

Privacy-Utility Trade-off of Linear Regression under Random Projections and Additive Noise

Mehrdad Showkatbakhsh
UCLA, Los Angeles, CA
mehrdadsh@ucla.edu

Can Karakus
UCLA, Los Angeles, CA
karakus@ucla.edu

Suhas Diggavi
UCLA, Los Angeles, CA
suhasdiggavi@ucla.edu

Abstract—Data privacy is an important concern in machine learning, and is fundamentally at odds with the task of training useful learning models, which typically require acquisition of large amounts of private user data. One possible way of fulfilling the machine learning task while preserving user privacy is to train the model on a transformed, noisy version of the data, which does not reveal the data itself directly to the training procedure. In this work, we analyze the privacy-utility trade-off of two such schemes for the problem of linear regression: additive noise, and random projections. In contrast to previous work, we consider a recently proposed notion of differential privacy that is based on conditional mutual information (MI-DP), which is stronger than the conventional (ϵ, δ) -differential privacy, and use relative objective error as the utility metric. We find that projecting the data to a lower-dimensional subspace before adding noise attains a better trade-off in general. We also make a connection between privacy problem and (non-coherent) SIMO, which has been extensively studied in wireless communication, and use tools from there for the analysis. We present numerical results demonstrating the performance of the schemes.

I. INTRODUCTION

High-complexity models are needed to solve modern learning problems, which require large amounts of data to achieve low generalization error. However, acquiring such data from users directly compromises the user privacy. Training useful machine learning models without compromising user privacy is an important and challenging research problem. One natural way to tackle this problem is to keep the data itself private, and reveal only a processed, noisy version of the data to the training procedure. Ideally, such processing would completely hide the content of the data samples, while still providing useful information to the training objective. In this paper, we analyze the privacy-utility trade-off of two such schemes for the linear regression problem: additive noise, where training is performed on the data samples with additive Gaussian noise; and random projections, where each data sample is randomly projected to a lower-dimensional subspace through Johnson-Lindenstrauss Transform (JLT) [1] before adding Gaussian

noise. We explore guarantees for a model that is trained on such transformed data for a given privacy constraint.

Differential privacy is perhaps the most well-known notion for privacy [2], and has been applied to a variety of domains (we refer reader to [3] and [2] and references therein). It assumes a strong adversary which has access to all data samples except one, thereby ensuring robustness of the privacy guarantee to adversaries with side-information about the database. Moreover differential privacy makes no distributional assumption on the data.

In this work we use the recently proposed notion of mutual information-differential privacy (MI-DP) to analyze the privacy performance of the schemes. This connects to the natural information-theoretic notion of privacy, as well as enabling the use of more standard tools for analysis. Moreover it is shown in [4] that MI-DP directly implies (ϵ, δ) -differential privacy.

Our contributions are as follows: First, we derive closed-form expressions on the relative objective error achievable by additive noise (Theorem 1) and random projection schemes (Theorem 2), under a privacy constraint, and show that in general random projections achieve better privacy-utility trade-off. We use results from randomized linear algebra [5] to prove the utility guarantees. Second, using the MI-DP measure, and using the fact that the random projection matrix is private, we make a connection between the MI-DP and SIMO channel, and show that non-coherent SIMO bounds do not give a stronger scaling guarantee than their coherent counterparts. Third, we present numerical results demonstrating the performance of the two schemes.

Related work. The works in [6]–[8] propose perturbing the objective to provide privacy guarantees on the trained model, where the training procedure is trusted and has access to the full database, and the adversary can only access the resulting trained model. In contrast, we assume that the training procedure itself may be adversarial, and is not given access to the raw data samples. In the context of linear regression and related problems, the works in [9] and [5] propose random projections to provide privacy, by showing that the mutual information between the raw and projected data samples grows sublinearly with dimensions. However, this does not necessarily translate to a formal differential privacy guarantee on the data samples. Random projection as a tool to provide differential privacy has also been considered in [10] and [11]. The main difference of these works with ours is that they

The work of the authors was partially supported by the NSF award 1740047, by the Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196, and by the UC-NL grant LFR-18-548554. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

project each data vector individually to a lower-dimensional subspace, whereas we consider mixing samples across the database, such that the effective number of “mixed” samples is fewer than original.

