F < \omega \qquad \forall s \in [J]. \qquad (39)$$

Since by equation 33, $\phi(M_1), \ldots \phi(M_J)$ are all in the unitary orbit $\mathcal{O}_\Omega$, equation 38 and equation 39 imply that $\phi(M_1), \ldots \phi(M_J)$ is a ζ packing for $B(\Lambda, \omega) \cap \mathcal{O}_\Lambda$. Therefore, the packing number of $B(\Lambda, \omega) \cap \mathcal{O}_\Lambda$ is

$$P(B(\Lambda, \omega) \cap \mathcal{O}_\Lambda, \|\cdot\|, \zeta) \geqslant J. \tag{40}$$

But since $\|\cdot\|_F$ is unitarily invariant, we have that

$$P(B(\Lambda, \omega) \cap \mathcal{O}_\Lambda, \|\cdot\|_F, \zeta) = P(B(X, \omega) \cap \mathcal{O}_\Lambda, \|\cdot\|, \zeta) \qquad \forall X \in \mathcal{O}_\Lambda. \tag{41}$$

Therefore, equation 40 and equation 41 together imply that

$$P(B(X, \omega) \cap \mathcal{O}_\Lambda, \|\cdot\|_F, \zeta) \leqslant \left(\frac{\min(\omega, \lambda_1\sqrt{i}, \lambda_1\sqrt{d-j+1})c \times (\lambda_i - \lambda_j)}{2\lambda_1 \zeta C}\right)^{2i \times (d-j+1)} \tag{42}$$

for every $1 \leqslant i < j \leqslant d$. Therefore, we have that

$$\begin{aligned} &\log P(B(X, \omega) \cap \mathcal{O}_\Lambda, \|\cdot\|_F, \zeta) \\ &\qquad \geqslant \max_{1 \leqslant i < j \leqslant d} 2i \times (d-j+1) \times \log\left(\frac{\min(\omega, \lambda_1\sqrt{i}, \lambda_1\sqrt{d-j+1})c \times (\lambda_i - \lambda_j)}{2\lambda_1 \zeta C}\right). \end{aligned}$$

This completes the proof of equation 2.

Bounding the packing number (entire orbit, with slightly stronger bound): We can get a slightly better bound when bounding the entire unitary orbit, using the following argument. Since

$$P(\mathcal{P}_{i,d-j+i+1}, \|\cdot\|_F, \frac{\zeta}{\lambda_i - \lambda_j}) \geqslant \left(\frac{cD_{\|\cdot\|_F}(\mathcal{G}_{d-j+i+1,\, i}) \times (\lambda_i - \lambda_j)}{\zeta}\right)^{2i \times (d-j+1)},$$

there exits a subset $\{M_1, \ldots M_n\} \subseteq \mathcal{P}_{i,d-j+i+1}$ of size $n = \left(\frac{cD_{\|\cdot\|}(\mathcal{G}_{d-j+i+1,\, i}) \times (\lambda_i - \lambda_j)}{\zeta}\right)^{2i \times (d-j+1)}$ such that $\|M_r - M_s\|_F > \frac{\zeta}{\lambda_i - \lambda_j}$ for all $r, s \in [n]$. Thus, by equation 34 we have that

$$\|\phi(M_r) - \phi(M_s)\|_F > \zeta, \qquad \forall r, s \in [n]. \tag{43}$$

Since we have a subset $\{\phi(M_1), \ldots \phi(M_n)\} \subseteq \mathcal{O}_\Lambda$ such that $\|\phi(M_r) - \phi(M_s)\| > \zeta$ for all $r, s \in [n]$, the packing number of $\mathcal{O}_\Lambda$ satisfies

$$P(\mathcal{O}_\Lambda, \|\cdot\|_F, \zeta) \geqslant n = \left(\frac{cD_{\|\cdot\|_F}(\mathcal{G}_{d-j+i+1,\, i}) \times (\lambda_i - \lambda_j)}{\zeta}\right)^{2i \times (d-j+1)}. \tag{44}$$

