

Differentially Private Linear Algebra in the Streaming Model

Jalaj Upadhyay
Center for Applied Cryptographic Research
David R. Cheriton School of Computer Science
University of Waterloo.
jalaj.upadhyay@uwaterloo.ca

Abstract

The focus of this paper is on differential privacy of streaming data using sketch-based algorithms. Previous works, like Dwork *et al.* (ICS 2010, STOC 2010), explored random sampling based streaming algorithms. We work in the well studied streaming model of computation, where the database is stored in the form of a matrix and a curator can access the database row-wise or column-wise.

Dwork *et al.* (STOC 2010) gave impossibility result for any non-trivial query on a streamed data with respect to the user level privacy. Therefore, in this paper, we restrict our attention to the event level privacy. We provide optimal, up to logarithmic factor, space differentially private mechanism in the streaming model for three basic linear algebraic tasks: matrix multiplication, linear regression, and low rank approximation, while incurring significantly less additive error.

Our approach for matrix multiplication and linear regression has some similarities with Blocki *et al.* (FOCS 2012) and Upadhyay (ASIACRYPT 2013) on the superficial level, but there are some subtle differences. For example, they perform an affine transformation to convert the private matrix in to a set of $\{\sqrt{w/n}, 1\}^n$ vectors for some appropriate w , while we perform an input perturbation that raises the singular value of the private matrix. In order to get a streaming algorithm for low rank approximation, we have to reuse the random Gaussian matrix in a specific way. We prove that the resulting distribution also preserve differential privacy.

We do not make any assumptions, like singular value separation, as made in the earlier works of Hardt and Roth (STOC 2013) and Kapralov and Talwar (SODA 2013). Further, we do not assume normalized row as in the work of Dwork *et al.* (STOC 2014). All our mechanisms, in the form presented, can also be computed in the distributed setting of Biemel, Nissim, and Omri (CRYPTO 2008).

Keywords. Differential Privacy, Linear Algebra, Random Projection.

1 Introduction

In the setting of large data analysis, one of the desired goals is to design a sub-linear space algorithm to perform computation while receiving the data online. There are many non-private algorithms in this setting. Such algorithms are called *online algorithms* or *streaming algorithms*. However, these data often contain sensitive information and privacy is as important as correct computation. Considering these two issues, a natural problem that a database curator faces is to generate a data structure that could be used to provide useful information without leaking sensitive information about an individual. The focus of this paper is space restricted streaming algorithms for answering linear algebraic queries in a private manner.

For an input matrix A , the standard techniques used in (non-private) streaming algorithms either does *random sampling* or compute a *sketch*, which has the form ΩA for some choice of random matrix Ω . Dwork

et al. [20] gave a private analogue of streaming algorithms that uses *sampling based approach* for various statistical queries, and gave an impossibility result for private analogues of *sketch based approaches* for specific “statistical queries.” This raises doubts over the applicability of sketch based approach in privacy. These doubts, if true, would be unfortunate because the sketch based approach is one of the major techniques (and often provide better utility guarantee) in the non-private setting. In this paper, *we show the first set of positive results for sketch based approach.* We deviate from the traditional mechanisms of differential privacy: instead of perturbing the output of the query by adding noise, we reversibly perturb the input and then multiply noise (analogous to Blocki *et al.* [6, 7] and Upadhyay [60]). We give almost optimal (in terms of space required) differentially private sketch based streaming algorithms for three basic algorithmic problems in linear algebra.

A natural question one might ask is, why should one care about private streaming algorithms for linear algebra? To answer this, let us consider the following scenario. A large database is streamed to the curator such that, unless we store the data, it is irretrievably gone. On the other hand, the curator has limited memory and it cannot store the whole database, but expects queries on the streamed data. In such a setting, the curator would like to store a data-structure with enough information about the database to (approximately) answer the queries. This idea has been used in the streaming model without any privacy concern. However, if a curator is handling a confidential data-base, it has to store a data-structure that, in addition to providing useful information to the query-maker, does not leak any information about the individual entry of the database as per the specific requirements of a robust privacy guarantee.

The scenario mentioned above is not an artificial problem. An $n \times d$ real-valued matrix is a natural structure for storing data about n entities, each of which is described by d features. For example, Hardt and Roth [34, 35] motivated the problem of differentially private low-rank approximation (LRA) by citing the Netflix competition. In addition to the Netflix type scenario, there are many other areas of large data analysis that have natural privacy concerns, like, genetics engineering finance. Computations in financial market often use various linear algebraic tasks as subroutine, like matrix multiplication or linear regression. Likewise, in genetic engineering, it is common to perform a procedure that computes a full or partial singular value decomposition (SVD) of the covariance matrix corresponding to the input data matrix, and then appeal to standard statistical model selection criterion, like getting the top half of the spectrum of the matrix, to quantify its significance.

On the other side, the traditional methods of *Krylov subspace iteration* (on which some of the recent works like [32, 35, 39] are based) and *rank revealing factorization method* requires a lot of space and are slow when matrices have high dimension. In fact, even storing the whole data during the computation is not always possible. Therefore, in such scenarios, computations are done in the streaming model. Many of these tasks have been studied in the non-private setting [3, 14, 25, 48, 56]. However, recent privacy violations in health-care and genetic studies have exemplified how sensitive these data are. This raises the question of whether one can perform all these tasks on a streamed data while giving a robust guarantee of privacy, like differential privacy. We do a principled study of private analogues of the known streaming algorithms for such tasks.

PRIVACY MODEL USED IN THIS PAPER. In this paper, we consider differential privacy in the streaming model where the data is streamed either row-wise or column-wise and the space available to the curator is sub-linear in the size of the dataset. There are two notions of differential privacy: *event level privacy*, where guarantees are at the granularity of individual records in the datasets, and *user level privacy*, where guarantees are at the granularity of each user whose data is present in the dataset. Dwork *et al.* [19] showed that it is impossible to obtain any non-trivial result with respect to the user level privacy when the data is

streamed online. Therefore, in this paper, we restrict our attention to the event level privacy.

1.1 Problem statements and our contributions.

In this section, we give the formal description of the problems we investigate in this paper. We also state the best known space lower bounds for non-private streaming algorithms and present our results. We note that stronger lower bounds are achievable for each of the problems considered because of the non-zero additive error; however, the bounds stated below are presently the best known space lower bounds that we can use to make any sort of optimality comparison.

We reserve the letter n for number of rows and d for number of columns of a private matrix. We assume $d < n$. The performance of a streaming algorithm is measured by three basic factors: the number of passes over the data stream, the memory used by the algorithm, and the total time taken by the algorithm. All our differentially private mechanisms for performing linear algebraic tasks are single-pass and achieve almost optimal space bound for one-pass algorithms. We do not make any assumption on the input matrix to compute low-rank approximation as in the previous works [22, 34, 35, 39]. For *bit complexity*, we use the convention used by Clarkson-Woodruff [14], i.e., the entries of an $n \times d$ matrix are $\kappa = \log(nd)$ -bit integers.

MATRIX MULTIPLICATION. The first problem we consider is matrix multiplication of two conforming matrices. It is one of the most important tools in numerical analysis (for example, every linear differential equations solver uses matrix product), and, therefore, wherever private data are analyzed numerically.

Problem 1. ((α, β, τ) -Matrix Multiplication). An $n \times d$ matrix A and $d \times n$ matrix B are given. Output a matrix C so that $\|AB - C\|_F \leq \alpha\|A\|_F \cdot \|B\|_F + \tau$ with probability at least $1 - \beta$.

Theorem 1.1. [14] Suppose $n \geq c\kappa/\alpha^2$ for an absolute constant $c > 0$. Then any randomized one-pass algorithm which solves matrix multiplication with probability at least $4/5$ uses $\Omega(d\alpha^{-2}\kappa)$ bits of space.

Clarkson-Woodruff [14, Theorem 2.4] extended the idea of Sarlos [56] to show the following in the non-private setting.

Theorem 1.2. Given $\alpha, \beta > 0$, and conforming matrices A and B . The matrices are presented row-wise and column-wise, respectively, with integer entries having κ bits. There is a data structure so that, at a given time, AB can be estimated, so that with probability at least $1 - \beta$, the Frobenius norm of the error is at most $\alpha\|A\|_F \cdot \|B\|_F$. There is an $r = O(1/\alpha^2)$ so that for d large enough, the data structure requires $O(rd \log(1/\beta) \log \kappa + \log(1/\alpha))$ bits of space and $O(r)$ update time.

Theorem 5.1 gives a data-structure that uses $O(d\alpha^{-2}\kappa \log(1/\beta) \log n)$ bits of space and provides (ϵ, δ) -differential privacy with $\tau = O(\alpha\sqrt{n})$ and $O(r)$ update time. If we agree to pay for update time, then we can remove the extra $\log n$ factor in the space used and achieve the space bound of Theorem 1.2.

LINEAR REGRESSION. The second problem we consider is *linear regression*, an important tool in statistical analysis. One of its major application is in finance, a domain where data are extremely sensitive. One example is *capital asset pricing model* which is used to predict demands [15], supplies [45], and investment [24].

Problem 2. ((α, β, τ) -Linear Regression). Given an $n \times d$ matrix A and a m set of $n \times 1$ column vectors $\{b_1, \dots, b_m\}$, output a set of vectors $X = \{x_1, \dots, x_m\}$ so that $\|AX - B\|_F \leq (1 + \alpha) \min_{Y \in \mathbb{R}^{d \times m}} \|AY - B\|_F + \tau$ with probability at least $1 - \beta$, where $B = \{b_1 | \dots | b_m\}$.

Theorem 1.3. [14] Suppose $n \geq \kappa d/36\alpha$ and d is sufficiently large. Then any randomized one-pass algorithm which solves the *Linear Regression* problem with probability at least $7/9$ needs $(d^2\alpha^{-1}\kappa)$ bits of space.

Clarkson-Woodruff [14, Theorem 3.2] showed the following in the non-private setting.

Theorem 1.4. Given $\alpha, \beta > 0$, an $n \times d$ matrix A , and an n -vector b . There is a data structure so that, at a given time, can be used to estimate the linear regression problem. There is $r = O(d/\alpha)$ so that for d large enough, the data structure requires $O(rd \log(1/\beta) \log \kappa + \log(1/\alpha))$ bits of space and update time $O(d/\alpha)$.

Theorem 5.3 gives a data-structure that uses $O(d^2 \alpha^{-1} \kappa \log(1/\beta) \log n)$ bits of space and provides (ϵ, δ) -differential privacy with $\tau = O(\alpha \sqrt{n})$ while maintaining an update time $O(d/\alpha)$. If we agree to pay for update time, then we can remove the extra $\log n$ factor in the space used and achieve the space bound of Theorem 1.4.

LOW-RANK APPROXIMATION. The last problem we consider is *low rank approximation* [16]. A partial list of its applications includes principal component analysis [36], fast multipole methods [27], and \mathcal{H} -matrices [26].

Problem 3. ((α, β, τ) -Low-rank approximation). Given an $n \times d$ matrix A , a target rank k , and an over-sampling parameter p , construct a matrix Ψ with $k + p$ orthonormal columns such that $\|A - \Psi \Psi^\top A\| \leq (1 + \alpha) \min_{\text{rank}(A_k) \leq k} \|A - A_k\| + \tau$ with probability at least $1 - \beta$ for both Frobenius and spectral norm.

Theorem 1.5. [14] Suppose $n \geq ck/\epsilon$ for an absolute constant $c > 0$, . Then any randomized 1-pass algorithm which solves *k-rank approximation* with probability at least $5/6$ uses $\Omega(nk\alpha^{-1})$ bits of space.

Clarkson and Woodruff [14] also gave a (non-private) one-pass algorithm for LRA. They used Rademacher matrices, i.e., matrices with entries ± 1 with probability $1/2$, to produce the sketch and an observation that the *projection* step can be emulated by the information gathered during the first stage if one uses Rademacher matrices. They showed the following [14, Theorem 4.5] in the non-private setting.

Theorem 1.6. Suppose input A is given as a sequence of columns or rows. There is an $r = O(k \log(1/\beta)/\alpha)$, such that with probability at least $1 - \beta$, a matrix \tilde{A}_k satisfying the condition of Problem 3 can be obtained under Frobenius norm. The space needed is $O(k\alpha^{-1}(n + d)\kappa \log(1/\beta))$

Unfortunately, we do not know a way to prove differential privacy using their algorithm, except to use additive noise mechanisms, much like Hardt and Roth [34]. In that case, when we try to emulate the second step, a simple analysis shows an error bound of order $k^{3/2}$. In this paper, we show how to emulate the projection step when using random Gaussian matrices. We give an $O(k\alpha^{-1}(n + d)\kappa)$ bits data structure (Theorem 5.5) that can be used to publish *k-rank approximation* of an $n \times d$ input matrix in a single pass with (ϵ, δ) -differential privacy while incurring an additive error $\tau \leq O(k\sqrt{n \ln(2/\delta)}/\epsilon)$ for the Frobenius and $\tau \leq O(\sqrt{nk \ln(2/\delta)}/\epsilon)$ for the spectral norm. Our multiplicative approximation factor is almost optimal because of Eckart-Young bound [23] in the case of Frobenius norm and Mirsky [46] in the case of spectral norm.

Note that our data-structure has an improved space requirement by a factor of $\log(1/\beta)$; thereby, just off by a factor of $1/\alpha$ from the optimal space bound. This is not the first time that methods used in private setting has helped to improve the results in non-private setting. Dwork *et al.* [22] has recently showed one such case. The reason behind this is simple. The bound on the single pass algorithm of Clarkson and Woodruff [14] and the two-pass algorithm of Sarlos [56] use the earlier result for matrix-multiplication, and, therefore, had to rely on the number of rows in the projection matrix used to prove the bound in matrix multiplication. On the other hand, as we discuss later, we use perturbation theory along the line of Halko *et al.* [31].

We next compare our results with the earlier known results. Let $\lambda_1, \dots, \lambda_{\text{rank}(A)}$ be the singular values of a matrix A . In Table 1, we compare our result stated for low-rank approximation with the previous works. For effective comparison, since we do not make any coherence assumption, one should put $\mu = n$ in Table 1. A detail comparison is done in Section 5.3.3 taking into account the difference in privacy model and assumptions made. In Table 2, we give the best known lower bound for the space required for each of

Method	Norm	Additive noise	Privacy Notion	Streaming	# Passes
Hardt-Roth [34]	F	$\frac{\sqrt{kn} \log(k/\delta)}{\varepsilon} + \sqrt{\frac{\mu \ A\ _F \log(k/\delta)}{\varepsilon}}$	Event level	No	2
Subspace Iteration [35]	S	$O\left(\frac{k^2}{\varepsilon} \sqrt{(\text{rk}(A)\mu + k \log n) \log\left(\frac{1}{\delta}\right) \log n}\right)$	Event level	No	$k\sqrt{\log \lambda}$
Kapralov-Talwar [39]	S	$O(dk^3/(\varepsilon\gamma^2\delta^2))$	User level	No	k
Hardt [32]	S	$\frac{\lambda_1 \sqrt{kn\mu \log(1/\delta) \log\left(\frac{n}{\gamma}\right) \log \log\left(\frac{n}{\gamma}\right)}}{\varepsilon\gamma^{1.5}\lambda_k}$	User level	No	$k\sqrt{\log \lambda}$
Dwork <i>et al.</i> [22]	S	$O((k\sqrt{n} \ln(1/\delta))/\varepsilon) + \tilde{O}(\sqrt{k^3 n^3/2}/\varepsilon^2)$	User level	Yes	1
This paper	F	$O(k\sqrt{n} \ln(2/\delta)/\varepsilon)$	Event level	Yes	1
This paper	S	$O(\sqrt{nk} \ln(2/\delta)/\varepsilon)$	Event level	Yes	1

Table 1: Comparison Between our Mechanism and Previous Mechanisms for k -Rank Approximation of an $n \times d$ matrix A . μ denotes the coherence of A , $\gamma = (\lambda_k/\lambda_{k+1}) - 1$, **S** stands for spectral norm and **F** for Frobenius norm.