In terms of motivation and techniques, the works in [12], [13] are the most closely related to ours. These works consider JLT in the context of linear regression, and prove that it guarantees differential privacy for well-conditioned data matrices. However, no explicit guarantee on the achievable empirical risk is given. In contrast, we directly analyze the privacy-utility trade-off of additive noise and random projections, where utility is measured by the objective value achieved by the trained model under the privacy scheme, normalized by the true minimum of the objective. We also use the stronger MI-DP privacy¹, instead of the traditional (ϵ, δ) -differential privacy. We emphasize that the main novelty of our work lies in the analysis of the algorithms and the resulting theoretical guarantees, and not in the algorithms themselves.

Paper organization. In Section II we give a brief overview on different privacy metrics followed by the precise problem formulation. Section III includes the main theoretical results of this work. The proof outlines are given in Section IV. Section V gives the numerical results.

II. FORMULATION AND BACKGROUND

In this paper we consider the quadratic optimization

$$\min_{\theta} g(\theta) := \min_{\theta} \|X\theta - y\|_2^2, \quad (1)$$

where $X \in \mathbb{R}^{n \times d}$ is the data matrix that each row corresponds to one user and $y \in \mathbb{R}^n$ are the response variables. We denote a solution of this optimization problem as θ^* . We use $X_{i,j}$ to denote the j -th feature of the i -th user data point for $i \in \{1, \dots, n\}, j \in \{1, \dots, d\}$. We assume the number of data points is greater than the number of features and X is full column rank. We assume that $|X_{i,j}| \leq 1$. Throughout this paper, we use bold letters for random variables to distinguish them from deterministic quantities.

Consider a database $D^N := (D_1, \dots, D_N)$ that returns a query according to a *randomized* mechanism $q(\cdot)$. Let D^{-i} denote the set of database entries excluding D_i .

Definition 1 (ϵ -mutual-information). *A randomized mechanism $q(\cdot)$ satisfies ϵ -mutual-information (MI-DP) if*

$$\sup_{i, P(D^N)} I(\mathbf{D}_i; q(D^N) | D^{-i}) \leq \epsilon \text{ bits}, \quad (2)$$

where the supremum is taken over all distribution on D^N .

We aim to preserve the privacy of each entry of X , therefore, in the context of our work, $D^N := (X_{1,1}, \dots, X_{1,d}, X_{2,1}, \dots, X_{2,d}, \dots, X_{n,d})$.

The notion of ϵ -MI-DP is closely related to (ϵ, δ) differential privacy [14]. We first define the notion of neighbor in databases:

¹In [4], it is shown that for discrete alphabets, the two notions are equivalent; however MI-DP is strictly stronger for continuous alphabets.

Definition 2 (Neighbor). *Two databases D^N and \bar{D}^N are called neighbor if they differ only in one entry.*

In the context of our problem, two data matrices are neighbors if they only differ in one entry. Now we are ready to define (ϵ, δ) differential privacy.

Definition 3 ((ϵ, δ) differential privacy). *A randomized mechanism $q(\cdot)$ satisfies (ϵ, δ) differential privacy (DP) if for all neighboring databases D^N and \bar{D}^N and all $S \subseteq \text{Range}(q(\cdot))$,*

$$\Pr(q(D^N) \in S) \leq e^\epsilon \Pr(q(\bar{D}^N) \in S) + \delta. \quad (3)$$

We say $q(\cdot)$ satisfies (δ) -DP if it satisfies $(0, \delta)$ -differential privacy.

Note that neither of MI-DP nor DP impose distributional assumptions on the database and the probabilities arise completely from the randomization of the mechanism.

Proposition 1 (Theorem 1 in [4]). *ϵ -MI-DP is stronger than (ϵ, δ) -DP in the sense that for all $\epsilon > 0$ if a mechanism is ϵ -MI-DP, there exists ϵ', δ' such that the mechanism satisfies (ϵ', δ') -DP. We denote this relation with ϵ -MI-DP $\succeq (\epsilon, \delta)$ -DP. Furthermore, we have the following relation:*

$$\epsilon\text{-MI-DP} \stackrel{(a)}{\succeq} (\delta)\text{-DP} \stackrel{(b)}{\equiv} (\epsilon, \delta)\text{-DP}, \quad (4)$$

where \succeq is interpreted as being stronger and (b) means (δ) -DP $\succeq (\epsilon, \delta)$ -DP and (ϵ, δ) -DP $\succeq (\delta)$ -DP.