Finally, since equation 44 holds for every choice of $1 \leqslant i < j \leqslant d$, and $D_{\|\cdot\|_F}(\mathcal{G}_{d-j+i+1,\, i}) \geqslant \Omega(\min(\sqrt{i}, \sqrt{d-j+1}))$, we have that

$$P(\mathcal{O}_\Lambda, \|\cdot\|_F, \zeta) \geqslant \max_{1 \leqslant i < j \leqslant d} \left(\frac{c \min(\sqrt{i}, \sqrt{d-j+1}) \times (\lambda_i - \lambda_j)}{\zeta}\right)^{2i \times (d-j+1)}.$$

This completes the proof of equation 3. ■

Appendix F. Lower Bound on Utility: Proof of Theorem 2.4

To prove the lower bound on the utility, we will also use the following lemma:

Lemma F.1 *For any $U, V \in \mathcal{U}(d)$ we have*

$$\|U\Lambda U^* - V\Lambda V^*\|_F^2 = 2\langle U\Lambda U^*,\ U\Lambda U^* - V\Lambda V^*\rangle$$

Proof

$$\begin{aligned}
\|U\Lambda U^* - V\Lambda V^*\|_F^2 &= \langle U\Lambda U^* - V\Lambda V^*,\ U\Lambda U^* - V\Lambda V^*\rangle \\
&= \mathrm{tr}((U\Lambda U^* - V\Lambda V^*)^*(U\Lambda U^* - V\Lambda V^*))) \\
&= \mathrm{tr}((U\Lambda U^* - V\Lambda V^*)^2) \\
&= \mathrm{tr}(U\Lambda^2 U^* - U\Lambda U^* V\Lambda V^* - V\Lambda V^* U\Lambda U^* + V\Lambda^2 V^*) \\
&= 2\mathrm{tr}(\Lambda^2) - 2\mathrm{tr}(U\Lambda U^* V\Lambda V^*) \\
&= 2\mathrm{tr}(U\Lambda^2 U^* - U\Lambda U^* V\Lambda V^*) \\
&= 2\mathrm{tr}(U\Lambda U^*(U\Lambda U^* - V\Lambda V^*)).
\end{aligned}$$

■

Lemma F.2 (Lower utility bound for unitary orbit, as a function of packing number) *Suppose, for some $\alpha > \eta > 0$, that $\lambda_1 \geqslant \cdots \geqslant \lambda_d \geqslant 0$ and $\varepsilon > 0$ are such that*

$$\sum_{\ell=1}^{d} \lambda_\ell^2 < \frac{1}{16\varepsilon\delta\alpha^2} \log(P(B(W, 2\alpha r) \cap \mathcal{O}_\Lambda, \|\cdot\|_F, 2\eta r))) - \frac{d}{16\delta\alpha^2}$$

for some $\delta > 0$, where we define $r := \sqrt{\delta \sum_{\ell=1}^{d} \lambda_\ell^2}$ and $W \in \mathcal{O}_\Lambda$ is any matrix in the unitary orbit.[3] *Then for any ε-differentially private algorithm $\mathcal{A}$ which takes as input a Hermitian matrix and outputs a matrix in the orbit $\mathcal{O}_\Lambda$, there exists a Hermitian matrix M with eigenvalues $\lambda_1 \geqslant \cdots \geqslant \lambda_d$ such that the output $\mathcal{A}(M)$ of the algorithm satisfies*

$$\sum_{\ell=1}^{d} \lambda_\ell^2 - \langle M, \mathcal{A}(M)\rangle \geqslant \eta^2\delta \sum_{\ell=1}^{d} \lambda_\ell^2 \tag{45}$$

with probability at least $\frac{1}{2}$.

Proof Let $W \in \mathcal{O}_\Lambda$ be such that $\langle M, W\rangle = \sum_{\ell=1}^{d} \lambda_\ell^2$. Define $r := 2\sqrt{\delta \sum_{i=1}^{d} \lambda_i^2}$. By the definition of packing number, there exists a $2\eta r$-packing $E = \{U_i \Lambda U_i^a st\}_{i=1}^n$ of $B(W, 2\alpha r)$ for $n = P(B(W, 2\alpha r), \|\cdot\|_F, 2\eta r))$ such that for every $i, j \in [n]$, $i \neq j$, we have

$$B(U_i\Lambda U_i^*, \eta r) \cap B(U_j\Lambda U_j^*, \eta r) = \varnothing.$$

3. The choice of W does not matter since unitary invariance means that the packing bound depend only on the radius of the ball we are packing, not its center.