	Space Bound			Additive Noise
	Lower Bound [14]	Non-private	Private	
Linear Regression	$\Omega(d^2\alpha^{-1}\kappa)$	$O(d^2\alpha^{-1}\kappa \log(1/\beta))$	$O(d^2\alpha^{-1}\kappa \log(1/\beta))$	$O(\sqrt{n}\alpha)$
Matrix Product	$\Omega(d\alpha^{-2}\kappa)$	$O(d\alpha^{-2}\kappa \log(1/\beta))$	$O(d\alpha^{-2}\kappa \log(1/\beta))$	$O(\sqrt{n}\alpha)$
Frobenius Low-rank	$\Omega(nk/\alpha)$	$O(k\alpha^{-2}(n+d)\kappa)$	$O(k\alpha^{-2}(n+d)\kappa)$	$O(k\sqrt{n} \ln(2/\delta)/\varepsilon)$
Spectral Low-rank	–	–	$O(k\alpha^{-2}(n+d)\kappa)$	$O(\sqrt{nk} \ln(2/\delta)/\varepsilon)$

Table 2: Optimality of our Results with Respect to the Best Known Space Bounds in Non-private Setting.

the problems stated above to compare the space optimality of our results. The various parameters used are as defined above.

Our mechanism for LRA can be easily compiled to give differentially private *principal component analysis* using standard algorithms that use LRA in the first step. Another important application of our mechanism for private sketch generation is in manifold learning. We do not formally state these mechanisms as there are standard algorithms for these applications that only use private matrix for one of the problems stated above and rest of the steps are deterministic function of these computations. One can also implement our mechanisms *as distributed algorithms*, a desirable feature as argued by [4]. This is because our mechanism uses operations that have efficient distributed algorithms. For example, one could use Jacobi method for SVD [41, Chapter 4], Cannon’s algorithm for multiplication [12], and GMRES for residual method [55].

OUR TECHNIQUES. The two standard techniques for streaming algorithms are (i) random sampling of the rows or columns of the streamed matrix and (ii) generating a random *sketch* of the matrix. In this paper, we use the sketch based approach. A sketch of a matrix A has the form ΩA for some appropriate choice of random matrix Ω . The known sketch based streaming algorithms for matrix product and linear regression, to our knowledge, use Rademacher matrix [14] or tug-of-war matrices [56]. Adding Gaussian noise to it to ensure differential privacy amounts to a large additive error.

In order to get a better utility bound, we use an idea analogous to Blocki *et al.* [6, 7] and Upadhyay [60]. We devise a private-sketch generation (PSG) mechanism to generate a private sketch of the private matrix. Our PSG uses random Gaussian matrix. We prove the privacy of PSG under certain spectral property of the input matrix. At a high level, all our mechanisms use this basic mechanism while maintaining the

spectral property of the input matrix (to guarantee privacy). The algorithm for both matrix product and linear regression involves deterministic use of the sketch generated to give an approximate matrix product and linear regression. This gives the same bound on the space used by the data-structure as achieved by Clarkson-Woodruff [14].

We still have not reached the required update time. The standard Gaussian matrices takes $O(nr)$ time for update. This has $O(n)$ gap with the bound we wish to achieve. For this, we use a projection matrix that mimics the action of standard Gaussian matrix. For this, we use the projection matrix proposed by Upadhyay [61], which is a generalization of the construction by Vybiral [62]. Vybiral [62] first pick n random Gaussian samples to form the first row and then construct the remaining $r - 1$ rows by shifting the vector left-wise relative to the previous row. This matrix is also known as *partial circulant matrices* and satisfies the *Restricted Isometry Property* [11], which we define next. For any set $T \subseteq \{1, \dots, r\}$, we say that an $r \times n$ matrix Φ satisfies the *Restricted Isometry Property* of order k if there exists an $0 < \varepsilon < 1$ such that, for all set T with $|T| < k$,

$$\Pr \left[(1 - \varepsilon) \|\mathbf{x}_T\|_2^2 \leq \|\Phi_T \mathbf{x}_T\|_2^2 \leq (1 + \varepsilon) \|\mathbf{x}_T\|_2^2 \right] \geq 1 - \eta \quad (1)$$

holds, where Φ_T (\mathbf{x}_T , respectively) is the restriction of Φ (\mathbf{x} , respectively) to the indices in T .

Rauhut *et al.* [51] proved that a partial circulant matrix formed as above satisfies the Restricted Isometry Property for values of $r \geq \max(\varepsilon^{-1} \sqrt{(k \log n)^3}, \varepsilon^{-2} k \log^4 n)$. Vybiral then multiply this matrix P to a diagonal matrix whose entries are ± 1 with probability $1/2$. By Theorem 1.7, this construction satisfies the Johnson-Lindenstrauss bound for suboptimal value of $r = O(\varepsilon^{-2} \log^2 m)$, where m is the number of vectors on which the transform is to be applied.

Theorem 1.7. (*Krahmer-Ward [42, Proposition 3.2]*) Let ε be an arbitrary constant. Let Φ be a matrix of order k and dimension $r \times n$ that satisfies the relation $k \leq c_1 \delta_k^2 r / \log(n/r)$ and equation (1). Then the matrix ΦD , where D is an $n \times n$ diagonal matrix whose entries are ± 1 with probability $1/2$ (also known as *Rademacher matrix*), is a Johnson-Lindenstrauss transform with r rows.

Therefore, unless there is a significant improvement in the understanding of concentration properties of partial circulant matrices, the dimension bound achieved by Vybiral [62] is hard to beat. The key observation here is that the diagonal Rademacher matrices does not produce a proper “mixing” of the entries of the partial circulant matrices to facilitate a strong concentration bound. For this, we need to compose it with a matrix that allows fast matrix-vector multiplication, preserves the Euclidean norm of the input vectors, and does not introduce more randomness. Our key observation is that, instead of only preconditioning by a diagonal matrix formed by a Rademacher sequence, if we also compose a Walsh-Hadamard matrix, then we achieve good enough mixing that translates to a better concentration result¹. This in turn helps us to strengthen the Vybiral’s bound [62].

To generalize this construction, a key point to note is that partial circulant matrices are nothing special. They are simply the first r rows of a fully circulant matrix, and, therefore, can be seen as a result of applying a truncated permutation matrix from the left to a fully circulant matrix. Therefore, combined with the symmetry of circulant matrices, intuitively, any r rows restricted circulant matrix should not effect the final concentration result. This intuition infact turns out to be true. This increases our sampling complexity by an additive factor of r because we need to sample r rows of a circulant matrix, but we now have a family of random projection matrices that satisfies the Johnson-Lindenstrauss lemma. Sampling r rows independently was an idea used by Rudelson and Vershynin [54], where they showed that a matrix formed by sampling r rows of a deterministic matrices with bounded orthonormal rows satisfies the Restricted Isometry Property.

¹This composed matrix is the isometry matrix of Ailon and Chazelle [1].

Here, we are sampling r rows of a random matrix. Therefore, at the cost of oversimplification, an intuitive way to see this general class of Johnson-Lindenstrauss transform is as a hybrid of the known constructions of projection matrix with Restricted Isometry Property and known constructions of the Johnson-Lindenstrauss transform.

At a high level, our mechanisms for matrix product and linear regression are private analogues of Ref. [14] with the utility proof using the analogous result guaranteed by our projection matrix instead of the variance bounds for 4-wise independent *tug-of-war* matrices used by earlier works [14, 56] (note that differential privacy would not necessarily hold if the random matrix has dependent entries in its row). One approach could be to prove concentration bound on the corresponding distribution of product matrices. We use the approach used by Sarlos [56] and Clarkson-Woodruff [14] combined with our bound for projection matrices.

The mechanism for LRA is more complicated and markedly different from the recent works [22, 32, 34, 35, 39] (the online version of Dwork *et al.* [22] uses binary tree technique [19] and assumes a lower bound condition on the optimal value, see [22, Theorem 8]). All the previous works perturb the output by adding noise to it, while we perturb the input matrix and then multiply noise matrix. We follow the general prototype of algorithms to compute a LRA, i.e., first computes a projection matrix (*range finding* step) and then computes a low rank matrix by operating the projection matrix (*projection* step) on the input matrix. In the most naive form, both the steps require the input matrix. However, Ref. [14] showed that one can emulate the projection step by using the matrices formed in the range finding step, eliminating the need of input matrix in the projection step. Unfortunately, it only works for Rademacher matrices. On the other hand, our projection matrices are generated using a random Gaussian matrix.

The first key observation is that, by a clever use of linear algebra, information gathered in the range-finding step can be used to emulate the projection step without using the input matrix explicitly. However, we need to reuse the random Gaussian matrix. Therefore, the privacy is not as straightforward as for the other two problems. Fortunately, the Gaussian matrix is reused in a specific manner for which one can prove privacy under certain spectral property of the input matrix. We believe that this could be of independent interest. The second observation, also done by Blocki *et al.* [6], is that the mechanism for PSG already gives considerable improvement in the range finding step. On top of that, we use an oversampling parameter p . This extra oversampling parameter helps us in getting much sharper bounds for both spectral as well as Frobenius norm.

We now mention the difference between our analysis and the analyses of Clarkson-Woodruff [14] and Sarlos [56]. The two results that used Johnson-Lindenstrauss matrices to give low rank approximation uses the trick that a good bound on matrix multiplication allows them to give a good bound on low-rank approximation. This only give them a bound when the approximation metric is stated in terms of Frobenius norm. Moreover, the algorithm of Sarlos [56] requires two-pass over the private input matrix. Our utility proof uses perturbation theory along the line of [31]. This allows us to give bounds for both the norms in an unified manner. We note that all the previous proofs were tailored for specific norm.

RELATED WORKS. The first formal definition of Differential Privacy was given by Dwork *et al.* [18]. They used Laplacian distribution to guarantee differential privacy for bounded *sensitivity* query functions. The Gaussian variant of this basic sanitizer was proven to preserve differential privacy by Dwork *et al.* [17] in a follow-up work. Since then, many mechanisms for preserving differential privacy have been proposed in the literature [9, 21, 28, 29, 40, 33, 34, 44, 53]. All these sanitizers have a common theme: they perturb the output before responding to queries. Blocki *et al.* [6, 7] and Upadhyay [60] took a complementary approach. They perturb the input reversibly and then perform a random projection of the perturbed matrix.

There are some recent works on differentially private low-rank approximation and differentially private streaming algorithm for statistical queries. Blum *et al.* [8] first studied this problem and gave a simple “input perturbation” algorithm that adds noise to the covariance matrix, an approach also taken recently by Dwork *et al.* [22]. This was improved by Hardt and Roth [34] who studied the low rank approximation in Frobenius norm under the low coherence assumption. Kapralov and Talwar [39] and Chaudhary *et al.* [13] studied the spectral low rank approximation of a matrix by giving a matching upper and lower bounds for privately computing the top k eigenvector of a matrix. Recently, Hardt and Roth [35] improved their noise bound by proposing robust private subspace iteration mechanism. All these works [13, 35, 39] uses some eigenvalue separation assumption, i.e., the top eigenvalue and the k -th eigenvalue has some separation, and Hardt and Roth [35] used the low-coherence assumption. Recently, Dwork *et al.* [22] revisited randomized mechanism to give a tighter bound. They also gave an online version of their mechanism under a normalized row assumption.

The literature of performing (non-private) linear algebra using streaming algorithms, started by Alon *et al.* [3], is so extensive that we cannot hope to cover it in any detail here. In the private setting, Dwork *et al.* [20] studied *pan-privacy*, where the internal state is known to the adversary, to answer various counting tasks, like estimating distinct elements, cropped means, number of heavy hitters, and frequency counts. All these mechanisms uses private version of various sampling based streaming algorithms. Subsequently, there have been some works on online differential privacy [22, 37] for various tasks.

ORGANIZATION OF THE PAPER. In Section 2, we cover the basic definitions and notations. In Section 3, we give our basic mechanism that we use for intuition of the utility proof. We give the mechanism which achieves a better update time in Section 4 and the mechanisms for solving various linear algebra task in Section 5. We conclude the paper by stating some open problems in Section 6.

2 Notations and Basic Preliminaries

NOTATIONS. We reserve the letters A and B for private input matrices, Ω for Gaussian matrix (i.e., matrix whose entries are picked independently from a Gaussian distribution $\mathcal{N}(0, 1)$). For an $n \times d$ matrix A , we let $A_{i\cdot}$ to denote the i -th row of A , $A_{\cdot j}$ to denote the j -th column of A , and A' to denote the symmetric matrix $\begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}$ corresponding to A . We let A_t denote the matrix received after t time epochs. When we wish to refer to both the Frobenius as well as the spectral norm, we overload the symbol $\|\cdot\|$ and drop the subscript. We let e_1, \dots, e_d denote the standard basis vectors in \mathbb{R}^d . We use the Dirac notation to denote vectors, i.e., $|\cdot\rangle$ for column-vector and $\langle \cdot|$ for row-vector. We use $\langle 0^n|$ to denote the transpose of an n -dimensional 0-vector. For a matrix M , we write $M \succ 0$ if all its eigenvalues are positive.

PRIVACY. We work with the relaxed notion of privacy, known as *approximate differential privacy*. We consider two data-sets D_1 and D_2 *neighbouring* if $\|D_1 - D_2\| \leq 1$. This notion was used in many of the earlier works [34, 35]

Definition 2.1. A randomized mechanism, \mathcal{K} , gives (ϵ, δ) -*differential privacy*, if for all neighbouring data-sets D_1 and D_2 , and all range $S \subset \text{Range}(\mathcal{K})$, $\Pr[\mathcal{K}(D_1) \in S] \leq \exp(\epsilon)\Pr[\mathcal{K}(D_2) \in S] + \delta$, where the probability is over the coin tosses of \mathcal{K} . When $\delta = 0$, we get the traditional definition of *differential privacy*.

We use the following in our analysis explicitly or implicitly.

Theorem 2.1. (*Composition Theorem [21]*). Let $\varepsilon, \delta \in (0, 1)$, and $\delta' > 0$. If $\mathcal{K}_1, \dots, \mathcal{K}_\ell$ are each (ε, δ) -differential private mechanism, then the mechanism $\mathcal{K}(D) := (\mathcal{K}_1(D), \dots, \mathcal{K}_\ell(D))$ releasing the concatenation of each algorithm is $(\varepsilon', \ell\delta + \delta')$ -differentially private for $\varepsilon' < \sqrt{2\ell \ln(1/\delta')} \varepsilon + 2\ell\varepsilon^2$.

Lemma 2.2. Let $M(D)$ be a (ε, δ) -differential private mechanism for a database D , and let h be any function, then any mechanism $M' := h(M(D))$ is also (ε, δ) -differentially private for the same set of queries.

In the course of this paper, we use few standard results from the theory of linear algebra, random matrices, and perturbation theory. We follow up with the key concepts and results we would need in our proofs.

LINEAR ALGEBRA. Our analysis makes extensive use of linear algebra and statistical properties of the Gaussian distribution. We give an exposition to the level required to understand this paper. Let A be an $n \times d$ matrix. The singular value decomposition (SVD) of A is $A = V\Lambda U^\top$, where U and V are left and right eigenvectors of A , and Λ is a diagonal matrix. The entries of Λ are called the *singular values* of A . Since U and V are unitary matrices, one can write $A^i = V\Lambda^i U^\top$ for any real value i . We let $rk(A)$ denote the rank of the matrix and $\lambda_i(A)$ its singular values. Where it is clear from context, we simply write λ_i for the singular values.