Proposition 2 (See Lemma 2 in [4]). *If a mechanism is ϵ -MI-DP then it also satisfies $(0, \sqrt{\frac{2}{\log(e)}\epsilon})$ -DP.*

Let us denote a solution to the original problem (1) with θ^* . Let us denote the the cost function of the transformed problem with $\hat{g}(\theta)$ with a minimum of $\hat{\theta} \in \text{argmin}_{\theta} \hat{g}(\theta)$. We define the relative error of this transformed problem as the smallest $\eta \geq 1$ such that,

$$g(\hat{\theta}) \leq \eta g(\theta^*). \quad (5)$$

In this paper we consider the achievable relative error for linear regression given ϵ -MI-DP requirement.

Notation. We denote the *condition number* of X with

$$\kappa(X) := \|X\|_2 \|X^\dagger\|_2 = \frac{\sigma_{\max}(X)}{\sigma_{\min}(X)}, \quad (6)$$

where X^\dagger is the Moore-Penrose pseudoinverse of X and $\|X\|_2$ is the spectral norm of X .

We denote that ratio of l_2 norm of the projection of y onto the column space of X over the l_2 norm of the residual with:

$$r(y) := \frac{\|X\theta^*\|_2}{\|X\theta^* - y\|_2}. \quad (7)$$

where $\|X\theta^*\|_2$ is the l_2 norm of the vector $X\theta^*$.

We define $f_j(X) := \sqrt{\sum_{i=1}^n |X_{i,j}^2| - \max_i |X_{i,j}^2|}$ for $j \in \{1, \dots, d\}$, and $f(X) := \min_j f_j(X)$. In order to give guarantees on the privacy of the projection method the amount of additional noise is expressed in terms of $f(X)$.

III. PRIVACY-UTILITY TRADE OFF

In this section we analyze the utility-privacy trade-off for both an additive noise mechanism as well as a scheme with random projection. We compare their utility-guarantees for the same level of ε -MI-DP privacy.

A. Additive Gaussian Noise

In order to satisfy privacy, we add Gaussian noise directly to the data,

$$X_{AN}(X) := X + \sigma_{AN}\mathbf{N}, \quad (8)$$

where $\mathbf{N} \in \mathbb{R}^{n \times d}$ with i.i.d. entries drawn from $\mathcal{N}(0, 1)$ and

$$\sigma_{AN}^2(\varepsilon) := \frac{1}{2^{2\varepsilon} - 1}, \quad (9)$$

is the variance of the noise.

$$\theta_{AN} := \operatorname{argmin}_{\theta} \underbrace{\|X_{AN}\theta - y\|_2^2}_{g_{AN}(\theta)}. \quad (10)$$

Theorem 1 (Privacy-Utility for Additive Noise). *Given a data set X and the randomized mechanism $X_{AN}(X)$ with ε -MI-DP constraint, with probability at least $1 - 2e^{-\frac{\sigma_{\max}(X)^2}{2\sigma_{AN}^2(\varepsilon)}\delta^2}$ we have the following bound on the relative error of the transformed problem:*

$$\eta_{AN} \leq \left(1 + \frac{\kappa(X)(\Delta(X, \varepsilon) + \delta)}{1 - \kappa(X)(\Delta(X, \varepsilon) + \delta)}(\kappa(X) + r(y))\right)^2, \quad (11)$$

if $\kappa(X)\Delta(\varepsilon, X) < 1$, where $\Delta(X, \varepsilon) = \frac{\sigma_{AN}(\varepsilon)}{\sigma_{\max}(X)}(\sqrt{n} + \sqrt{d})$ and $\delta > 0$ is a free parameter².

Note that if $\sigma_{\max}^2(X)$ scales linearly with n then Δ converges to a constant term. Based on Proposition 2, additive noise also satisfies (δ) -DP.