Consider the matrices $M_i = U_i \Lambda U_i^*$ for each $i \in [n]$. We would like to show that equation 45 holds for one of these matrix M_i. Suppose, on the contrary, that for every $i \in [n]$, the output of the algorithm, $\mathcal{A}(M_i) = V_i \Lambda V_i^*$ (where we denote by $V_i \in \mathcal{U}(d)$ a unitary matrix which diagonalizes M_i), satisfies

$$\langle M_i, \mathcal{A}(M_i) \rangle > (1 - \eta^2 \delta) \sum_{\ell=1}^{d} \lambda_\ell^2, \tag{46}$$

with probability at least $\frac{1}{2}$. Let E_i be the event that equation 46 is satisfied. The $\mathbb{P}(E) \geqslant \frac{1}{2}$.

Suppose that the event E_i occurs. Then, since $\langle M_i, M_i \rangle = \sum_{\ell=1}^{d} \lambda_\ell^2$, we have that

$$\langle M_i, \mathcal{A}(M_i) - M_i \rangle = \langle M_i, \mathcal{A}(M_i) \rangle - \langle M_i, M_i \rangle > -\eta^2 \delta \sum_{\ell=1}^{d} \lambda_\ell^2,$$

Therefore by Lemma F.1 we have that

$$\begin{aligned}
\delta \sum_{\ell=1}^{d} \lambda_\ell^2 &> \langle M_i,\ M_i - \mathcal{A}(M_i) \rangle \\
&= \langle U_i \Lambda U_i^*,\ U_i \Lambda U_i^* - V_i \Lambda V_i^* \rangle \\
&= \frac{1}{2} \| U_i \Lambda U_i^* - V_i \Lambda V_i^* \|_F^2 \\
&= \frac{1}{2} \| M_i - \mathcal{A}(M_i) \|_F^2.
\end{aligned} \tag{47}$$

That is,

$$\| M_i - \mathcal{A}(M_i) \|_F^2 < 2\eta^2 \delta \sum_{\ell=1}^{d} \lambda_\ell^2.$$

whenever the event E_i occurs. Thus, since $\mathbb{P}(E_i) \geqslant \frac{1}{2}$, for every $i \in [n]$, we have that

$$\mathbb{P}\left(\| M_i - \mathcal{A}(M_i) \|_F^2 < 2\eta^2 \delta \sum_{\ell=1}^{d} \lambda_\ell^2 \right) \geqslant \frac{1}{2}.$$

Hence, for every $i \in [n]$,

$$\mathbb{P}\left(\| M_i - \mathcal{A}(M_i) \|_F < 2\eta \sqrt{\delta \sum_{\ell=1}^{d} \lambda_\ell^2} \right) \geqslant \frac{1}{2}. \tag{48}$$

Inequality 48 implies that the output $\mathcal{A}(M_i)$ of Algorithm $\mathcal{A}$ falls inside the Frobenius-norm ball $B(M_i, r)$ with probability at least $\frac{1}{2}$.

For any $i, j \in [n]$, we have that $M_i - M_j = \sum_{s=1}^{m} x_s x_s^*$ for some $m \leqslant \|M_i - M_j\|_F^2 + d$ and data vectors $x_1, \ldots, x_m \in \mathbb{C}^d$ for which $\|x_s\|_2 \leqslant 1$. Thus, one can modify the data matrix M_i into any other data matrix M_j by replacing at most $\|M_i - M_j\|_F^2 + d$ points in the dataset.

Since by assumption, Algorithm $\mathcal{A}$ is ε-differentially private, we have that for any $i, j \in [n]$,

$$\begin{aligned}
\frac{\mathbb{P}(\mathcal{A}(M_i) \in B(M_j, \eta r))}{\mathbb{P}(\mathcal{A}(M_i) \in B(M_i, \eta r))} &\geqslant e^{-\varepsilon(\|M_i - M_j\|_F^2 + d)} \\
&\geqslant e^{-\varepsilon(d + 16\alpha^2 r^2)},
\end{aligned} \tag{49}$$

since $M_i, M_j \in B(W, 2\alpha r)$. But by equation 48 we have that $\mathbb{P}(\mathcal{A}(M_i) \in B(M_i, \eta r)) > \frac{1}{2}$. Therefore, equation 49 implies that

$$\mathbb{P}(\mathcal{A}(M_i) \in B(M_j, \eta r)) \geqslant \frac{1}{2} e^{-\varepsilon(d+16\alpha^2 r^2)} \qquad \forall i \in [n]. \tag{50}$$