We use various matrix norms. We use the notation $\|\cdot\|_F$ for Frobenius norm. the Frobenius norm for a matrix $A = (a_{ij})_{i \in [n], j \in [d]}$ is defined as following $\|A\|_F = \sqrt{\sum_{ij} |a_{ij}|^2}$. For a matrix A , we let $\|A\|_2$ denote the 2-norm, i.e., $\max_{x \in \mathbb{R}^d} \|Ax\|_2 / \|x\|_2$. We use the symbol $\|\cdot\|$ when we wish to refer to both Frobenius as well as 2-norm. We explicitly or implicitly use the fact that matrix norms are Lipschitz. We let e_1, \dots, e_n denote the standard basis vectors in \mathbb{R}^n . We denote by $a|b$ the vector formed by appending the vectors a and b . A matrix M is *positive semi-definite* if all its eigenvalues are non-negative, i.e., if for all $\mathbf{x} \in \mathbb{R}^n$, we have $\mathbf{x}^\top M \mathbf{x} \geq 0$. For two $n \times n$ matrices M and N , we denote by $M \succeq N$ if $M - N$ is a positive semi-definite matrix. We write $M \succ 0$ if all its eigenvalues are positive. We note few of the key lemmata of linear algebra used in this paper.

Lemma 2.3. If matrix A and B are conforming, then $\|AB\|_F \leq \|A\|_2 \|B\|_F$.

Lemma 2.4. Let $A = (a_{ij})$ be a symmetric $n \times n$ matrix. If $|a_{ii}| \geq |a_{ij}|$ for all $1 \leq i, j \leq n$, then A is a positive semi-definite matrix.

Lemma 2.5. A matrix A has $\|A\|_2 \leq 1$ if and only if $\begin{pmatrix} \mathbb{I} & A \\ A^\top & \mathbb{I} \end{pmatrix}$ is positive semi-definite.

Lemma 2.6. Let $A = (a_{ij})$ be a symmetric $n \times n$ matrix. If $|a_{ii}| \geq |a_{ij}|$ for all $1 \leq i, j \leq n$, then A is a positive semi-definite matrix.

Lemma 2.7. Let A and B be Hermitian matrices with only 0, 1 entries. Then $\text{Tr}(AB) \leq \text{Tr}(A)\text{Tr}(B)$. Moreover, if $\|A - B\| \leq 1$, then $\text{Tr}(A^\top A) - \text{Tr}(B^\top B) \leq 2$.

Lemma 2.8. [30] For matrices A, B, C, D, E , and X of appropriate dimensions, we have

1. $\text{vec}(AXB) = (B^\top \otimes A)\text{vec}(X)$,
2. $\text{Tr}(CXB) = (\text{vec}(C^\top))^\top (\mathbb{I}_q \otimes X)\text{vec}(B)$,
3. $\text{Tr}(DX^\top EXB) = (\text{vec}(X))^\top (D^\top B^\top \otimes E)\text{vec}(X) = (\text{vec}(X))^\top (BD \otimes E^\top)\text{vec}(X)$.

GAUSSIAN DISTRIBUTION. A random variable, X , distributed according to a Gaussian distribution has the probability density function, $\text{PDF}_X(x) = \frac{1}{\sqrt{2\pi\sigma}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$. We denote it by $X \sim \mathcal{N}(\mu, \sigma^2)$. The

Gaussian distribution is invariant under affine transformation, i.e., if $X \sim \mathcal{N}(\mu_x, \sigma_x)$ and $Y \sim \mathcal{N}(\mu_y, \sigma_y)$, then $Z = aX + bY$ has the distribution $Z \sim \mathcal{N}(a\mu_x + b\mu_y, a\sigma_x^2 + b\sigma_y^2)$.

The multivariate Gaussian distribution is a generalization of univariate Gaussian distribution. Given a m dimensional multivariate random variable, $X \sim \mathcal{N}(\mu, \Sigma)$, the PDF of a multivariate Gaussian is given by $\text{PDF}_{\mathbf{X}}(\mathbf{x}) := \frac{1}{\sqrt{(2\pi)^{\text{rank}(\Sigma)} \Delta(\Sigma)}} \exp\left(-\frac{1}{2}\mathbf{x}^T \Sigma^{-1} \mathbf{x}\right)$ with mean $\mu \in \mathbb{R}^m$ and covariance matrix $\Sigma = \mathbb{E}[(X - \mu)(X - \mu)^T]$. It is easy to see from the description of the PDF that, in order to define the PDF corresponding to a multivariate Gaussian distribution, Σ has to have full rank. If Σ has a non-trivial kernel space, then the PDF is undefined. However, in this paper, we only need to compare the probability distribution of two random variables which are defined over the same subspace. Therefore, wherever required, we restrict our attention to the (sub)space orthogonal to the kernel space of Σ .

Multivariate Gaussian distribution maintains many key properties of univariate Gaussian distribution. For example, any (non-empty) subset of multivariate Gaussians is a multivariate Gaussian and linear functions of multivariate Gaussian random variables are multivariate Gaussian random variables, i.e., if $\mathbf{y} = A\mathbf{x} + \mathbf{b}$, where $A \in \mathbb{R}^{n \times m}$ is a non-singular matrix and $\mathbf{b} \in \mathbb{R}^n$, then $\mathbf{y} \sim \mathcal{N}(A\mu + \mathbf{b}, A\Sigma A^T)$.

RANDOM MATRIX THEORY. Our analysis for utility guarantee relies heavily on the theory of random matrices. We enumerate few of the key lemmata that we use in our analysis of low-rank approximation. For more details, we refer the readers to excellent books on random matrix theory [59] and multivariate statistics [47].

Lemma 2.9. Let Ω be a Gaussian matrix. Then for any fixed matrices A and B , $\mathbb{E}[\|A\Omega B\|_2] \leq \|A\|_2 \|B\|_F + \|A\|_F \|B\|_2$.

Lemma 2.10. Let Ω be a $n \times (k + p)$ Gaussian matrix. Then

$$\mathbb{E}[\|\Omega^{-1}\|_F^2] = \sqrt{\frac{k}{p-1}} \quad \text{and} \quad \mathbb{E}[\|\Omega^{-1}\|_2] \leq \frac{e\sqrt{k+p}}{p}.$$

Lemma 2.11. Let Ω be a random $n \times (k + p)$ Gaussian matrix whose entries are picked from the distribution $\mathcal{N}(0, 1)$. Then $\mathbb{E}[\text{Tr}((\Omega^T \Omega)^{-1})] = k/(p - 1)$.

Theorem 2.12. (*Johnson-Lindenstrauss lemma*) Fix any $\alpha < 1/2$ and let m be a positive integer. Let Ω be a $k \times n$ matrix whose entries are picked from a Gaussian distribution $\mathcal{N}(0, 1)$, where $k \geq 4(\alpha^2/2 - \alpha^3/3)^{-1} \ln m$. Then for any m unit vector set S in \mathbb{R}^n

$$\forall x \in S, \Pr_M [\|\Omega x\|_2 \in (1 \pm \alpha)\|x\|_2] \geq 2/3.$$

3 Space Efficient Differentially Private Sketch Generation

We study differential privacy in the well known streaming model of computation [3]. We present it at the level required to understand this paper. A more formal definition appears in Alon *et al.* [3]. It has three entities: a stream generator (database owner) \mathcal{S} , a (database) curator \mathcal{K} , and a query maker \mathcal{Q} . \mathcal{S} starts the process at time $t = 0$. The curator initializes its data structure to \mathcal{D}_0 . From $t = 0$, the curator is allowed only one-pass over the input matrix, i.e., it can copy any entry of the data-base to its memory space during exactly one time epoch. It updates its data structure to \mathcal{D}_t using \mathcal{D}_{t-1} and newly accessed data-points of the matrix. At certain time, $t = t_q$, the query maker \mathcal{Q} generates a query function q . The curator responds with the response, $q(\mathcal{D}_{t_q})$.

On input a streamed vector v , **flag**, parameters ε, δ , and an $r \times O(n)$ random Gaussian matrix Ω , the mechanisms does the following:

Variante 1: If **flag** = 0, compute $Y_v = \langle v | \Omega^\top$; else compute $Y_v = \Omega | v \rangle$. In the end, return Y_v .

Variante 2: If **flag** = 0, compute $Y_v = \langle v | \Omega^\top \Omega$; else compute $Y_v = \Omega^\top \Omega | v \rangle$. In the end, return Y_v .

Figure 1: Private Sketch Generation (PSG) Algorithm

The streaming model has a resource bound on the curator. A curator is only allowed to use time *polynomial in the size of the data base* to construct the data structure, and memory *sub-linear in the size of the data base*. For *differential privacy*, we further require that the response of \mathcal{K} to the query of \mathcal{Q} should satisfy Definition 2.1 with the two neighbouring streams differing in at most one entry.

In all our mechanisms, the curator uses sketch of the private matrix as the data-structure. The sketch of the matrix is generated row (or column) wise using one of the variants presented in Figure 1.

The first variant has some resemblance to the mechanism of Blocki *et al.* [6] and Upadhyay [60] in the sense that we multiply an appropriate Gaussian matrix, while the second variant can be seen as its extension in the sense that two successive application of Gaussian matrix in a defined form also preserve privacy. However, the analogy ends here. For example, [6, 60] perform an affine transformation to convert the private matrix in to a set of $\{\sqrt{w/n}, 1\}^n$ vectors, while we perform the perturbation to raise the singular value before invoking PSG (see Section 5). As argued by Blocki *et al.* [6], their mechanism does not give a guarantee that singular values of $A^\top A$ and their published matrix is close or their eigenvalues are comparable. In other words, it does not give a LRA. Apart from these major differences, there are couple of subtle differences. First of all, their mechanism is not a single pass in the way it is presented (they require at least two-pass over the input matrix even with the streaming algorithms for computing the SVD [52, 58]: first to subtract the mean of the entries of the matrix and second to compute the projection on the altered SVD). Our first observation is that we do not need to subtract the entries of the matrix because of the type of queries we are dealing with. Secondly, they project the entries of the columns of the private matrix to a higher dimensional space; here, we perform embedding to a lower dimensional subspace in the similar vein as other applications of JL-transform.

Theorem 3.1. If the singular values of the streamed matrix to the first variant of the PSG algorithm are at least $\sigma_{\min} := \frac{4\sqrt{r \log(2/\delta) \log(r/\delta)}}{\varepsilon}$ and for second variant are at least $\sigma_{\min} := \frac{4r \log(r/\delta)}{\varepsilon}$. Then PSG using $r \times O(n)$ Gaussian matrix preserves (ε, δ) -differential privacy.

The proof of the first variant follows the idea of Blocki *et al.* [6] taking into account the subtle differences mentioned above; however, the proof of variants 2 is more involved. We show that, for a streamed matrix A , the probability density function of the published matrix when using the second variant is

$$\frac{\exp(-\text{Tr}((A^\top A)^{-1} \Phi)/2) \Delta(\Phi)^{(n-r-1)/2}}{2^{rn/2} \pi^{r(r-1)/4} \Delta(A^\top A)^{r/2} \prod_{i=1}^r \Gamma((n-i+1)/2)}, \quad (2)$$

where $\Delta(A)$ is the product of the singular values of A and $\Phi = \sum_{i=1}^r |a_i\rangle\langle a_i|$ for n -variate Gaussians, a_1, \dots, a_r . This is the technical part of the proof; rest of the proof follows by evaluating the pdf for neighbouring matrices. We give a detail proof for variant 1 in Appendix B.

Proof. We now prove that the second variant preserves privacy if the singular values of the streamed matrix follows the hypothesis of the theorem. We start by computing the probability density function when the underlying multivariate Gaussian distribution is $\mathcal{N}(0, \mathbb{I})$. The case for arbitrary positive definite covariance

matrix follows just like the transition from identity to arbitrary positive definite covariance matrices in the multivariate Gaussian distribution. Let $\alpha_1, \dots, \alpha_r$ be i.i.d. $\mathcal{N}(0, \mathbb{I})$ be r multivariate Gaussian distribution, i.e., $\alpha_{ij} \sim \mathcal{N}(0, 1)$ for $1 \leq i \leq r, 1 \leq j \leq n$. The distribution we are interested in is $\Phi = \sum_{i=1}^r |\alpha_i\rangle\langle\alpha_i|$. We use the notation PDF($\Phi; \mathbb{I}$) to denote the probability density function of Φ when each random variable is picked using a normal distribution, i.e., when covariance matrix of the random variables is \mathbb{I} .

Using the chain rule, the joint distribution of the entries of Φ is as follows.

$$\begin{aligned} \text{PDF}(\Phi; \mathbb{I}) &= \text{PDF}(\langle\alpha_1, \alpha_1\rangle; \mathbb{I}) \text{PDF}(\langle\alpha_2, \alpha_1\rangle, \langle\alpha_2, \alpha_2\rangle | \langle\alpha_1, \alpha_1\rangle; \mathbb{I}) \cdots \\ &\quad \text{PDF}(\langle\alpha_r, \alpha_1\rangle, \dots, \langle\alpha_r, \alpha_r\rangle | \Phi_{[r-1]}; \mathbb{I}). \end{aligned} \quad (3)$$

There are $r(r+1)/2$ distinct entries, $\langle\alpha_1, \alpha_1\rangle, (\langle\alpha_2, \alpha_1\rangle, \langle\alpha_2, \alpha_2\rangle), \dots, (\langle\alpha_r, \alpha_1\rangle, \dots, \langle\alpha_r, \alpha_r\rangle)$. Our aim is to compute each individual term in the product form of the above chain rule. For this, we use the facts we first analyze the distribution of $(\langle\alpha_i, \alpha_1\rangle, \dots, \langle\alpha_i, \alpha_{i-1}\rangle)$. Let $\beta_{i-1}^\top = (\langle\alpha_i, \alpha_1\rangle, \dots, \langle\alpha_i, \alpha_{i-1}\rangle)$. Then we use the fact that there is a transformation of Jacobian one from $(\langle\alpha_i, \alpha_1\rangle, \dots, \langle\alpha_i, \alpha_{i-1}\rangle, \langle\alpha_i, \alpha_i\rangle - \beta_{i-1}^\top \Phi_{[i-1]}^{-1} \beta_{i-1})$ to $(\langle\alpha_i, \alpha_1\rangle, \dots, \langle\alpha_i, \alpha_{i-1}\rangle)$ to compute each term in the chain rule [47, 50].

Distribution of β_{i-1} . We first prove that β_{i-1} is $(i-1)$ variate Gaussian distribution. Since $\Sigma = \mathbb{I}$, and $\alpha_{11}, \dots, \alpha_{1n}, \dots, \alpha_{r1}, \dots, \alpha_{rn}$ are i.i.d. $\mathcal{N}(0, 1)$, from the elementary property of linear functions of normal variables, conditional on α_{kj} for $1 \leq k \leq i-1$ and $1 \leq j \leq n$, β_{i-1} is $(i-1)$ variate Gaussian distribution with

$$\Phi_{[i]} = \begin{pmatrix} \langle\alpha_1, \alpha_1\rangle & \cdots & \langle\alpha_1, \alpha_i\rangle \\ \vdots & \ddots & \vdots \\ \langle\alpha_i, \alpha_1\rangle & \cdots & \langle\alpha_i, \alpha_i\rangle \end{pmatrix}$$

Now $\alpha_{11}, \dots, \alpha_{rn}$, for every $j = 1, \dots, n$ are mutually independent; therefore, we have

$$\text{COV}(\beta_{i-1}, \alpha_{ij}) = [\text{COV}(\langle\alpha_i, \alpha_1\rangle \alpha_{ij}), \dots, \text{COV}(\langle\alpha_i, \alpha_{i-1}\rangle \alpha_{ij})] = (\alpha_{1j}, \dots, \alpha_{i-1,j})^\top$$

and $\mathbb{E}[|\beta_{i-1}\rangle\langle\beta_{i-1}| | \alpha_{kj}] = \Phi_{[i-1]}$ for $1 \leq j < i$. This implies

$$\text{COV} \left[\beta_{i-1}, \alpha_{ij} - \beta_{i-1}^\top \Phi_{[i-1]}^{-1} (\alpha_{1j}, \dots, \alpha_{i-1,j})^\top | \alpha_{kj} \right] = 0 \quad \forall 1 \leq k \leq i-1, \quad (4)$$

as the left hand side equals $(\alpha_{1j}, \dots, \alpha_{i-1,j})^\top - \Phi_{[i-1]} \Phi_{[i-1]}^{-1} (\alpha_{1j}, \dots, \alpha_{i-1,j})^\top$.