B. Gaussian Random Projections

We encode the data matrix using JLT to a lower dimensional space n' and we add Gaussian noise, when necessary, to guarantee ε -MI-DP. We denote the encoded data by $X_{RP} \in \mathbb{R}^{n' \times d}$ and $y_{RP} \in \mathbb{R}^{n'}$:

$$X_{RP}(X) := \mathbf{S}X + \sigma_{RP}\mathbf{N}, \quad y_{RP} := \mathbf{S}y, \quad (12)$$

where $\mathbf{S} \in \mathbb{R}^{n' \times n}$ represents the random projection with i.i.d. $\mathcal{N}(0, 1)$ entries and $\mathbf{N} \in \mathbb{R}^{n' \times d}$ is the noise added to ensure the privacy with i.i.d. entries drawn from $\mathcal{N}(0, 1)$, and

$$\sigma_{RP}^2(X, \varepsilon) := \left(\frac{n'}{2^{2\varepsilon} - 1} - f^2(X)\right)_+, \quad (13)$$

is the variance of the additive noise³.

We solve the following problem to estimate the model:

$$\theta_{RP} := \operatorname{argmin}_{\theta} \underbrace{\|X_{RP}\theta - y_{RP}\|_2^2}_{g_{RP}(\theta)}, \quad (14)$$

²Note that support of δ is restricted to the set where $\kappa(X)(\Delta + \delta) \leq 1$.

³Note that our algorithm does not reveal this quantity explicitly avoiding an extra privacy leakage.

Theorem 2 (Privacy-Utility for Random Projection). *Given a dataset X and the randomized mechanism $X_{RP}(X)$ with ε -MI-DP constraint and a projection dimension of $n' < n$, with probability at least $1 - c_1 e^{-c_2 n' \delta^2}$, we have the following bound on the relative error of the transformed problem:*

$$\eta_{RP} \leq (1 + \delta)^2 (1 + l_1(X, \varepsilon))(1 + l_2(X, \varepsilon))^2, \quad (15)$$

where $l_1(X, \varepsilon) := \sigma_{RP}^2(X, \varepsilon) \left(\max_i \frac{\sigma_i(X)}{\sigma_i^2(X) + \sigma_{RP}^2}\right)^2 r^2(y)$, $l_2(X, \varepsilon) := \frac{\sigma_{RP}^2(X, \varepsilon)}{\sigma_{\min}^2(X) + \sigma_{RP}^2(X, \varepsilon)} r(y)$, $\delta \geq \sqrt{c_0 \frac{d}{n'}}$ is a free parameter, and c_0, c_1 and c_2 are constants.

Corollary 1. *The random projection methods also satisfies (δ) -DP for $\delta := \sqrt{\frac{2}{\log(e)}}\varepsilon$.*

Corollary 2. *Note that the amount of noise added to the projected data is $\sigma_{RP}^2(\varepsilon, X) = \left(\frac{n'}{2^{2\varepsilon} - 1} - f^2(X)\right)_+$. If f^2 scales linearly with n and $n' = o(n)$, asymptotically the noise variance goes to zero, i.e., random projection itself guarantees the privacy. Furthermore, for a given δ , $\eta_{RP} \leq (1 + \delta)^2$ asymptotically as two other terms in (15) vanish.*

Remark 1. *In the proof of Theorem 2 in order to derive an upper bound for (2) we make a connection to the SIMO non-coherent channel. We used the coherent SIMO bound for upper bounding this quantity. One may ask if we can get a tighter bound by using the tighter non-coherent bounds (see for example [15]). The known non coherent bound,*

$$C \leq \frac{n'}{2} \log\left(1 + \frac{1}{\sigma_{RP}^2 + f^2(X)}\right). \quad (16)$$

does not give any improvement. This bound (16) is known to be tight for the low-SNR regime [15]. Therefore when $f^2 = \Omega(n)$ asymptotically both bounds yield the same result.