Since $M_1, \ldots, M_n$ is a $2\eta r$-packing of $B(W, 2\alpha r)$, the balls $B(M_j, \eta r)$, $j \in [n]$, are pairwise disjoint. Thus,

$$1 \geqslant \sum_{j=1}^{n} \mathbb{P}(\mathcal{A}(M_i) \in B(M_j, \eta r)) \geqslant n \times e^{-\varepsilon(d+16\alpha^2 r^2)} \tag{51}$$

Rearranging equation 51, we have that

$$\log(n) \leqslant \varepsilon(d + 16\alpha^2 r^2)$$

and hence that

$$\frac{1}{16\varepsilon\alpha^2} \log(n) - \frac{d}{16\alpha^2} \leqslant r^2 = \delta \sum_{\ell=1}^{d} \lambda_\ell^2.$$

In other words,

$$\sum_{\ell=1}^{d} \lambda_\ell^2 \geqslant \frac{1}{16\varepsilon\delta\alpha^2} \log(P(B(W, 2\alpha r), \|\cdot\|_F, 2\eta r))) - \frac{d}{16\delta\alpha^2} \tag{52}$$

Inequality 52 contradicts the theorem statement. Thus, our assumption that

$$\mathbb{P}\left(\langle M_i, \mathcal{A}(M_i)\rangle \geqslant (1 - \eta^2\delta) \sum_{\ell=1}^{d} \lambda_\ell^2\right) \geqslant \frac{1}{2}$$

for every $i \in [n]$ is false, and we therefore have that for some $i \in [n]$ the utility for the matrix M_i satisfies

$$\langle M_i, \mathcal{A}(M_i)\rangle < (1 - \eta^2\delta) \sum_{\ell=1}^{d} \lambda_\ell^2$$

with probability at least $\frac{1}{2}$.

∎

Proof (of Theorem 2.4) Consider any $1 \leqslant i < j \leqslant d$. Set $\omega, \zeta, \alpha, \eta, r, \delta$ as follows

$$\zeta = \min(\omega, \lambda_1\sqrt{i}, \lambda_1\sqrt{d-j+1}) \times \frac{(\lambda_i - \lambda_j)}{4C\lambda_1}, \tag{53}$$

and $\alpha = \frac{\omega}{2r}$, $\eta = \frac{\zeta}{2r}$, and $r = \sqrt{\delta \sum_{\ell=1}^{d} \lambda_\ell^2}$. Then we have $\delta = \frac{r^2}{\sum_{\ell=1}^{d} \lambda_\ell^2} = \frac{\zeta^2}{4\eta^2 \sum_{\ell=1}^{d} \lambda_\ell^2}$ Thus, by Lemma 2.7, we have that for any $W \in \mathcal{O}_\Lambda$,

$$
\begin{aligned}
&\frac{1}{16\varepsilon\delta\alpha^2}\log(P(B(W,2\alpha r)\cap\mathcal{O}_\Lambda,\|\cdot\|_F,2\eta r))) - \frac{d}{16\delta\alpha^2}\\
&\geqslant \frac{\sum_{\ell=1}^d \lambda_\ell^2}{4\varepsilon\omega^2}\log(P(B(W,\omega)\cap\mathcal{O}_\Lambda,\|\cdot\|_F,\zeta))) - \frac{d}{\omega^2}\sum_{\ell=1}^d\lambda_\ell^2\\
&\geqslant \frac{\sum_{\ell=1}^d \lambda_\ell^2}{4\varepsilon\omega^2} i\times(d-j+1)\log(2) - \frac{d}{\omega^2}\sum_{\ell=1}^d\lambda_\ell^2 \qquad (54)
\end{aligned}
$$

Consider the following equation for any $\delta > 0$:

$$
\sum_{\ell=1}^d \lambda_\ell^2 < \frac{1}{16\varepsilon\delta\alpha^2}\log(P(B(W,2\alpha r)\cap\mathcal{O}_\Lambda,\|\cdot\|_F,2\eta r))) - \frac{d}{16\delta\alpha^2}. \qquad (55)
$$