Therefore, β_{i-1} is independent of $\sum_{j=1}^k \left(\alpha_{ij} - (\alpha_{1j}, \dots, \alpha_{i-1,j})^\top \Phi_{[i-1]}^{-1} (\alpha_{1j}, \dots, \alpha_{i-1,j}) \right)^2$. Rao [50] proved that

$$\sum_{j=1}^k \left(\alpha_{ij} - (\alpha_{1j}, \dots, \alpha_{i-1,j})^\top \Phi_{[i-1]}^{-1} (\alpha_{1j}, \dots, \alpha_{i-1,j}) \right)^2 \sim \chi_{n-i+1}^2, \quad (5)$$

the standard χ^2 distribution.

Computing every term in the chain rule. From the fact that β_{i-1} is a $(i-1)$ -variate Gaussian distribution, equation (4), equation (5), and the identity

$$\Delta(\Phi_{[i]}) = \Delta(\Phi_{[i-1]}) \sum_{j=1}^k \left(\alpha_{ij} - (\alpha_{1j}, \dots, \alpha_{i-1,j})^\top \Phi_{[i-1]}^{-1} (\alpha_{1j}, \dots, \alpha_{i-1,j}) \right)^2,$$

we first calculate the joint pdf of

$$\begin{aligned} & \left(\langle \alpha_i, \alpha_1 \rangle, \dots, \langle \alpha_i, \alpha_{i-1} \rangle, \langle \alpha_i, \alpha_i \rangle - \langle \beta_{i-1} | \Phi_{[i-1]}^{-1} | \beta_{i-1} \rangle \right)^\top \\ &= \frac{\exp \left(-\frac{1}{2} \left(\langle \alpha_i, \alpha_1 \rangle, \dots, \langle \alpha_i, \alpha_{i-1} \rangle \right)^\top \Phi_{[i-1]}^{-1} \left(\langle \alpha_i, \alpha_1 \rangle, \dots, \langle \alpha_i, \alpha_{i-1} \rangle \right) \right)}{(2\pi)^{(i-1)/2} \Delta(\Phi_{[i-1]})^{1/2}} \\ & \quad \times \frac{\exp \left(-\frac{\langle \alpha_i, \alpha_i \rangle - \beta_{i-1}^\top \Phi_{[i-1]}^{-1} \beta_{i-1}}{2} \right) \left(\langle \alpha_i, \alpha_i \rangle - \beta_{i-1}^\top \Phi_{[i-1]}^{-1} \beta_{i-1} \right)^{(n-i+1)/2-1}}{2^{(n-i+1)/2} \Gamma((n-i+1)/2)} \\ &= \frac{\exp \left(-\frac{\beta_{i-1}^\top \Phi_{[i-1]}^{-1} \beta_{i-1}}{2} \right) \Delta(\Phi_i)^{(n-i-1)/2}}{2^{rn/2} \pi^{(i-1)/2} \Gamma((n-i+1)/2) \Delta(\Phi_{[i-1]})^{(n-i)/2}} \\ &= \text{PDF}(\langle \alpha_i, \alpha_1 \rangle, \dots, \langle \alpha_i, \alpha_{i-1} \rangle, \langle \alpha_i, \alpha_i \rangle | \Phi_{[i-1]}), \end{aligned} \tag{6}$$

where the last step used the fact that there is a one-to-one transformation of Jacobian 1 to $(\langle \alpha_i, \alpha_1 \rangle, \dots, \langle \alpha_i, \alpha_{i-1} \rangle)$ from $(\langle \alpha_i, \alpha_1 \rangle, \dots, \langle \alpha_i, \alpha_{i-1} \rangle, \langle \alpha_i, \alpha_i \rangle - \beta_{i-1}^\top \Phi_{[i-1]}^{-1} \beta_{i-1})$.

Computing the joint distribution of Φ . A simple arithmetic followed by plugging equation (6) in equation (3) gives the closed formed expression of the pdf of Φ as

$$\frac{\exp(-\text{Tr}(\Phi)/2) \Delta(\Phi)^{(n-r-1)/2}}{2^{rn/2} \pi^{\sum_i (i-1)/2} \prod_{i=1}^r \Gamma((n-i+1)/2)} \times \prod_{i=1}^r \left(\frac{\Delta(\Phi_{[i]})^{(n-i-1)/2}}{\Delta(\Phi_{[i-1]})^{(n-i)/2}} \right) = \frac{\exp(-\text{Tr}(\Phi)/2) \Delta(\Phi)^{(n-r-1)/2}}{2^{rn/2} \pi^{n(n-1)/4} \prod_{i=1}^r \Gamma((n-i+1)/2)}.$$

Using standard techniques (that could be found in any standard textbook, including Rao [50]) of the transformation method and the factorization theorem yields the PDF for arbitrary linear translation. That is, if $X \sim \text{PDF}(\Phi; \mathbb{I})$, then $Y = AX$ is distributed as $\text{PDF}(\Phi; A^\top A)$. The proof is similar to the similar transformation for multivariate Gaussian distribution. It is easy to verify that it does not matter if we multiply A from right or left of vectors $\alpha_1, \dots, \alpha_r$, i.e., $\sum_{i=1}^r A |\alpha_i\rangle \langle \alpha_i|$, $\sum_{i=1}^r |\alpha_i\rangle \langle \alpha_i| A^\top$, and $\sum_{i=1}^r |\alpha_i\rangle A^\top \langle \alpha_i|$ have the same distribution. More concretely, for $\sum_{i=1}^r A |\alpha_i\rangle \langle \alpha_i|$, the distribution is

$$\begin{aligned} \text{PDF}(\Phi; A^\top A) &= \text{PDF}(\Phi; \mathbb{I}) \frac{\text{PDF}(\langle \alpha_1, \dots, \alpha_r \rangle; A^\top A)}{\text{PDF}(\langle \alpha_1, \dots, \alpha_r \rangle; I)} \\ &= \frac{\exp(-\text{Tr}((A^\top A)^{-1} \Phi)/2) \Delta(\Phi)^{(n-r-1)/2}}{2^{rn/2} \pi^{r(r-1)/4} \Delta(A^\top A)^{r/2} \prod_{i=1}^r \Gamma((n-i+1)/2)}. \end{aligned}$$

We can now prove the privacy guarantee. Let $\delta_0 = \delta/r$. Let A and \tilde{A} be the matrix such that $A - \tilde{A} = E = ve^\top$. The published matrices corresponding to the two neighboring matrices have the following

probability density function

$$\begin{aligned} \text{PDF}(\Phi; A^\top A) &= \frac{\exp(-\text{Tr}((A^\top A)^{-1}\Phi)/2)\Delta(\Phi)^{(n-r-1)/2}}{2^{rn/2}\pi^{n(n-1)/4}\Delta(A^\top A)^{r/2}\prod_{i=1}^r\Gamma((n-i+1)/2)} = C\frac{\exp(-\text{Tr}((A^\top A)^{-1}\Phi)/2)}{\Delta(A^\top A)^{r/2}}, \\ \text{PDF}(\Phi; \tilde{A}^\top \tilde{A}) &= \frac{\exp(-\text{Tr}((\tilde{A}^\top \tilde{A})^{-1}\Phi)/2)\Delta(\Phi)^{(n-r-1)/2}}{2^{rn/2}\pi^{n(n-1)/4}\Delta(\tilde{A}^\top \tilde{A})^{r/2}\prod_{i=1}^r\Gamma((n-i+1)/2)} = C\frac{\exp(-\text{Tr}((\tilde{A}^\top \tilde{A})^{-1}\Phi)/2)}{\Delta(\tilde{A}^\top \tilde{A})^{r/2}}, \end{aligned}$$

where $C = \Delta(\Phi)^{(n-r-1)/2} / (2^{rn/2}\pi^{n(n-1)/4}\prod_{i=1}^r\Gamma((n-i+1)/2))$. As in Blocki *et al.* [6], it is straightforward to see that combination of the following proves differential privacy of the published matrix:

$$\sqrt{\frac{\tilde{\Delta}(A^\top A)}{\tilde{\Delta}(\tilde{A}^\top \tilde{A})}} \in \exp(\pm\varepsilon/r) \quad \text{and} \quad \Pr\left[|\text{Tr}\left(\left((A^\top A)^{-1} - (\tilde{A}^\top \tilde{A})^{-1}\right)\Phi\right)| \leq \varepsilon\right] \geq 1 - \delta.$$

Let $\sigma_1 \geq \dots \geq \sigma_d \geq \sigma_{\min}$ be the singular values of A . Let $\lambda_1, \geq \dots \geq \lambda_d \geq \sigma_{\min}$ be the singular value for \tilde{A} . Since the singular values of $A - \tilde{A}$ and $\tilde{A} - A$ are the same, $\sum_{i \in G}(\sigma_i - \lambda_i) \leq 1$ using Linski's theorem, where G is the set of indices for which $\sigma_i > \lambda_i$. The first bound follows similarly as in Blocki *et al.* [6]. For the second bound required for the privacy, we first bound the following

$$\begin{aligned} \text{Tr}\left(\left((A^\top A)^{-1} - (\tilde{A}^\top \tilde{A})^{-1}\right)\Phi\right) &= \text{Tr}\left(\left((A^\top A)^{-1}(\tilde{A}^\top \tilde{A})(\tilde{A}^\top \tilde{A})^{-1} - (\tilde{A}^\top \tilde{A})^{-1}\right)\Phi\right) \\ &= \text{Tr}\left(\left((A^\top A)^{-1}(A + E)^\top(A + E)(\tilde{A}^\top \tilde{A})^{-1} - (\tilde{A}^\top \tilde{A})^{-1}\right)\Phi\right) \\ &= \text{Tr}\left(\left((A^\top A)^{-1}(A^\top E + E^\top \tilde{A})(\tilde{A}^\top \tilde{A})^{-1}\right)\Phi\right). \end{aligned}$$

Using the singular value decomposition of $A = U\Sigma V^\top$ and $\tilde{A} = \tilde{U}\Lambda\tilde{V}^\top$, and the fact that $E = ve_i^\top$ for some i , we can further solve the above expression.

$$\begin{aligned} \text{Tr}\left(\left((A^\top A)^{-1} - (\tilde{A}^\top \tilde{A})^{-1}\right)\Phi\right) &= \text{Tr}\left(\left(V\Sigma^{-1}U^\top|e_i\rangle\langle v|\tilde{V}\Lambda^{-2}\tilde{V}^\top + V\Sigma^{-2}V^\top|v\rangle\langle e_i|\tilde{U}\Lambda^{-1}\tilde{V}^\top\right)\Phi\right) \\ &= \text{Tr}\left(V\Sigma^{-1}U^\top|e_i\rangle\langle v|\tilde{V}\Lambda^{-2}\tilde{V}^\top\Phi\right) + \text{Tr}\left(V\Sigma^{-2}V^\top|v\rangle\langle e_i|\tilde{U}\Lambda^{-1}\tilde{V}^\top\Phi\right) \\ &= \sum_{j=1}^r \text{Tr}\left(\langle\alpha_j|V\Sigma^{-1}U^\top|e_i\rangle\langle v|\tilde{V}\Lambda^{-2}\tilde{V}^\top|\alpha_j\rangle\right) \\ &\quad + \sum_{j=1}^r \text{Tr}\left(\langle\alpha_j|V\Sigma^{-2}V^\top|v\rangle\langle e_i|\tilde{U}\Lambda^{-1}\tilde{V}^\top|\alpha_j\rangle\right). \end{aligned}$$

Fix a j . We bound the following.

$$|\text{Tr}\left(\langle\alpha_j|V\Sigma^{-1}U^\top|e_i\rangle\langle v|\tilde{V}\Lambda^{-2}\tilde{V}^\top|\alpha_j\rangle\right) + \text{Tr}\left(\langle\alpha_j|V\Sigma^{-2}V^\top|v\rangle\langle e_i|\tilde{U}\Lambda^{-1}\tilde{V}^\top|\alpha_j\rangle\right)|. \quad (7)$$

We now look at each term in the above expression. $\langle\alpha_j|V\Sigma^{-2}V^\top|v\rangle$ is distributed as $\mathcal{N}(0, \|V\Sigma^{-2}V^\top|v\rangle\|^2)$, $\langle v|\tilde{V}\Lambda^{-2}\tilde{V}^\top|\alpha_j\rangle$ as $\mathcal{N}(0, \|\langle v|\tilde{V}\Lambda^{-2}\tilde{V}^\top\|^2)$, $\langle\alpha_j|V\Sigma^{-1}U^\top|e_i\rangle$ as $\mathcal{N}(0, \|V\Sigma^{-1}U^\top|e_i\rangle\|^2)$, and $\langle e_i|\tilde{U}\Lambda^{-1}\tilde{V}^\top|\alpha_j\rangle$ as $\mathcal{N}(0, \|\langle e_i|\tilde{U}\Lambda^{-1}\tilde{V}^\top\|^2)$. Since v and e_i are unit vectors, the norm of the above four quantities are less than $1/\sigma_{\min}$, $1/\sigma_{\min} + 1/\sigma_{\min}^2$, 1 , and $1 + 1/\sigma_{\min}$, respectively.

Therefore, from the concentration inequality of Gaussian distribution, we have

$$\Pr\left[(7) \leq 2\left(\frac{1}{\sigma_{\min}} + \frac{1}{\sigma_{\min}^2}\right)\ln(4/\delta_0) \leq \varepsilon\right] \geq 1 - \delta_0.$$

Taking union bound, we have with probability $1 - \delta$, $-\varepsilon \leq \text{Tr} \left(\left((A^\top A)^{-1} - (\tilde{A}^\top \tilde{A})^{-1} \right) \Phi \right) \leq \varepsilon$. Adjusting and renaming the value of ε , we have the result. \square

4 Update Time Efficient Private Sketch Generation

In this section, we use the projection matrix of Upadhyay [61] to get the improvement over naive private sketch generation mechanism as far as update time is concerned. One can also use the fast-Johnson Lindenstrauss transform which also preserves differential privacy [60]. Let $\alpha := \langle \alpha_1, \dots, \alpha_n \rangle$ be n i.i.d. Gaussian samples and C be a circulant matrix formed using the vector α , i.e, for $1 \leq i \leq n$, $C_{i:} = \langle \alpha_i, \dots, \alpha_n, \alpha_1, \dots, \alpha_{i-1} \rangle$. Then the matrix P corresponding to a permutation matrix Π truncated to r rows and circulant matrix C is formed by choosing r rows of the circulant matrix (i.e., $P = \Pi_{1..r} C$). This approach has been used by Rudelson and Vershynin [54] to prove that certain matrix satisfies the Restricted Isometry Property. This choice of sampling r rows combined by sampling a Gaussian vector gives us a family of matrices $\mathcal{P} = \{P\}_\Pi$, the size of this family being $\binom{n}{r}$. Of this, a special case is the partial circulant matrix, which was used by Vybiral [62], in which the first r rows of the matrix C is chosen deterministically to form the matrix P . Note that this is not the only option. One can also pick r rows deterministically by specifying some known fixed permutations, like Affine transformation or combinatorial designs, like r random rows of Latin squares, but they are also a special case of the family \mathcal{P} . and the re-defined mechanism for private sketch generation uses $\frac{1}{\sqrt{r}} P W^{(n)} D$ instead of Ω in Figure 1. Upadhyay [61] proved the following.

Theorem 4.1. Let $\varepsilon \in (0, 1)$ be a constant. Let Ω be an $r \times n$ matrix as constructed above, where $r = c\varepsilon^{-2} \log m \log n$ for a large global constant c . Then for any set \mathcal{S} of m vectors $\mathbf{x} \in \mathbb{R}^n$, the following holds with probability at least $2/3$,

$$(1 - \varepsilon) \|\mathbf{x}\|_2^2 \leq \|\Omega \mathbf{x}\|_2^2 \leq (1 + \varepsilon) \|\mathbf{x}\|_2^2. \quad (8)$$

Moreover, if the singular values of an $n \times d$ input matrix A is at least $\sqrt{\frac{16n \log(2/\delta)}{\varepsilon}}$, then publishing $A^\top \Omega$ is (ε, δ) -differentially private. The time taken to compute $A^\top \Omega$ is $O(nd \log n)$.