IV. PROOF OUTLINES

A. Theorem 1

Proof Outline. The proof consists of two steps. First we derive the minimum amount of noise needed to ensure ε -MI-DP for X_{AN} with respect to any feature of users, which is stated in the following lemma:

Lemma 1 (Privacy Guarantee for the additive noise). *If $\sigma_{AN}^2 = \frac{1}{2^{2\varepsilon} - 1}$ then X_{AN} is ε -MI-DP with respect to any entry of X .*

Proof. We show that (2) is bounded by ε for this choice of σ_{AN}^2 and $q(X) := X_{AN}$. Due to the symmetry of the problem, we fix \mathbf{D}_i to be the first feature of the first data point without loss of generality. Note that

$$I(\mathbf{X}_{1,1}; \mathbf{X}_{AN} | \mathbf{X}^{-(1,1)}) = I(\mathbf{X}_{1,1}; \mathbf{X}_{1,1} + \sigma_{AN}\mathbf{N}_{1,1} | \mathbf{X}^{-(1,1)}). \quad (17)$$

By expanding the mutual information:

$$\begin{aligned} I(\mathbf{X}_{1,1}; \mathbf{X}_{1,1} + \sigma_{AN}\mathbf{N}_{1,1} | \mathbf{X}^{-(1,1)}) \\ \stackrel{(a)}{=} h(\mathbf{X}_{1,1} + \sigma_{AN}\mathbf{N}_{1,1} | \mathbf{X}^{-(1,1)}) - h(\sigma_{AN}\mathbf{N}_{1,1}), \end{aligned} \quad (18)$$

where (a) holds because the noise is independent of the data. Now we need to take the maximization over all possible

distribution on \mathbf{X} . Note that the absolute value of each entry is bounded by 1 therefore we need to take the supremum over all distribution inside this ball. The absolute value constraint implies the second moment constraint for all distribution defined on it, therefore by using the maximum entropy bound the result follows:

$$\sup_{P(\mathbf{X})} I(\mathbf{X}_{1,1}; \mathbf{X}_{AN} | \mathbf{X}^{-(1,1)}) \leq \frac{1}{2} \log\left(1 + \frac{1}{\sigma_{AN}^2}\right) = \varepsilon. \quad (19)$$

□

The second step bounds the relative error. We use perturbation theory in the least square setup (see Theorem 5.1 in [16]) and probabilistic bounds on the maximum singular value of an i.i.d. Gaussian to derive the result [17]. The details of the proof are provided in [18]. □

B. Theorem 2

Proof Outline. The proof consists of two steps. First we find the variance of noise needed to add to satisfy ε -MI-DP that results to ε -MI-DP model, θ_{RP} . Following lemma characterizes the amount of noise sufficient to make the mechanism ε -MI-DP.

Lemma 2. *If $\sigma_{RP}^2 := (\frac{n'}{2^{2\varepsilon}-1} - f^2(X))_+$ then X_{RP} is ε -MI-DP with respect to any entry of X .*

Proof. We show that the conditional mutual information (2) is bounded by ε for this choice of σ_{RP}^2 . Due to the symmetry of the problem, we fix D_i to be the first feature of the first data point.

$$\begin{aligned} & \max_{P(\mathbf{X}) \in \mathbb{P}} I(\mathbf{X}_{1,1}; \mathbf{X}_{RP} | \mathbf{X}^{-(1,1)}) \quad (20) \\ &= \max_{P(\mathbf{X}^{-(1,1)}) \rho(\mathbf{X}_{1,1} | \mathbf{X}^{-(1,1)})} \mathbb{E}_{\mathbf{X}^{-(1,1)}} [I(\mathbf{X}_{1,1}; \mathbf{X}_{RP} | \mathbf{X}^{-(1,1)} = \mathbf{X}^{-(1,1)})], \\ &\stackrel{(a)}{\leq} \max_{P(\mathbf{X}^{-(1,1)})} \mathbb{E}_{\mathbf{X}^{-(1,1)}} \left[\max_{P(\mathbf{X}_{1,1} | \mathbf{X}^{-(1,1)})} I(\mathbf{X}_{1,1}; \mathbf{X}_{RP} | \mathbf{X}^{-(1,1)} = \mathbf{X}^{-(1,1)}) \right], \end{aligned}$$

where \mathbb{P} is the set of distributions which assign non-zero measure to \mathbf{X} only if the absolute value of each entry is upper bounded by 1 and $f(\mathbf{X})$ is lower bounded by the $f(\cdot)$ evaluated for the original database, (a) follows from the Jensen's Inequality and the fact that maximization over the conditional distribution is a convex function. Now we find upper bounds on the inside of the expectation, note that columns of \mathbf{X}_{RP} rather than first one does not have any term associated with $\mathbf{X}_{1,1}$ and they are conditionally independent, therefore we can write