By plugging equation 54 into equation 55, we get that equation 55 holds for any $\omega > 0$ such that

$$
\sum_{\ell=1}^d \lambda_\ell^2 < \frac{\sum_{\ell=1}^d \lambda_\ell^2}{4\varepsilon\omega^2} i\times(d-j+1)\log(2) - \frac{d}{\omega^2}\sum_{\ell=1}^d\lambda_\ell^2 \qquad (56)
$$

Rearranging equation 56 we get

$$
\omega^2 < \frac{1}{4\varepsilon} i\times(d-j+1)\log(2) - d.
$$

Thus, by Lemma F.2 we have that for any ε-differentially private algorithm $\mathcal{A}$ which takes as input a Hermitian matrix and outputs a matrix in the orbit $\mathcal{O}_\Lambda$, there exists a Hermitian matrix M with eigenvalues $\lambda_1 \geqslant \cdots \geqslant \lambda_d$ such that, with probability at least $\frac{1}{2}$, the output $\mathcal{A}(M)$ of the algorithm satisfies

$$
\begin{aligned}
\sum_{\ell=1}^d \lambda_\ell^2 - \langle M, \mathcal{A}(M)\rangle &\geqslant \eta^2\delta\sum_{\ell=1}^d\lambda_\ell^2\\
&= \frac{\zeta^2}{4}\\
&\overset{\text{Eq. 53}}{=} \min(\omega^2, \lambda_1^2 i, \lambda_1^2(d-j+1)) \times \frac{(\lambda_i-\lambda_j)^2}{64C^2\lambda_1^2}\\
&= \min\left(\frac{1}{4\varepsilon\lambda_1^2} i\times(d-j+1)\log(2) - \frac{d}{\lambda_1^2},\ i,\ d-j+1\right)\times\frac{(\lambda_i-\lambda_j)^2}{64C^2}. \qquad (57)
\end{aligned}
$$

Since equation 57 holds for any $1 \leqslant i < j \leqslant d$, we must have that any ε-differentially private algorithm $\mathcal{A}$ which takes as input a Hermitian matrix and outputs a matrix in the orbit $\mathcal{O}_\Lambda$, there exists a Hermitian matrix M with eigenvalues $\lambda_1 \geqslant \cdots \geqslant \lambda_d$ such that, with probability at least $\frac{1}{2}$, the output $\mathcal{A}(M)$ of the algorithm satisfies

$$\sum_{\ell=1}^{d} \lambda_\ell^2 - \langle M \mathcal{A}(M) \geqslant \max_{1 \leqslant i < j \leqslant d} \frac{(\lambda_i - \lambda_j)^2}{64C^2} \times \min\left(\frac{1}{4\varepsilon\lambda_1^2} i \times (d-j+1)\log(2) - \frac{d}{\lambda_1^2},\ i,\ d-j+1\right). \tag{58}$$

Plugging Lemma F.1 into equation 58, we get

$$\|M - \mathcal{A}(M)\|_F^2 \geqslant \max_{1 \leqslant i < j \leqslant d} \frac{(\lambda_i - \lambda_j)^2}{64C^2} \times \min\left(\frac{1}{4\varepsilon\lambda_1^2} i \times (d-j+1) - \frac{d}{\lambda_1^2},\ i,\ d-j+1\right) \tag{59}$$

Taking the maximum over only pairs (i, j) where $1 \leqslant i < j \leqslant d$ and either $j = \frac{d}{2}$, or $i = \frac{d}{2}$, and adjusting the universal constant C, we get

$$\|M - \mathcal{A}(M)\|_F^2 \geqslant \max_{1 \leqslant i \leqslant \frac{d}{2}} \frac{(\lambda_i - \lambda_{d-1})^2}{C^2} \times \min\left(\frac{1}{\varepsilon\lambda_1^2} i \times d - \frac{d}{\lambda_1^2},\ i,\ \frac{d}{2}\right)$$

and hence that

$$\|M - \mathcal{A}(M)\|_F^2 \geqslant \Omega\left(\max_{1 \leqslant i \leqslant \frac{d}{2}} (\lambda_i - \lambda_{d-1})^2 \times \min\left(\frac{1}{\varepsilon\lambda_1^2} i \times d,\ i\right)\right). \tag{60}$$

Inequality 60 implies that

$$\|M - H\|_F^2 \geqslant \Omega\left(\frac{d}{\max(\lambda_1\sqrt{\varepsilon}, \sqrt{d})^2} \max_{1 \leqslant i \leqslant \frac{d}{2}} i \times (\lambda_i - \lambda_{d-i+1})^2\right) \tag{61}$$

with probability at least $\frac{1}{2}$.