The sketch of the matrix is generated row (or column) wise using one of the variants presented in Figure 1.

5 Applications of Private Sketch Generation

Our mechanisms assume that the matrix A is provided as the symmetric $\text{rank}(A)$ matrix A' corresponding to A (note the steps 2 and 3 in Figure 2 and step 2 in Figure 3), and stop updating the data structure once d rows (or columns) are streamed. This simplifies the presentation as well as the analysis. By the argument of Hardt and Roth [34], this leads to a depreciation of the privacy guarantee by at most half. The utility proof does not change for matrix multiplication and linear regression if we use A or A' . However, the analysis for utility gets more complicated in the case of LRA because right and left singular vectors are different. Keeping this in mind, we present its analysis in the most general form by working with its SVD.

We need more care to design the mechanism for matrix product and linear regression. For example, if we naively use the streamed matrix in PSG and then use the product estimates, we end up getting an additive error proportional to the Frobenius norm of input matrices. On the other hand, after the transformation to A' and B' of any conforming matrices A and B , respectively, we can use the identity,

Initialization: On input parameters $\alpha, \beta, \varepsilon, \delta$, set $r = O(\log(1/\beta)/\alpha^2)$. Pick an $r \times 2(n+d)$ matrix as in Section 4, Ω . Set $s = \frac{16r \ln(2/\delta)}{\varepsilon} \ln(16r/\delta)$. Set Y_{A_0} and Y_{B_0} to be all zero matrix.

On input the row r of a $d \times n$ matrix A and column c of an $n \times d$ matrix B , the mechanism does the following:

1. Compute $\|A_{r,:}\|_F$ and $\|B_{:,c}\|_F$ and update the values of $\|A_t\|_F$ and $\|B_t\|_F$.
2. Set $\tilde{A}_{r,:} = s \begin{pmatrix} e_r & \langle 0^{n+d} | & A_{r,:} \end{pmatrix}$ and $\tilde{B}_{:,c} = s \begin{pmatrix} e_c & \langle 0^{n+d} | & B_{:,c} \end{pmatrix}$.
3. Invoke variant 1 of PSG with inputs $(\tilde{A}_{r,:}, \text{flag} = 0, r, \varepsilon, \Omega)$ and $(\tilde{B}_{:,c}, \text{flag} = 1, r, \varepsilon, \Omega)$. Append the rows and columns returned to update the sketch Y_{A_r} and Y_{B_c} .

On query to compute the matrix product at time t , compute the product using $Y_{A_t}/s, Y_{B_t}/s$.

Figure 2: The Mechanism for Matrix Product

$(\mathbb{I} \ A') \ (\mathbb{I} \ B')^\top = (\mathbb{I} + A'B'^\top)$, to perturb the input matrix with a careful choice of parameters. Intuitively, the identity term of the published matrix causes the additive error. We use the same idea for linear regression as well. To prove the privacy of linear regression and matrix product, we use the first variant, while for low-rank approximation, we need both the variants defined in Figure 1. The details are in the respective sections.

We note that this is not the only method that could be used for the first two problems. However, an alternative would require some changes in step 2: compute the SVD using one of the standard algorithms [52, 58] and then add scaled e_t to the matrix. Surprisingly, there are streaming algorithms for computing the SVD of a matrix [52, 58]—by say, setting the target rank in the algorithm of [52] to be $\text{rank}(A)$. However, this makes the mechanisms more complicated with added utility loss. For the sake of coherence and simplicity, we do not explore this idea any further.

5.1 Matrix Multiplication

In this section, we present our mechanism to compute the matrix product (see Figure 2).

Theorem 5.1. Let Ω be an $r \times 2(n+d)$ random matrix as described in Section 4. Then, the mechanism in Figure 2 preserves (ε, δ) -differential privacy while using $O(d\alpha^{-2}\kappa \log(1/\beta))$ bits of space to compute the matrix product of two conforming matrices A and B with an additive error, $\tau \leq \sqrt{n}\alpha$.

Proof. It is easy to see that the space required to store the sketch is at most $O(d^2\kappa \log(1/\beta)/\alpha)$ bits. Similarly, for storing the norms and other things in the data structure, we use at most rd entries; therefore, it does not change the space complexity in terms of $O(\cdot)$ bound.

The privacy guarantee follows because the mechanism maintains the singular values of input matrix above the threshold of Theorem 3.1 ($\tilde{A}_i \tilde{A}_i^\top \succeq U(32r \ln(2/\delta)/\varepsilon) \ln(16r/\delta)^2 \mathbb{I} U^\top = \sigma_{\min}^2 \mathbb{I}$ for both $i = \{1, 2\}$). The proof of utility readily follows from the proof of Clarkson-Woodruff [14]. To prove the utility guarantee of our mechanism, we need a variance bound on $\|A\Omega^\top \Omega B - AB\|_F^2$. Using this, the rest of the proof is just arranging the terms, bounding $\|\Omega^\top \Omega - \mathbb{I}\|_2$, and using sub-additivity of norm. The following lemma corresponding to our projection matrix in Section 4 follows immediately from using the proof of the corresponding lemma of Sarlos [56, Corollary 4].

Lemma 5.2. Let Ω be a $k \times d$ matrix as constructed in Section 4 with every entries picked from the distribution $\mathcal{N}(0, 1)$, then for a set of m vectors, $v_1, \dots, v_m \in \mathbb{R}^n$, with probability at least $1 - 2 \exp(-k\alpha^2/8)$, for any pair v_i, v_j , we have $|\langle \Omega v_i, \Omega v_j \rangle - \langle v_i, v_j \rangle| \leq \alpha \|v_i\| \cdot \|v_j\|$.

Initialization: On input parameters $\alpha, \beta, \varepsilon, \delta$, set $r = O(d \log(1/\beta)/\alpha)$. Pick an $r \times 2(n+d)$ matrix, Ω , as constructed in Section 4. Set $s = \frac{32 \log(1/\beta) \ln(2/\delta)}{\alpha \varepsilon} \ln(16r/\delta)$ and Y_{A_0} and Y_{B_0} to be all zero matrix. On input the column c of an $n \times d$ matrix A , the mechanism does the following:

1. Set $\tilde{A}_{:c} = s \begin{pmatrix} e_c & \langle 0^{n+d} | & A_{:c} \end{pmatrix}$.
2. Invoke variant 1 of PSG with inputs $(\tilde{A}_{:c}, \text{flag} = 0, r, \varepsilon, \Omega)$. Append the column returned to get an updated sketch Y_{A_c} .
3. On being queried with a set of query vectors $B = \{b_1, \dots, b_m\}$, invoke PSG as in step 1 with input matrix $(B, \text{flag} = 1, r, \varepsilon, \Omega)$ to get the sketch Y_B by appending the new column.

Compute an X satisfying $\min_X \|Y_{A_t} X - Y_B\|$.

Figure 3: The Mechanism for Linear Regression

The proof of the utility of mechanism in Figure 2 follows readily from Lemma 5.2. We need to upper bound the quantity $\|AB - \tilde{A}\Omega^\top\Omega\tilde{B}/s^2\|_F$. Now

$$\tilde{A}\Omega^\top\Omega\tilde{B} = s^2 \begin{pmatrix} \mathbb{I} & 0 & 0 & A \end{pmatrix} \Omega^\top \Omega \begin{pmatrix} \mathbb{I} & 0 & 0 & B^\top \end{pmatrix}^\top = s^2 (\mathbb{I}\Omega^\top\Omega\mathbb{I} + A\Omega^\top\Omega B).$$

Therefore,

$$\|AB - \tilde{A}\Omega^\top\Omega\tilde{B}/s^2\|_F = \|AB - A\Omega^\top\Omega B - \mathbb{I}\Omega^\top\Omega\mathbb{I}\|_F \leq \|AB - A\Omega^\top\Omega B\|_F + \|\mathbb{I} - \Omega^\top\Omega\|_F. \quad (9)$$

To bound the first term, let random variable X_{ij} denote $(AB)_{ij} - (A\Omega^\top\Omega B)_{ij}$. Then, with probability at least $1 - 2 \exp(-k\alpha^2/8)$, we have $|X_{ij}| \leq \alpha \|A_{i:\cdot}\|_2 \cdot \|B_{:j}\|_2$, and using Lemma 5.2,

$$\|AB - A\Omega^\top\Omega B\|_F^2 = \sum |X_{ij}|^2 \leq \sum \alpha^2 \|A_{i:\cdot}\|_2^2 \|B_{:j}\|_2^2 \leq \alpha^2 \|A\|_F^2 \|B\|_F^2. \quad (10)$$

For the second term, we need to bound the variance on unitaries. For any unitaries U_1 and U_2 , the analysis of Kane and Neilson [38] gives $\|U_1^\top \Omega^\top \Omega U_2 - U_1 U_2\|_2 \leq \alpha$. Plugging this and equation (10) in equation (9), using the fact that $\|X\|_2 \leq \|X\|_F \leq \sqrt{n} \|X\|_2$ for any $n \times n$ matrix X , and adjusting the value of α , the result follows. \square

5.2 Linear Regression

The mechanism for linear regression follows from the observation that PSG is a low distortion embedding of the vectors. Therefore, if one solves the linear regression problem in the lower dimensional space, it should translate to the higher dimension. The technical part is to prove that this intuition is in fact correct.

Theorem 5.3. Let Ω be an $r \times 2(n+d)$ matrix, where $r = O(d \log(1/\beta)/\alpha)$. Then for every constant $c > 32$, with probability at least $1 - \beta$, if \tilde{X} is the solution of $\min_X \|\Omega(AX - B)\|^2$ for an $n \times d$ input matrix, then the mechanism in Figure 3 uses $O(d^2 \alpha^{-1} \kappa \log(1/\beta))$ bits and solves the linear regression problem in an (ε, δ) -differentially private manner with additive error $\tau \leq O(\sqrt{n}\alpha)$.

This is better than the bound achieved by Jain, Kothari, and Thakurta [37] for $\alpha < 1$, which is a more practical choice of parameter. Jain, Kothari, and Thakurta [37] gave a private learning algorithm for linear regression and achieve a bound of $\tilde{O}((R^6 \log(1/\delta) \sqrt{n} \log^{1.5} T) / \sqrt{\varepsilon} \alpha^3)$, where R is the maximum 2-norm

On input parameters $\alpha, \beta, \varepsilon, \delta$, the target rank k , set $w = \frac{16k \ln(2/\delta)}{\varepsilon}$. Pick a $2n \times k$ standard Gaussian matrix Ω . On input an $n \times n$ matrix A of rank r , the mechanism does the following:

Range Finding. Compute $Y_A = (w\mathbb{I} \ A) \Omega$ as in Figure 2 and Figure 3. Let $\Pi_{Y_A} = \Psi\Psi^\top$ be the projection matrix corresponding to the range of Y_A .

Projection: When the whole matrix is streamed, the curator does the following:

1. Let the (unknown) matrix $B = \Psi^\top A_t \Psi$.
2. Use the minimal residual method to find a solution to $B\Psi_t^\top \Omega = \Psi_t^\top Y_t$.
3. Compute the decomposition of $B_t = \bar{U}_t \Lambda_t \bar{U}_t^\top$, form the product $\hat{U}_t = \Psi_t \bar{U}_t$, and publish $\hat{U}_t \Lambda_t \hat{U}_t^\top$.

Figure 4: The Mechanism for k -Rank Approximation

of query input and T is the number of training data set². Our approach to prove the utility bound for linear regression follows from straightforward application of the result of Kane and Nelson [38] to the proof of Clarkson and Woodruff [14]: first bound $\|U^\top(A\tilde{X} - X^*)\|_F$, where X^* is the right value of the linear regression, and \tilde{X} is the value of regression for the sketch of A and B using Lemma 5.4, and then apply Pythagorus theorem and an observation that A and U have the same column-space.

Lemma 5.4. (Sarlos [56], Kane and Neilson [38]) Given $r = O(d \log(1/\beta)/\alpha)$. Let U be any unitary matrix. If Ω satisfies the Johnson-Lindenstrauss bound, then with probability at least $1 - \beta$, we have $\|U\Omega^\top \Omega U^\top - \mathbb{I}\|_2 \leq \alpha$.

For the sake of completion, we give the detailed proof in Appendix A.

5.3 Low Rank Approximation

The mechanism of Hardt and Roth [34] is a privacy-preserving version of the prototype mentioned in [31]. They improve the worst case lower bound under a *low coherence* assumption. The main prototype in [31] is the following: construct a low-dimensional subspace that captures the action of the matrix (*range-finding*) and restrict the matrix to that subspace to compute the required factorization (*projection*). More concretely, range finding finds a measurement matrix $Y = A\Omega$, where Ω is a Gaussian matrix and computes the orthonormal projection matrix Π_Y corresponding to the range defined by Y ; projection then computes a rank k matrix $B = \Pi_Y A$. From this exposition, it seems that privacy preserving algorithms are required for both the stages.

We show that the two-step prototype mentioned in [31] can be replaced by a two-step algorithm in which the input matrix is explicitly needed only in the first step. Let $\hat{U}\hat{\Lambda}\hat{V}^\top$ be the factorization of the low rank approximation of A we aim to achieve. The crucial observation for the single pass LRA when Ω is a Gaussian matrix is that Ω , Y , and the basis for the range of Y contains enough information to compute the matrix B . Since we are using Ω twice, we have to rely on privacy of the second variant of PSG algorithm. We first note that if $\Psi\Psi^\top A$ is a LRA of A , i.e., $\|A - \Psi\Psi^\top A\| \leq \eta$, then so is $\Psi\Psi^\top A\Psi\Psi^\top$. This is because $\|A - \Psi\Psi^\top A\Psi\Psi^\top\| = \|A - \Psi\Psi^\top A + \Psi\Psi^\top A - \Psi\Psi^\top A\Psi\Psi^\top\| \leq \|A - \Psi\Psi^\top A\| + \|\Psi\Psi^\top A - \Psi\Psi^\top A\Psi\Psi^\top\| \leq 2\eta$. In order to simplify the presentation, we state our mechanism for symmetric matrices in Figure 4 (note that [35, 39] also made this assumption).

In Figure 4, we give a mechanism for symmetric matrices. To construct a mechanism for non-symmetric

²It should be noted though that the setting of private learning algorithm is stricter than streaming model and might be one of the reason why we achieve such significant improvement.

matrices, we construct two sketches Y_t and \tilde{Y}_t corresponding to A_1 and A_2 using a single pass over A and using two Gaussian matrices Ω and $\tilde{\Omega}$ of appropriate dimension, where $A_1 = (s\mathbb{I} \ A)$ and $A_2 = (A^\top \ s\mathbb{I})$ for appropriate dimension identity matrices in each of A_1 and A_2 . Basically, we do the following using a single pass over the matrix A :

$$\begin{pmatrix} Y \\ \tilde{Y} \end{pmatrix} = \begin{pmatrix} s\mathbb{I}_n & A \\ A^\top & \end{pmatrix} \begin{pmatrix} \Omega \\ \tilde{\Omega} \end{pmatrix},$$

where \mathbb{I}_n is an $n \times n$ identity matrix. Since $(Y \ \tilde{Y})^\top$ corresponds to a symmetric matrix, we can use the steps used in the projection stage in Figure 4. We state our result for the general case.