$$\begin{aligned} & \max_{P(\mathbf{X}_{1,1} | \mathbf{X}^{-(1,1)})} I(\mathbf{X}_{1,1}; \mathbf{X}_{RP} | \mathbf{X}^{-(1,1)} = \mathbf{X}^{-(1,1)}) \\ &= \underbrace{\max_{P(\mathbf{X}_{1,1} | \mathbf{X}^{-(1,1)})} I(\mathbf{S}\mathbf{X}^{(:,1)} + \sigma_{RP}\mathbf{N}^{(:,1)}; \mathbf{X}_{1,1} | \mathbf{X}^{-(1,1)} = \mathbf{X}^{-(1,1)})}_{(\star)}, \end{aligned}$$

where $\mathbf{X}^{(:,1)}$ denotes the first column of X . We find an upper bound on (\star) for a fixed set of $\mathbf{X}^{-(1,1)}$. We observe that (\star) is same as the capacity of the following non-coherent SIMO channel with Rayleigh fading with a unit power constraint:

$$z = \underbrace{\mathbf{S}^{(:,1)} \mathbf{X}_{1,1} + \sum_{i \neq 1} \mathbf{S}^{(:,i)} X_{i,1} + \sigma_{RP} \mathbf{N}^{(:,1)}}_{\mathbf{v}}, \quad (21)$$

where $z \in \mathbb{R}^{n'}$ is the first column of \mathbf{X}_{RP} which we treat here as the output of the channel. Note that $X_{i,1}$ ($i \neq 1$) are treated as constants for this channel and therefore \mathbf{v} is effectively a zero mean i.i.d. Gaussian noise with the covariance of

$$\mathbb{E}[\mathbf{v}\mathbf{v}^T] = (\sigma_{RP}^2 + \sum_{i \neq 1} (X_{i,1})^2) \mathbf{I}_{n'} = \sigma_{\mathbf{v}}^2 \mathbf{I}_{n'}, \quad (22)$$

Now we bound the capacity of this channel, We use the coherent upper bound for the capacity of this channel:

$$\begin{aligned} \max_{P(\mathbf{X}_{1,1})} I(\mathbf{X}_{1,1}; z) &\leq \max_{P(\mathbf{X}_{1,1})} I(\mathbf{X}_{1,1}; z, \mathbf{S}^{(:,1)}) \\ &\stackrel{(a)}{\leq} \mathbb{E}_{\mathbf{S}^{(:,1)}} \left[\frac{1}{2} \log\left(1 + \frac{\|\mathbf{S}^{(:,1)}\|^2}{\sigma_{\mathbf{v}}^2}\right) \right] \\ &\stackrel{(b)}{\leq} \frac{1}{2} \log\left(1 + \mathbb{E}_{\mathbf{S}^{(:,1)}} \left[\frac{\|\mathbf{S}^{(:,1)}\|^2}{f(X)^2 + \sigma_{RP}^2} \right] \right) \stackrel{(c)}{\leq} \varepsilon. \quad (23) \end{aligned}$$

Note that the absolute value constraint implies the second moment constraint for all distribution defined on it and (a) follows from the maximum entropy bound, (b) follows directly from the Jensen's Inequality, (c) comes from the fact that the outer maximization is over distributions that assign non-zero measure to \mathbf{X} only if $f(\mathbf{X}) \geq f(X)$. □

Now we derive the utility guarantee for this method. Note that by rewriting $\tilde{X}_{PR} = [S \ N] \begin{bmatrix} X \\ \sigma_{RP} I \end{bmatrix} = \tilde{S} \begin{bmatrix} X \\ \sigma_{RP} I \end{bmatrix}$ we observe that adding direct noise to the projected data can be interpreted as the random projection of the l_2 regularized least square problem (Ridge Regression), i.e.,

$$\begin{aligned} \theta_{RP} &= \operatorname{argmin}_{\theta} \|\mathbf{X}_{RP}\theta - y_{RP}\|_2^2 \quad (24) \\ &= \operatorname{argmin}_{\theta} \|\underbrace{\tilde{S} \begin{bmatrix} X \\ \sigma_{RP} I \end{bmatrix}}_{\text{RR}} \theta - \begin{bmatrix} y \\ 0 \end{bmatrix}\|_2^2. \quad (25) \end{aligned}$$

Therefore we can split the utility analysis into two parts,

- 1) What is the utility loss for the l_2 regularized least square?
- 2) What is the utility loss for the randomized sketching (JLT)?