Inequality 61 proves Theorem 2.4 when the output is in the unitary orbit $\mathcal{O}_\Lambda$. The bound for the setting when the output is a rank-k matrix is a special case of Corollary 2.5, and we defer the proof of this fact to the proof of Corollary 2.5.

■

Proof (of Corollary 2.5) Define $\Delta := \sqrt{\frac{cd}{\max(\lambda_1\sqrt{\varepsilon}, \sqrt{d})^2} \max_{1 \leqslant i \leqslant \frac{d}{2}} i \times (\lambda_i - \lambda_{d-i+1})^2}$. Let $\mathcal{A}$ be any ε-differentially private algorithm which takes as input a Hermitian matrix M and outputs a matrix $H = \mathcal{A}(M)$.

Consider the algorithm $\hat{\mathcal{A}}$ defined by

$$\hat{\mathcal{A}}(M) = \mathrm{argmin}_{Z \in \mathcal{O}_\Lambda} \|Z - \mathcal{A}(M)\|_F.$$

Since $\hat{\mathcal{A}}$ is just a post-processing of the output of the ε-differentially private algorithm $\mathcal{A}$, $\hat{\mathcal{A}}$ must also be ε-differentially private. Thus, by Theorem 2.4 there is a matrix $M \in \mathcal{O}_\Lambda$ such that the output $\hat{H} := \hat{\mathcal{A}}(M)$ of the projection of H onto $\mathcal{O}_\Lambda$ satisfies

$$\|M - \hat{H}\|_F^2 \geqslant \Delta^2, \tag{62}$$

with probability at least $\frac{1}{2}$. Let E be the event that equation 62 holds. Then $\mathbb{P}(E) \geqslant \frac{1}{2}$.

In the remainder of the proof, we suppose that the event E occurs. Then equation 62 holds and we have

$$\|M - \hat{H}\|_F \geqslant \Delta. \tag{63}$$

We consider the following two cases: $\|\hat{H} - H\|_F \geqslant \frac{\Delta}{2}$, and $\|\hat{H} - H\|_F < \frac{\Delta}{2}$. In the first case where $\|\hat{H} - H\|_F \geqslant \Delta$, we must have that, since $M \in \mathcal{O}_\Lambda$ and $\hat{H} = \mathrm{argmin}_{Z \in \mathcal{O}_\Lambda} \|Z - H\|_F$,

$$\begin{aligned} \|M - H\|_F &\geqslant \|\hat{H} - H\|_F \\ &\geqslant \frac{\Delta}{2}. \end{aligned} \tag{64}$$

Next, we consider the second case where $\|\hat{H} - H\|_F < \frac{\Delta}{2}$: By equation 63 we have that $\|M - \hat{H}\|_F \geqslant \Delta$. Thus, by the triangle inequality we have

$$\begin{aligned} \|M - H\|_F &\geqslant \|M - \hat{H}\|_F - \|\hat{H} - H\|_F \\ &\geqslant \Delta - \frac{\Delta}{2} \\ &\geqslant \frac{\Delta}{2}. \end{aligned} \tag{65}$$

Therefore, from equation 64 and equation 65, we have that

$$\|M - H\|_F \geqslant \frac{\Delta}{2}$$

and hence that

$$\begin{aligned} \|M - H\|_F^2 &\geqslant \frac{\Delta^2}{4} \\ &= \frac{cd}{4 \max(\lambda_1 \sqrt{\varepsilon}, \sqrt{d})^2} \max_{1 \leqslant i \leqslant \frac{d}{2}} i \times (\lambda_i - \lambda_{d-i+1})^2 \end{aligned}$$

whenever the event E occurs. Thus, since c is a universal constant, we can choose a slightly different universal constant c such that

$$\|M - H\|_F^2 \geqslant \frac{cd}{\max(\lambda_1 \sqrt{\varepsilon}, \sqrt{d})^2} \max_{1 \leqslant i \leqslant \frac{d}{2}} i \times (\lambda_i - \lambda_{d-i+1})^2, \tag{66}$$

with probability at least $\frac{1}{2}$ since $\mathbb{P}(E) \geqslant \frac{1}{2}$.