Theorem 5.5. Let $\lambda_1 \geq \dots \geq \lambda_{\text{rk}(A)}$ be the singular values of A . Then for an over-sampling parameter p , there is a single-pass mechanism that compute LRA using $O(k(n+d)\alpha^{-2} \log(nd))$ bits while preserving (ε, δ) -differential privacy such that

$$\|A - \tilde{A}\|_F \leq \left(1 + \frac{k}{p-1}\right)^{1/2} \min_{\text{rk}(A') < k} \|A - A'\|_F + \frac{2k\sqrt{(n+d)\ln(2/\delta)}}{\varepsilon}, \quad \text{and} \quad (11)$$

$$\|A - \tilde{A}\|_2 \leq \left(1 + \frac{k}{p-1}\right)^{1/2} \lambda_{k+1} + \frac{e\sqrt{(k+p)\sum_{j>k}\lambda_j^2}}{p} + \frac{2\sqrt{k(n+d)\ln(2/\delta)}}{\varepsilon}. \quad (12)$$

We need both the variants of PSG algorithm for the privacy proof—we use the first variant for range finding step and the second variant because we reuse Ω in the second step of the projection stage. Note that the first step pushes the singular values above the threshold of Theorem 3.1. Since rest of the computations are deterministic, privacy follows from Lemma 2.2, Theorems 2.1 and 3.1, and our choice of w . The space guarantee is also straightforward. Now, Π_Y has a decomposition, $\Psi\Psi^\top$, for some orthonormal basis Ψ , which also forms an orthonormal basis for the approximated matrix \tilde{A} after the run of the algorithm. Thus, B must satisfy $B\Psi^\top\Omega \approx \Psi^\top Y$. This is what step 2 does. Therefore, $\tilde{A}_t = \Pi_Y A_t$ for the projection operator Π_Y with the same range as Ψ . The rest of the utility proof relies heavily on the fact that $(\Omega\Omega^\top)^{-1}$ exists, left singular vectors does not play any essential role in the concentration bound, and the rotational invariance of multivariate Gaussian distribution. We bound the norm of $\|(\mathbb{I} - \Pi_Y)A\|$ for both the spectral and Frobenius norm. Once we have the bound on $\|(\mathbb{I} - \Pi_Y)A\|$, we use standard results in random matrix theory to get the final bounds. We note that our proof is very different from the proof of Sarlos [56] which is for two-pass algorithm and Clarkson-Woodruff [14], and more in line with the proof of Halko *et al.* [31]. The proof of Sarlos [56] and Clarkson-Woodruff [14] uses the error bound computed in the estimate of matrix multiplication, while that used here is based on perturbation theory. The details follows.

Proof. The space complexity is easy to follow. The privacy guarantee follows from Theorem 3.1 and noting that all the singular-values of the matrix on which Ω is operated from the right is greater than the threshold required for the statement of the Theorem 3.1, and the distribution of $\Psi^\top\Omega = \Psi^\top\Pi_Y\Omega$ is the same as that of the second variant as we reuse Ω . The latter statement can be shown using the proof of Bura and Pfeiffer [10]. They showed that if a random matrix has a normal distribution in limits, then its left singular matrix has a normal limit distribution. Using their idea for normally distributed matrix, the left singular matrix has a normal distribution³. A simplified argument could be the following. Since the entries of $\Omega \in \mathbb{R}^{n \times r}$ is $\mathcal{N}(0, 1)$, then for any orthogonal matrices $G \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{r \times r}$, the entries of $G\Omega R^\top$

³One gets a different covariance matrix, scaled by an orthonormal matrix, from that of $A^\top\Omega$, but since in our proof, we are only concerned with the singular values of the covariance matrices, this does not effect our analysis.

is also i.i.d. normal. This can be seen by translating to the vector form of the matrix, i.e., $v = \text{vec}(V)$ represent rn vector with entries $v_{i+nj} = V_{ij}$. Then $\text{vec}(G\Omega R^T) = (G \otimes R)\text{vec}(\Omega)$ using Lemma 2.8. Now $G \otimes R$ is also an orthogonal matrix, and multivariate Gaussian distribution is preserved if one multiply by an orthogonal matrix. Therefore, the distribution of the left singular vectors of Ω is the same as $G\Omega R^T$ (owing to the invariance under rotation). Consequently the distribution of each singular vector is also spherically distributed. It is easy to see that Ψ is the left singular matrix of Y ; therefore, from the above argument, it is also spherically distributed.

Now it follows from the proof of the second variant (Theorem 3.1) that it does not matter if we multiply A (or A^T) from left (or right, respectively) of vectors $\alpha_1, \dots, \alpha_r$, i.e., $\sum_{i=1}^r A|\alpha_i\rangle\langle\alpha_i|$, $\sum_{i=1}^r |\alpha_i\rangle\langle\alpha_i|A^T$, and $\sum_{i=1}^r |\alpha_i\rangle A^T \langle\alpha_i|$ have the same distribution. Combining all these arguments, we have the distribution of the second step of projection stage is identical to the second variant in Figure 1, modulo some deterministic computation. Since, any arbitrary preprocessing preserves differential privacy, we can now complete the proof by invoking Theorem 3.1. The privacy guarantee due to Theorem 3.1 requires the minimum singular value to be greater than $\frac{4\sqrt{k \log(2/\delta) \log(k/\delta)}}{\epsilon}$ for the first variant and $\frac{k \log(k/\delta)}{\epsilon}$ for the second variant. By our choice of w , the singular values of the streamed matrix to the algorithm for PSG are at least the eigen-values of $\sqrt{w^2 \mathbb{I} + A^T A}$, which are all greater than $\frac{16k \ln(2/\delta)}{\epsilon}$. Since $\frac{4\sqrt{k \log(2/\delta) \log(k/\delta)}}{\epsilon} \ll 16 \frac{k \log(k/\delta)}{\epsilon}$, the privacy guarantee follows from Theorem 3.1.

For the utility guarantee, we need to show that the mechanism does what [31, Section 1.2] prototype algorithm achieves. The main prototype in [31] and that used by [34] is the following: construct a low-dimensional subspace that captures the action of the matrix (*range-finding*), restrict the matrix to that subspace, and compute a standard factorization (*projection*). More concretely, range finding finds a measurement matrix $Y = A\Omega$, where Ω is a Gaussian matrix and computes the orthonormal projection matrix $\Pi_Y = Y(Y^T Y)^{-1} Y^T$; projection then computes a rank k matrix $B = \Pi_Y A$. Note that $\Pi_Y = \Psi \Psi^T$ for some orthonormal basis Ψ .

The utility guarantee needs a slight tweak from [31] though overall structure remains the same. Mainly, we need to worry about the computation of the minimal residual method and the requirement on the input to PSG algorithm. We include it here to make sure that we achieve the bound as claimed.

Keeping the most general case in mind, we first show that the left singular value has hardly any role to play in bounding the perturbation. We assume that we perform SVD – the case for eigenvalue decomposition follows similarly. Let the SVD of A be $U\Sigma V^T$. In the following discussion, we compute the approximation of $(w\mathbb{I} - A)$ and denote it by \tilde{A} . This is because this matrix is more manageable for computation and any upper bound on the approximation of this matrix is an upper bound on the approximation of the original matrix. The actual bound on the approximation of the private matrix as computed in the main text can be obtained by simply performing the computation on the singular value decomposition as performed in the last step of mechanism and using the sub-additivity of norms.

From the discussion above, we know that $\tilde{A} = \Pi_Y A$; therefore, we need to bound $\|(\mathbb{I} - \Pi_Y)A\|$, where $\|\cdot\|$ in this section refers to both the Frobenius as well as the spectral norm. From the Hölder's inequality on the second moment, we have

$$\mathbb{E}[\|(\mathbb{I} - \Pi_Y)A\|_F] \leq \left(\mathbb{E} [\|(\mathbb{I} - \Pi_Y)A\|_F^2] \right)^{1/2}. \quad (13)$$

We now bound $\|(\mathbb{I} - \Pi_Y)A\|$. Let $\Lambda = \sqrt{\Sigma^2 + w^2}$. Let Λ_1 denote the diagonal matrix formed by the first k singular values and Λ_2 be the diagonal matrix for the other singular values. We decompose V^T similarly. Let the matrix formed by the first k rows of V^T be V_1^T and by the rest of the rows be V_2^T .

Left singular vectors have essentially no role in the approximation bound. Recall that Ω is a $n \times k$ matrix; therefore, $Y = A\Omega$ and $Y = U (\Lambda_1 V_1^\top \Omega \quad \Lambda_2 V_2^\top \Omega)^\top$. It would be useful to consider the first k rows of Y to be the one that mimics the action of A and the rest of the rows of Y as a small perturbation that we wish to bound. We first prove that the left singular vectors have essentially no role to play in bounding the error. Let $A' = UA$, then the following chain of equalities are straightforward.

$$\|(\mathbb{I} - \Pi_Y)A\| = \|U^\top(\mathbb{I} - \Pi_Y)A\| = \|U^\top(\mathbb{I} - \Pi_Y)UA'\| = \|\mathbb{I}A' - U^\top\Pi_Y UA'\|. \quad (14)$$

Now note that for an orthogonal projector corresponding to a matrix Y is uniquely defined by $\text{range}(Y)$, the range of Y . Therefore,

$$\text{range}(U^\top\Pi_Y U) = U^\top\text{range}(\Pi_Y) = \text{range}(U^\top\Pi_Y). \quad (15)$$

This gives us that $\|A' - U^\top\Pi_Y UA'\| = \|(\mathbb{I} - \Pi_{U^\top Y})A'\|$. A useful way to understand the above expression is to view this geometrically and recall that unitary are just rotation in the space: projection by an unitary, followed by any projection operator, followed by the inverse of unitary brings us to the same space as projection by an operator followed by the inverse of the unitary.

Finding and bounding an appropriate perturbed matrix. Our next idea is to use the identity that for two operators O_1 and O_2 , if the range of O_1 is a subset of the range of O_2 , then the projection of any matrix using O_1 will have all its norm smaller than that by O_2 . More concretely, we find a matrix C such that its range is a strict subset of the range of $U^\top Y$. We obtain this matrix by flattening out the first k rows of $U^\top Y$. This is in correspondence with our earlier observation that the first k rows mimic the action of A and rest of the rows are the perturbation that we wish to bound. Since the first k rows of $U^\top Y$ is $\Lambda_1 V_1^\top \Omega$, let $C := U^\top Y \Omega^{-1} V_1 \Lambda_1^{-1}$. The rows corresponding to the perturbation are $\Lambda_2 V_2 \Omega$. Thus, $C = (\mathbb{I} \quad \Lambda_2 V_2^\top V_1 \Lambda_1^{-1})^\top$.

Let us denote by $S = \Lambda_2 V_2^\top V_1 \Lambda_1^{-1}$. It is not difficult to see that $\text{range}(C) \subset \text{range}(U^\top Y)$. Moreover, $\Pi_C \preceq \mathbb{I}$, $\Pi_{U^\top Y} \Pi_C \Pi_{U^\top Y} \preceq \Pi_{U^\top Y}$. This follows from the fact that $\text{range}(C) \subset \text{range}(U^\top Y)$ and the following derivation

$$\Pi_{U^\top Y} \succeq \Pi_{U^\top Y} \Pi_C \Pi_{U^\top Y} = \Pi_C \Pi_{U^\top Y} = (\Pi_{U^\top Y} \Pi_C)^\top = \Pi_C.$$

An immediate result of the above is the following:

$$\|(\mathbb{I} - \Pi_{U^\top Y})A'\| \leq \|(\mathbb{I} - \Pi_C)A'\|. \quad (16)$$

Since, $\Pi_C = C(C^\top C)^{-1}C^\top$, we have the following set of derivations.

$$\begin{aligned} \Pi_C &= \begin{pmatrix} \mathbb{I} \\ S \end{pmatrix} \left[(\mathbb{I} \quad S^\top) \begin{pmatrix} \mathbb{I} \\ S \end{pmatrix} \right]^{-1} (\mathbb{I} \quad S^\top) = \begin{pmatrix} \mathbb{I} \\ S \end{pmatrix} [(\mathbb{I} + S^\top S)]^{-1} (\mathbb{I} \quad S^\top) \\ &= \begin{pmatrix} \mathbb{I}(\mathbb{I} + S^\top S)^{-1} \\ S(\mathbb{I} + S^\top S)^{-1} \end{pmatrix} (\mathbb{I} \quad S^\top) = \begin{pmatrix} (\mathbb{I} + S^\top S)^{-1} & (\mathbb{I} + S^\top S)^{-1} S^\top \\ S(\mathbb{I} + S^\top S)^{-1} & S(\mathbb{I} + S^\top S)^{-1} S^\top \end{pmatrix} \\ &\succeq \begin{pmatrix} (\mathbb{I} - S^\top S) & (\mathbb{I} + S^\top S)^{-1} S^\top \\ S(\mathbb{I} + S^\top S)^{-1} & 0 \end{pmatrix}, \end{aligned}$$

where the last inequality uses the fact that $\mathbb{I} - S^\top S \preceq (\mathbb{I} + S^\top S)^{-1}$ and $S(\mathbb{I} + S^\top S)^{-1} S^\top \succeq 0$. Therefore,

$$\mathbb{I} - \Pi_C \preceq \begin{pmatrix} S^\top S & \mathbb{I} - (\mathbb{I} + S^\top S)^{-1} S^\top \\ \mathbb{I} - ((\mathbb{I} + S^\top S)^{-1} S^\top)^\top & \mathbb{I} \end{pmatrix}.$$

Conjugating $\mathbb{I} - \Pi_C$ with Λ , and applying the fact that for every positive definite matrix, $P = \begin{pmatrix} X & Y \\ Y^\top & Z \end{pmatrix}$, we have $\|P\| \leq \|X\| + \|Z\|$, we get

$$\|(\mathbb{I} - \Pi_C)A'\| \leq \|S^\top SA'\| + \|A'\| \quad (17)$$

for any norm. From here on, it is easy arithmetic to show that

$$\|(\mathbb{I} - \Pi_C)A'\| \leq \sqrt{\|\Lambda'_2\| + \|\Lambda'_2 V_2^\top \Omega (V_1^\top \Omega)^{-1}\|} \quad (18)$$

for both the required norms. Using equation (14), equation (16) and equation (17), this gives us a bound on the approximation of matrix A' . Till this point, our analysis closely follows the ideas of [31], accommodating the steps of our algorithm. Now, all that remains is to bound $\|\Lambda_2 V_2^\top \Omega (V_1^\top \Omega)^{-1}\|$, and for this, we have to analyze the matrix Ω .

5.3.1 Error bound for Frobenius norm

We now exploit the rotational invariance of a Gaussian distribution. An important point to note is that $(\Omega \Omega^\top)^{-1}$ exists and has a well defined trace. The first part of the right hand side of equation (18) is immediate. Thus, if we bound $\mathbb{E}[\|\Lambda'_2 V_2^\top \Omega (V_1^\top \Omega)^{-1}\|]$, we are done. This could be accomplished as below.

$$\begin{aligned} \mathbb{E}[\|\Lambda_2 V_2^\top \Omega (V_1^\top \Omega)^{-1}\|] &\leq \sqrt{\mathbb{E} \left[\sum_{ij} |(\Lambda_2)_{ij}' \Pi_{ij} (V_1 \Omega^{-1})_{jj}| \right]} \\ &\leq \sqrt{\|\Lambda'_2\|_F \|\Omega^{-1}\|_F} \\ &= \sqrt{\|\Lambda'_2\|_F \text{Tr} \left((\Omega \Omega^{-1})^\top \right)} = \sqrt{\|\Lambda'_2\|_F \text{Tr}(\Omega \Omega^\top)^{-1}} \\ &\leq \sqrt{\text{Tr}(\Omega \Omega^\top)^{-1}} \min_{rk(A') \leq k} \|A - A'\|_F + \sqrt{(n+d)w \text{Tr}(\Omega \Omega^\top)^{-1}}. \end{aligned}$$

The utility guarantee follows by plugging this value in equation (18), and combining equation (13) and the fact that $(\Omega \Omega^\top)^{-1}$ has a well defined trace $k/(p-1)$ [47] and the relation $\|\Lambda'\|_F = \sqrt{\|\Lambda^2 + w^2 \mathbb{I}\|_F} \leq \|\Lambda\|_F + w \|\mathbb{I}\|_F$.

5.3.2 Error bound for Spectral norm

In order to bound the second term, we use few well known facts in the theory of random matrices to simplify equation (18). In particular, using Lemma 2.9 and 2.10, and Holder's inequality, the statement of the theorem for spectral norm follows. The utility bound then follows using the same arithmetic of representing Λ' in terms of Λ as done in the case of Frobenius norm. In more details, we first bound

$$\begin{aligned} \mathbb{E} \left[\|\Lambda'_2 V_2^\top \Omega (V_1^\top \Omega)^{-1}\| \right] &\leq \|\Lambda'_2\| \left(\mathbb{E} \left[\|(V_1^\top \Omega)^{-1}\|_F^2 \right] \right)^{1/2} + \|\Lambda'_2\|_F \left(\mathbb{E} \left[\|(V_1^\top \Omega)^{-1}\| \right] \right) \\ &\leq \|\Lambda'_2\| \left(\mathbb{E} \left[\|\Omega^{-1}\|_F^2 \right] \right)^{1/2} + \|\Lambda'_2\|_F \left(\mathbb{E} \left[\|\Omega^{-1}\| \right] \right), \end{aligned}$$

and then invoke Lemma 2.10 followed by the sub-additivity of norms. Making these substitution and on simplification, we get the bound stated in equation (12) of Theorem 5.5.

$$(18) \leq \left(1 + \frac{k}{p-1}\right)^{1/2} \|\Lambda_2\|_2 + \frac{e\sqrt{(k+p)\sum_{j>k}\lambda_j^2}}{p} + \frac{2\sqrt{k(n+d)\ln(2/\delta)}}{\varepsilon}.$$

□

We compare our results with the best possible results in the non-private setting. Eckart and Young [23] have shown that the quantity $\min_{\text{rank}(A') < k} \|A - A'\|_F$ in the first term of equation (11) is optimal. Likewise, Mirsky [46] proved that λ_{k+1} is the minimum spectral error for k -rank approximation. The second term in equation (12) shows that we also pay for the Frobenius norm error when doing a unified analysis. However, when the oversampling parameter $p \approx k$, then the factor on λ_{k+1} is constant and that on the second term is of order $k^{-1/2}$. In fact, on closer analysis,

$$\|A - \tilde{A}\|_2 \leq \left(1 + \frac{k}{p-1} + \frac{e(\sqrt{(k+p)\min\{d,n\} - k})}{p}\right) \lambda_{k+1} + \frac{2\sqrt{k(n+d)\ln(2/\delta)}}{\varepsilon},$$

therefore, the error always lies within some polynomial factor of λ_{k+1} , modulo some additive error.

Kapralov-Talwar [39] showed a lower bound on additive error when $\delta = 0$ for neighbouring data differing by unit spectral norm. Our privacy proof depends strongly on the fact that $\delta \neq 0$. In fact, our bound is vacuous if $\delta = 0$. Though incomparable due to difference in the notion of neighbouring data, this separation gap further strengthen the belief that better bounds are possible for $\delta \neq 0$.

5.3.3 Comparison with Earlier Works

In the recent past, there have been five major works that give a tight bound on differentially private low-rank approximation. We now compare our results with these works in more detail.

Hardt-Roth [34]: This work uses two passes over the input matrix; therefore, it does not fall in our one-pass streaming model of computation. They also work in the case of event level privacy—two datasets are neighbouring if they differ in at most one entry of Euclidean norm at most one. This makes their coherence conditions and notion of neighbouring data sets rotationally invariant. We use the same concept of neighbouring data sets; however, as argued by Blocki *et al.* [6], we achieve a better utility bound in the range finding step. Intuitively, this could be seen due to the absence of additive Gaussian noise. Hardt and Roth [34] achieved an error bound of $\sqrt{kn} \log(k/\delta)/\varepsilon + \sqrt{\mu}\|A\|_F(n/d)^{1/2} \log(k/\delta)/\varepsilon$. Without any coherence assumption, their error bound depends on $\|A\|_F$, which can be as large as \sqrt{nd} for binary matrices. On the other hand, we achieve a bound that is independent of $\|A\|_F$.

Hardt-Roth [35]: In some sense, this paper is based on Krylov subspace iteration combined with powering method of Halko *et al.* [31]. The coherence definition used in this paper depends on the maximum value of the left or right singular vectors. This makes their coherence condition rotationally variant. Moreover, they also make an assumption regarding the singular value separation, i.e., the first and the k -th singular value has a non-trivial separation. They define two data-sets as neighbouring in the same manner as in Hardt and Roth [35]. They give LRA in spectral norm, and therefore, their work has application in problems like principal component analysis. Their bound, however, depends on the rank of the input matrix. Their mechanism uses k rounds of subspace generation,

each of which depends on the spectrum of the matrix and uses the power-iteration method of Halko *et al.* [31]; therefore, it cannot be implemented in a streaming fashion. Moreover, the error bound is $O(k^2 \varepsilon^{-1} \sqrt{(\text{rank}(A)\mu + k \log n) \log(1/\delta) \log n})$ compared to $\frac{2\sqrt{k(n+d) \ln(2/\delta)}}{\varepsilon}$ (equation (12)). Note that they have a better multiplicative bound that achieved by us. This is in lieu with the observation made by Halko *et al.* [31] that power iteration could be used to improve the multiplicative bound. Unfortunately, we cannot use power iteration in one-pass.

Kapralov-Talwar [39]: The only assumption this paper makes is that of singular value separation of the same form as in Hardt-Roth [35]. They also give low-rank approximation in the spectral norm. Additionally, they achieve $(\varepsilon, 0)$ -differential privacy, which is not achieved by any other work, including ours. Their definition of neighbouring data sets can be (arguably) considered the most general in the sense that they consider two data sets neighbouring if they differ by at most one in the spectral norm. On the negative side, their mechanism uses k rounds; therefore, it cannot be implemented in a streaming fashion. Since the notion of neighbouring data-sets and privacy guarantee achieved is different from that of ours, we believe our result is incomparable to that of Kapralov and Talwar [39]. However, if we just concentrate on the error bound, they achieve a bound of $O(dk^3/(\varepsilon\gamma^2\delta^2))$ compared to $\frac{2\sqrt{k(n+d) \ln(2/\delta)}}{\varepsilon}$ (equation (12)).

Hardt [32]: In this recent work, Hardt [32] gave a robust subspace iteration mechanism that allows to publish LRA with noise independent of the rank of the input matrix, thereby, resolving one of the open problems in Hardt-Roth [35]. They define two data-sets as neighbouring in the same manner as in Hardt and Roth [35]. However, they also make an assumption on the singular value separation—a separation between the k -th and $(k + 1)$ -th singular value of the input matrix. Their mechanism uses k rounds of subspace generation, each of which depends on the spectrum of the matrix; therefore, it cannot be implemented in a streaming fashion. We achieve a bound of $\frac{2\sqrt{k(n+d) \ln(2/\delta)}}{\varepsilon}$ (equation (12)) compared to $\lambda_1 \sqrt{kn\mu \log(1/\delta) \log(n/\gamma) \log \log(n/\gamma)}/\varepsilon\gamma^{1.5} \lambda_k$ of Hardt [32].

Dwork *et al.* [22]: Dwork *et al.* [22] gave the first single-pass online learning algorithm for private low-rank approximation under the assumption that the rows of the input matrix are normalized (note that, we do not make any such assumption). However, they assume that two data-sets are neighbouring if they differ by at most one row unlike our notion which is the same as Hardt and Roth [34, 35]. Therefore, we do not see any natural way to compare. They give a bound that assumes a lower bound of $k\sqrt{n} \log^2(m/\delta)/\varepsilon^2$ on the optimal value, for $\delta < 1/\text{poly}(n)$. More concretely, if OPT is the optimal value, then their error bound is $O(\sqrt{k\text{OPT}}n^{1/4} \log^2(m/\delta))$. We do not make any of the assumptions made by them and, if we just consider the end result, we achieve a bound which is factor $k\sqrt{n}$ better than theirs (see equation (12)). The case that we are able to bypass their lower bound gives a mathematical indication that the unit norm notion of neighbouring data is strictly weaker than user-level privacy.

An Alternate Range Finding Step. We mention an alternative, though we do not pursue it any further, to the range finding step in Figure 4 for the sake of completion of the argument given in Section 3. One can also compute $Y_A = U(\Sigma^2 + w^2 \mathbb{I}_{r \times r})^{1/2} V^T \Omega$ using standard single-pass streaming algorithms on SVD computation [52, 58]—by say, setting the target rank in [52] to be $\text{rank}(A)$. The projection step and our proof still remains the same. However, we lose on utility due to the SVD computation.

6 Discussions and Open Problems

In this paper, extending on a few of the earlier results [19, 22, 37], we explored differential privacy on streamed data. Ours is the first work that gives positive results for private analogue of sketch based streaming algorithms. Previously, there were known negative results for private analogue of sketch based approach for specific statistical queries [19]. On the backdrop of these two conflicting results, one important open problem is to characterize the class of queries that allow private sketch based mechanism. Another direction of future research is to study privacy of streamed data under different notion of neighbouring data. Another really interesting problem is to get a tighter space lower bound for answering these queries privately in the streaming model. In this paper, we only considered row-wise or column-wise stream. An interesting question is to consider turnstile stream. Another open problem, also suggested by Blocki *et al.* [6], is to get a better error bound by, maybe, using error correcting codes. It is also interesting to verify whether other variants of Johnson-Lindenstrauss, more specifically sparse and randomness efficient variants [2, 38, 42], also preserves privacy or not. Our results in Section 3 and Section 4 also raise a question as to whether there is any space-time tradeoff while performing these tasks privately. A particularly interesting problem is to characterize the class of query function that has a differentially private mechanism in the streaming model. We have a fair bit of understanding of the class of query functions when the whole database is provided to the curator. Our mechanism as provided is strictly non-interactive. An interesting domain of future research is what types of privacy in streaming model of computation one can achieve with interaction. Iterative multiplicative weight mechanisms seems to be the most natural candidate because of how it constructs the data-structure.

ACKNOWLEDGEMENTS. I would like to thank Prateek Jain for the insightful discussion and suggestion of matrices of the form used in Section 4. My discussion with him led to the first two problems studied in this paper. I would like to thank Shweta Agarwal and Ragesh Jaiswal for the discussions in the early stages of this work, Or Sheffet for his input on the initial draft of this work, Abhradeep Guha Thakurta for clarifying the difference between our model and the online learning model, and Shitikanth Kashyap for proof-reading an earlier draft. I would also like to thank the anonymous reviewers of CRYPTO 2014 for various suggestions that helped in improving the presentation of this paper. I am extremely thankful to the anonymous reviewers of SODA 2014 for various references (like, Munro-Paterson [48], Flajolet-Martin [25], and Kane-Nelson [38]) and specifically pointing the improvements for utility bounds for linear regression through Kane and Nelson [38].

References

- [1] Nir Ailon and Bernard Chazelle. The Fast Johnson–Lindenstrauss Transform and Approximate Nearest Neighbors. *SIAM J. Comput.*, 39(1):302–322, 2009. 6
- [2] Nir Ailon and Edo Liberty. An Almost Optimal Unrestricted Fast Johnson-Lindenstrauss Transform. *ACM Transactions on Algorithms*, 9(3):21, 2013. 24
- [3] Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 20–29. ACM, 1996. 2, 8, 10
- [4] Amos Beimel, Kobbi Nissim, and Eran Omri. Distributed private data analysis: Simultaneously solving how and what. In *Advances in Cryptology—CRYPTO 2008*, pages 451–468. Springer, 2008. 5

- [5] Lucien Birgé and Pascal Massart. *From Model Selection to Adaptive Estimation*, pages 55–87. Springer New York, 1997.
- [6] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. The Johnson-Lindenstrauss Transform Itself Preserves Differential Privacy. In *FOCS*, pages 410–419, 2012. 2, 5, 7, 11, 14, 23, 24, 29, 30
- [7] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. Differentially private data analysis of social networks via restricted sensitivity. In Robert D. Kleinberg, editor, *ITCS*, pages 87–96. ACM, 2013. 2, 5, 7
- [8] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: the sulq framework. In Chen Li, editor, *PODS*, pages 128–138. ACM, 2005. 8
- [9] Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to noninteractive database privacy. *J. ACM*, 60(2):12, 2013. 7
- [10] E Bura and R Pfeiffer. On the distribution of the left singular vectors of a random matrix and its applications. *Statistics & Probability Letters*, 78(15):2275–2280, 2008. 19
- [11] Emmanuel J. Candès and Terence Tao. Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies? *IEEE Transactions on Information Theory*, 52(12):5406–5425, 2006. 6
- [12] Lynn E Cannon. A cellular computer to implement the kalman filter algorithm. Technical report, DTIC Document, 1969. 5
- [13] Kamalika Chaudhuri, Anand D Sarwate, and Kaushik Sinha. Near-optimal differentially private principal components. In *NIPS*, pages 998–1006, 2012. 8
- [14] Kenneth L. Clarkson and David P. Woodruff. Numerical linear algebra in the streaming model. In *STOC*, pages 205–214, 2009. 2, 3, 4, 5, 6, 7, 16, 17, 19, 29
- [15] Angus Deaton. *Understanding consumption*. Oxford University Press, 1992. 3
- [16] Scott Deerwester. Improving information retrieval with latent semantic indexing. 1988. 4
- [17] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *EUROCRYPT*, pages 486–503, 2006. 7
- [18] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating Noise to Sensitivity in Private Data Analysis. In Shai Halevi and Tal Rabin, editors, *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006. 7
- [19] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N. Rothblum. Differential privacy under continual observation. In Schulman [57], pages 715–724. 2, 7, 24
- [20] Cynthia Dwork, Moni Naor, Toniann Pitassi, Guy N Rothblum, and Sergey Yekhanin. Pan-Private Streaming Algorithms. In *ICS*, pages 66–80, 2010. 2, 8
- [21] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and Differential Privacy. In *FOCS*, pages 51–60. IEEE Computer Society, 2010. 7, 9

- [22] Cynthia Dwork, Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Analyze Gauss: Optimal Bounds for Privacy-Preserving Principal Component Analysis. In *STOC*, pages 11–20, 2014. 3, 4, 5, 7, 8, 24
- [23] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936. 4, 22
- [24] Ronald G Ehrenberg and Robert S Smith. Modern labor economics. 2010. 3
- [25] Philippe Flajolet and G Nigel Martin. Probabilistic counting algorithms for data base applications. *Journal of computer and system sciences*, 31(2):182–209, 1985. 2, 25
- [26] Lars Grasedyck and Wolfgang Hackbusch. Construction and arithmetics of H-matrices. *Computing*, 70(4):295–334, 2003. 4
- [27] Leslie Greengard and Vladimir Rokhlin. A new version of the fast multipole method for the laplace equation in three dimensions. *Acta numerica*, 6:229–269, 1997. 4
- [28] Anupam Gupta, Moritz Hardt, Aaron Roth, and Jonathan Ullman. Privately Releasing Conjunctions and the Statistical Query Barrier. *SIAM J. Comput.*, 42(4):1494–1520, 2013. 7
- [29] Anupam Gupta, Aaron Roth, and Jonathan Ullman. Iterative Constructions and Private Data Release. In *TCC*, pages 339–356, 2012. 7
- [30] Arjun K Gupta and Daya K Nagar. *Matrix variate distributions*, volume 104. CRC Press, 1999. 9
- [31] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011. 4, 7, 18, 19, 20, 21, 23
- [32] Moritz Hardt. Robust subspace iteration and privacy-preserving spectral analysis. In *Allerton*, pages 1624–1626. IEEE, 2013. 2, 5, 7, 24
- [33] Moritz Hardt, Katrina Ligett, and Frank McSherry. A Simple and Practical Algorithm for Differentially Private Data Release. In Peter L. Bartlett, Fernando C. N. Pereira, Christopher J. C. Burges, Léon Bottou, and Kilian Q. Weinberger, editors, *NIPS*, pages 2348–2356, 2012. 7
- [34] Moritz Hardt and Aaron Roth. Beating randomized response on incoherent matrices. In *STOC*, pages 1255–1268, 2012. 2, 3, 4, 5, 7, 8, 15, 18, 20, 23, 24
- [35] Moritz Hardt and Aaron Roth. Beyond worst-case analysis in private singular vector computation. In *STOC*, pages 331–340, 2013. 2, 3, 5, 7, 8, 18, 23, 24
- [36] Trevor Hastie, Robert Tibshirani, Jerome Friedman, and James Franklin. The elements of statistical learning: data mining, inference and prediction. *The Mathematical Intelligencer*, 27(2):83–85, 2005. 4
- [37] Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. Differentially Private Online Learning. In Shie Mannor, Nathan Srebro, and Robert C. Williamson, editors, *COLT*, volume 23 of *JMLR Proceedings*, pages 24.1–24.34. JMLR.org, 2012. 8, 17, 24
- [38] Daniel M. Kane and Jelani Nelson. Sparser Johnson-Lindenstrauss Transforms. *J. ACM*, 61(1):4, 2014. 17, 24, 25