Note that, since the output of this algorithm $\mathcal{A}$ is allowed to be any matrix (either full rank or restricted to rank-k), the lower bound in equation 66 applies to both the setting when the output is rank-k for any $k \in [d]$ and when the output is full rank. Moreover, in the setting where the output has rank $k < d$, we have also have that $\|M - H\|_F^2 \geqslant \sum_{\ell=k+1}^{d} \lambda_\ell^2$ with probability 1. This fact, together with equation 66 imply that

$$\|M - H\|_F^2 \geqslant \Omega\left(\sum_{\ell=k+1}^{d} \lambda_\ell^2 + \frac{d}{4 \max(\lambda_1 \sqrt{\varepsilon}, \sqrt{d})^2} \max_{1 \leqslant i \leqslant \frac{d}{2}} i \times (\lambda_i - \lambda_{d-i+1})^2\right).$$

■

Proof (**of Corollary 2.6**) Theorem 2.4 guarantees that for any ε-differentially private algorithm $\mathcal{A}'$, with output eigenvalues $\gamma_1 \geqslant \cdots \geqslant \gamma_d \geqslant 0$, there exists an input matrix M with eigenvalues $\gamma_1 \geqslant \cdots \geqslant \gamma_d \geqslant 0$ such that the output $\mathcal{A}'(M)$ of this algorithm (with the same eigenvalues $\gamma_1 \geqslant \cdots \geqslant \gamma_d \geqslant 0$) satisfies

$$\|M - \mathcal{A}'(M)\|_F^2 \geqslant c^2,$$

where $c^2 := \Omega\left(\sum_{\ell=k+1}^d \gamma_\ell^2 + \frac{d}{\max(\gamma_1\sqrt{\varepsilon}, \sqrt{d})^2} \max_{1 \leqslant i \leqslant \frac{d}{2}} i \times (\gamma_i - \gamma_{d-i+1})^2\right)$.

Let $\Gamma = \mathrm{diag}(\gamma_1, \ldots, \gamma_d)$ and $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_d)$. Let $M = U\Gamma U^*$ be the spectral decomposition of M. And let $\tilde{M} = U\Lambda U^*$ be the matrix with eigenvalues λ_i, $i \in [d]$, and the same eigenvectors as M.

Recall from the statement of Corollary 2.6 that $H = \mathcal{A}(M)$ where $\mathcal{A}$ is an ε-differentially private algorithm $\mathcal{A}$ which takes as input a Hermitian matrix and outputs a rank-k Hermitian matrix with eigenvalues $\lambda_1, \ldots, \lambda_k$. Consider the following two cases:

$$(i) \quad \|M - \tilde{M}\|_F > \|\tilde{M} - H\|_F \text{ and } \quad (ii) \quad \|M - \tilde{M}\|_F \leqslant \|\tilde{M} - H\|_F.$$

On the one hand, if (i) holds, then we have

$$\|M - H\|_F \geqslant \|M - \tilde{M}\|_F > \|\tilde{M} - H\|_F \geqslant c, \tag{67}$$

where the first inequality holds since $\tilde{M} = \mathrm{argmin}_{Z \in \mathcal{O}_\Lambda} \|M - Z\|_F$.

On the other hand, if (ii) holds, we still have that

$$\|M - \tilde{M}\|_F \leqslant \|M - H\|_F$$

since $\tilde{M} = \mathrm{argmin}_{Z \in \mathcal{O}_\Lambda} \|M - Z\|_F$. Thus

$$\|\tilde{M} - H\|_F \leqslant \|M - H\|_F + \|M - \tilde{M}\|_F \leqslant 2\|M - H\|_F,$$

which implies that

$$\|M - H\|_F \geqslant \frac{1}{2}\|\tilde{M} - H\|_F \geqslant \frac{1}{2}c. \tag{68}$$

Thus, equation 67 and equation 68 together imply that

$$\|M - H\|_F^2 \geqslant \Omega\left(\sum_{\ell=k+1}^d \gamma_\ell^2 + \frac{d}{\max(\gamma_1\sqrt{\varepsilon}, \sqrt{d})^2} \max_{1 \leqslant i \leqslant \frac{d}{2}} i \times (\gamma_i - \gamma_{d-i+1})^2\right).$$

∎