- [39] Michael Kapralov and Kunal Talwar. On differentially private low rank approximation. In *SODA*, volume 5, page 1. SIAM, 2013. 2, 3, 5, 7, 8, 18, 23
- [40] Krishnaram Kenthapadi, Aleksandra Korolova, Ilya Mironov, and Nina Mishra. Privacy via the johnson-lindenstrauss transform. *arXiv preprint arXiv:1204.2606*, 2012. 7
- [41] Erricos John Kontoghiorghes. *Handbook of parallel computing and statistics*. CRC Press, 2010. 5
- [42] Felix Krahmer and Rachel Ward. New and Improved Johnson-Lindenstrauss Embeddings via the Restricted Isometry Property. *SIAM J. Math. Analysis*, 43(3):1269–1281, 2011. 6, 24
- [43] Pascal Massart and Jean Picard. *Concentration inequalities and model selection*, volume 1896. Springer, 2007.
- [44] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 94–103. IEEE, 2007. 7
- [45] Gerald M Meier et al. The international economics of development. Theory and policy. *The international economics of development. Theory and policy.*, 1968. 3
- [46] Leon Mirsky. Symmetric gauge functions and unitarily invariant norms. *The quarterly journal of mathematics*, 11(1):50–59, 1960. 4, 22
- [47] Robb J Muirhead. *Aspects of multivariate statistical theory*, volume 197. John Wiley & Sons, 2009. 10, 12, 22
- [48] J Ian Munro and Mike S Paterson. Selection and sorting with limited storage. *Theoretical computer science*, 12(3):315–323, 1980. 2, 25
- [49] Gilles Pisier. Some applications of the metric entropy condition to harmonic analysis. In *Banach Spaces, Harmonic Analysis, and Probability Theory*, pages 123–154. Springer, 1983.
- [50] C Radhakrishna Rao. *Linear statistical inference and its applications*, volume 22. John Wiley & Sons, 2009. 12, 13
- [51] Holger Rauhut, Justin K. Romberg, and Joel A. Tropp. Restricted Isometries for Partial Random Circulant Matrices. *CoRR*, abs/1010.1847, 2010. 6
- [52] Radim Rehurek. Subspace tracking for latent semantic analysis. In *Advances in Information Retrieval*, pages 289–300. Springer, 2011. 11, 16, 24
- [53] Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In Schulman [57], pages 765–774. 7
- [54] Mark Rudelson and Roman Vershynin. On sparse reconstruction from Fourier and Gaussian measurements. *Communications on Pure and Applied Mathematics*, 61(8):1025–1045, 2008. 6, 15
- [55] Youcef Saad and Martin H Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on scientific and statistical computing*, 7(3):856–869, 1986. 5

- [56] Tamas Sarlos. Improved approximation algorithms for large matrices via random projections. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 143–152. IEEE, 2006. 2, 3, 4, 5, 7, 16, 17, 19
- [57] Leonard J. Schulman, editor. *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*. ACM, 2010. 26, 28
- [58] Volker Strumpfen, Henry Hoffmann, and Anant Agarwal. A Stream Algorithm for the SVD. 2003. 11, 16, 24
- [59] Terence Tao. *Topics in random matrix theory*, volume 132. American Mathematical Soc., 2012. 10
- [60] Jalaj Upadhyay. Random Projections, Graph Sparsification, and Differential Privacy. In *ASIACRYPT (1)*, pages 276–295, 2013. 2, 5, 7, 11, 15
- [61] Jalaj Upadhyay. Circulant matrices and differential privacy. *IACR Cryptology ePrint Archive*, 2014:818, 2014. 6, 15
- [62] Jan Vybíral. A variant of the Johnson–Lindenstrauss lemma for circulant matrices. *Journal of Functional Analysis*, 260(4):1096–1105, 2 2011. 6, 15

A Utily Proof of Section 5.2

We give the formal proof of Theorem 5.3, which is similar to Clarkson and Woodruff [14] modulo the analysis to consider the lift of singular value. We include it for the sake of completion. Let $\tilde{A} = U\tilde{\Sigma}V^\top$. Since the columns of U is a set of orthonormal vectors, we have $UU^\top U = U$ and $\|U^\top UC\|_F = \|UC\|_F$ for any matrix C . Therefore, it suffices for the utility bound to prove a bound on $\|U^\top \tilde{A}(\tilde{X} - X^*)\|$. For this, we first prove that $U^\top \Omega^\top \Omega \tilde{A}(\tilde{X} - X^*)$ has a small norm. We have

$$\begin{aligned} U^\top \Omega^\top \Omega \tilde{A}(\tilde{X} - X^*) &= U^\top \Omega^\top \Omega \tilde{A}(\tilde{X} - X^*) + U^\top \Omega^\top \Omega (B - \tilde{A}\tilde{X}) \\ &= U^\top \Omega^\top \Omega (B - \tilde{A}X^*). \end{aligned}$$

This is because $U^\top \Omega^\top \Omega (\tilde{A}\tilde{X} - B) = \tilde{A}^\top \Omega^\top \Omega (\tilde{A}\tilde{X} - B) = 0$. Therefore, from Theorem 5.1 with $\alpha' = \sqrt{\alpha/d}$ (since we chose r in Theorem 5.1 which differs by a factor of α and $1/d$ with respect to that in Theorem 5.3), we have

$$\begin{aligned} \|U^\top \Omega^\top \Omega \tilde{A}(\tilde{X} - X^*)\|_F &= \|U^\top \Omega^\top \Omega (B - \tilde{A}X^*)\|_F \\ &\leq \sqrt{\alpha} \|B - AX^*\|_F + \sqrt{\tau} \end{aligned}$$

From the sub-additivity of the norm and property of conforming matrices, we have

$$\begin{aligned} \|U^\top A(\tilde{X} - X^*)\|_F &\leq \|U^\top \Omega^\top \Omega \tilde{A}(\tilde{X} - X^*)\|_F + \|U^\top \Omega^\top \Omega \tilde{A}(\tilde{X} - X^*) - U^\top \tilde{A}(\tilde{X} - X^*)\|_F \\ &\leq \sqrt{\alpha} \|B - AX^*\|_F + \sqrt{\tau} + \|U^\top \Omega^\top \Omega U - \mathbb{I}\|_2 \cdot \|U^\top \tilde{A}(\tilde{X} - X^*)\|_F \end{aligned}$$

Using Lemma 5.4 and rearranging the terms, we get $\|U^\top \tilde{A}(\tilde{X} - X^*)\|_F \leq 2\sqrt{\alpha} \|B - AX^*\|_F + \sqrt{\tau}$. The utility proof is now immediate by observing that the column-space of A and U are the same, and the Pythagorus theorem on the norms. More concretely, with probability at least $1 - 2\beta$

$$\begin{aligned} \|A\tilde{X} - B\|_F^2 &= \|AX^* - B\|_F^2 + \|A(\tilde{X} - X^*)\|_F^2 \\ &\leq (1 + 4\alpha) \|AX^* + B\|_F + \tau. \end{aligned}$$

Adjusting and renaming the values of α and β , we get the claim of the theorem.

B Differential Privacy of Variant 1

We give the proof Blocki *et al.* [6] for the sake of completion. We denote by \bar{A} the matrix that differs from A by at most one entry. In other word, if A and \bar{A} differs in row i , then there exists a unit vector v such that $A - \bar{A} = E = ve_i^T$. Let $U\Sigma V^T$ ($\bar{U}\bar{\Lambda}\bar{V}^T$, respectively) be the SVD of A (\bar{A} , respectively).

The PDF for the two distributions, corresponding to A and \bar{A} , is just a linear transformation of $\mathcal{N}(0, \mathbb{I}_{n \times n})$. Therefore,

$$\begin{aligned} \text{PDF}_{A^T Y}(\mathbf{x}) &= \frac{1}{\sqrt{(2\pi)^d \Delta(A^T A)}} \exp\left(-\frac{1}{2} \langle \mathbf{x} | (A^T A)^{-1} | \mathbf{x} \rangle\right) \\ \text{PDF}_{\bar{A}^T Y}(\mathbf{x}) &= \frac{1}{\sqrt{(2\pi)^d \Delta(\bar{A}^T \bar{A})}} \exp\left(-\frac{1}{2} \langle \mathbf{x} | (\bar{A}^T \bar{A})^{-1} | \mathbf{x} \rangle\right) \end{aligned}$$

We prove the result for a row of the published matrix; the theorem follows from Theorem 2.1. It is straightforward to see that combination of the following proves differential privacy for a row of published matrix:

$$\sqrt{\frac{\tilde{\Delta}(A^T A)}{\tilde{\Delta}(\bar{A}^T \bar{A})}} \in \exp(\pm \varepsilon_0) \quad \text{and} \quad \Pr \left[\left| \langle \mathbf{x} | \Omega^T (A^T A)^{-1} \Omega | \mathbf{x} \rangle - \langle \mathbf{x} | \Omega^T (\bar{A}^T \bar{A})^{-1} \Omega | \mathbf{x} \rangle \right| \leq \varepsilon_0 \right] \geq 1 - \delta_0,$$

where

$$\varepsilon_0 = \frac{\varepsilon}{\sqrt{4r \ln(2/\delta)}} \quad \delta_0 = \frac{\delta}{2r}.$$

The first part of the proof follows simply as in Blocki *et al.* [6]. More concretely, we have $\tilde{\Delta}(A^T A) = \prod_i \sigma_i^2$, where $\sigma_1 \geq \dots \geq \sigma_d \geq \sigma_{\min}$ are the singular values of A . Let $\lambda_1 \geq \dots \geq \lambda_d \geq \sigma_{\min}$ be its singular value for \bar{A} . Since the singular values of $A - \bar{A}$ and $\bar{A} - A$ are the same, $\sum_i (\sigma_i - \lambda_i) \leq 1$ using Linskii's theorem. Therefore,

$$\sqrt{\prod_i \frac{\lambda_i^2}{\sigma_i^2}} \leq \exp \left(\frac{\varepsilon}{32\sqrt{r \log(2/\delta) \log(r/\delta)}} \right) \sum_i (\lambda_i - \sigma_i) \leq e^{\varepsilon_0/2}.$$

The second part of the proof is slightly more involved. Each row i of the published matrix is distributed identically and is constructed by multiplying an n -dimensional vector Ω_i that has entries picked from a normal distribution $\mathcal{N}(0, 1)$. Note that $\mathbb{E}[\Omega_i] = 0$ and $\text{COV}(\Omega_i) = \mathbb{I}$. Let $\mathbf{y} = \Pi_i(\alpha)$. Then

$$\langle \mathbf{x} | \Omega^T (A^T A)^{-1} \Omega | \mathbf{x} \rangle - \langle \mathbf{x} | \Omega^T (\bar{A}^T \bar{A})^{-1} \Omega | \mathbf{x} \rangle = \langle \mathbf{x} | \Omega^T \left[(A^T A)^{-1} (A^T E + E^T \bar{A}) (\bar{A}^T \bar{A})^{-1} \right] \Omega | \mathbf{x} \rangle.$$

Using the singular value decomposition of $A = U\Sigma U^T$ and $\bar{A} = \bar{U}\bar{\Lambda}\bar{V}^T$, this simplifies as

$$\left[\langle \mathbf{x} | \Omega^T (V\Sigma^{-1}U^T)e_i \right] \left[v^T (\bar{V}\bar{\Lambda}^{-2}\bar{V}^T)\Omega | \mathbf{x} \right] + \langle \mathbf{x} | \Omega^T \left[(V\Sigma^{-2}V^T)v \right] \left[e_i^T (\bar{U}\bar{\Lambda}^{-1}\bar{V}^T)\Omega | \mathbf{x} \right].$$

Since $x \sim A^T \mathbf{y}$, where $\mathbf{y} \sim \mathcal{N}(0, 1)$, we can further simplify it as

$$\underbrace{\left[\mathbf{y}^T A \Omega^T (V\Sigma^{-1}U^T)e_i \right]}_{t_1} \underbrace{\left[v^T (\bar{V}\bar{\Lambda}^{-2}\bar{V}^T)\Omega A^T \mathbf{y} \right]}_{t_2} + \underbrace{\left[\mathbf{y}^T A \Omega^T (V\Sigma^{-2}V^T)v \right]}_{t_3} \underbrace{\left[e_i^T (\bar{U}\bar{\Lambda}^{-1}\bar{V}^T)\Omega A^T \mathbf{y} \right]}_{t_4}.$$

Now since $\|\Lambda\|_2, \|\Sigma\|_2 \geq w$, plugging in the SVD of A and $A - A' = ve_i^\top$, and that every term t_i in the above expression is a linear combination of a Gaussian, i.e., each term is distributed as per $\mathcal{N}(0, \|t_i\|^2)$, we calculate $\|t_i\|$ as below.

$$\begin{aligned} \|(U\Sigma V^\top)\Omega^\top(V\Sigma^{-1}U^\top)e_i\|_2 &\leq \|\Omega^\top\|_2 \leq 1, & \|(U\Sigma V^\top)\Omega^\top(V\Sigma^{-2}V^\top)v\|_2 &\leq \|\Omega^\top\|_2\|\Sigma^{-1}\|_2 \leq \frac{1}{\sigma_{\min}}, \\ \|v^\top(\bar{V}\Lambda^{-2}\bar{V}^\top)\Omega(\bar{V}\Lambda\bar{U}^\top - ve_i^\top)\|_2 &\leq \|v^\top(\bar{V}\Lambda^{-2}\bar{V}^\top)\Omega\bar{V}\Lambda\bar{U}^\top\|_2 + \|v^\top(\bar{V}\Lambda^{-2}\bar{V}^\top)\Omega ve_i^\top\|_2 \leq \frac{1}{\sigma_{\min}} + \frac{1}{\sigma_{\min}^2}, \\ \|e_i^\top(\bar{U}\Lambda^{-1}\bar{V}^\top)\Omega(\bar{V}\Lambda\bar{U}^\top - ve_i^\top)\|_2 &\leq \|e_i^\top(\bar{U}\Lambda^{-1}\bar{V}^\top)\Omega\bar{V}\Lambda\bar{U}^\top\|_2 + \|e_i^\top(\bar{U}\Lambda^{-1}\bar{V}^\top)\Omega ve_i^\top\|_2 \leq 1 + \frac{1}{\sigma_{\min}}, \end{aligned}$$

where $\sigma_{\min} = \left(\frac{\sqrt{r \log(2/\delta)} \log(r/\delta)}{\varepsilon}\right)$. Using the concentration bound on the Gaussian distribution, each term,

t_1, t_2, t_3 , and t_4 , is less than $\|t_i\| \ln(4/\delta_0)$ with probability $1 - \delta_0/2$. From the fact that $2\left(\frac{1}{\sigma_{\min}} + \frac{1}{\sigma_{\min}^2}\right) \ln(4/\delta_0) \leq \varepsilon_0$, we have the second part of the proof.