We use the standard SVD argument to bound the Ridge Regression relative error and by following Pilanci et. al. [5] (see Corollary 2) we give guarantees on the sketching step. The details of the proof are provided in [18]. □

V. NUMERICAL RESULTS

We numerically evaluate the relative error η achieved by the schemes in Section III subject to an ε -MI-DP constraint.

A. Random data

We generate the elements $X_{i,j}$ i.i.d. uniformly in the interval $[-1, 1]$, where $X \in \mathbb{R}^{n \times 800}$, and $n = 1000k$ with $k \in \{1, 2, \dots, 20\}$. For each case, the additive noise parameter σ_{AN} is computed according to (9). Similarly, the additive noise σ_{RP} is computed according to (13). Given k , we evaluate three choices of n' : logarithmic ($n'_1 := 1000(\log(k) + 1)$), linear ($n'_2 := 1000\frac{k+1}{2}$), and full ($n'_3 = n = 1000k$). The resulting relative error curves are given in Figure 1 for $\varepsilon = 0.5$, averaged




Fig. 1. Relative error of the schemes vs. n for $\epsilon = 0.5$, for random data




Fig. 2. Relative error of the schemes vs. ϵ , for $n = 10000$, for random data

over 5 trials. We note that random projection results in uniformly better privacy-utility trade-off compared to additive noise. Further, at this regime of ϵ , lower projection dimensions result in significantly better trade-off. Figure 2 plots the achieved relative error as a function of ϵ , for $n = 10000$. We note that the relative error decreases linearly until it saturates for all schemes, and for stricter privacy constraints (small ϵ), lower projection dimension achieves smaller relative error. As ϵ tends higher, the privacy constraint becomes less restrictive, and schemes with higher projection dimension perform better because of the additional rows of information.

B. MNIST Handwritten Digits Dataset

We consider a reduced version of the MNIST hand-written digits dataset [19], where we only take the digits 4 and 9, leading to 11791 data samples, and only consider the 300




Fig. 3. Test error of the schemes vs. n for $\epsilon = 0.2$, for MNIST

pixels that contain the most energy across these data samples. Mapping the digit labels to $+1$ and -1 , and vectorizing each data image, we solve the corresponding linear problem, which generates a model that classifies 4's versus 9's. Figure 3 gives the resulting test error (subject to a 80%/20% training/test set partition) for the logarithmic ($n'_1 := 500(\log(k) + 1)$), linear ($n'_2 := 500 \frac{k+1}{2}$), and full ($n' = n = 1000k$) random projections, additive noise, as well as (non-private) least squares. To evaluate values of n smaller than 11791, we randomly sample the dataset. The results are averaged over 10 trials. Similar to the random case, we observe that random projection with logarithmic dimensions result in the best performance, while preserving MI-DP with $\epsilon = 0.2$.

REFERENCES

- [1] S. S. Vempala, *The random projection method*, vol. 65. American Mathematical Soc., 2005.
- [2] C. Dwork, A. Roth, et al., "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [3] A. D. Sarwate and K. Chaudhuri, "Signal processing and machine learning with differential privacy: Algorithms and challenges for continuous data," *IEEE signal processing magazine*, vol. 30, no. 5, pp. 86–94, 2013.
- [4] P. Cuff and L. Yu, "Differential privacy as a mutual information constraint," in *ACM CCS*, pp. 43–54, ACM, 2016.
- [5] M. Pilanci and M. J. Wainwright, "Randomized sketches of convex programs with sharp guarantees," *IEEE Transactions on Information Theory*, vol. 61, pp. 5096–5115, Sept 2015.
- [6] R. Bassily, A. Smith, and A. Thakurta, "Private empirical risk minimization: Efficient algorithms and tight error bounds," in *Foundations of Computer Science (FOCS)*, pp. 464–473, Oct 2014.
- [7] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate, "Differentially private empirical risk minimization," *Journal of Machine Learning Research*, vol. 12, no. Mar, pp. 1069–1109, 2011.
- [8] D. Kifer, A. Smith, and A. Thakurta, "Private convex empirical risk minimization and high-dimensional regression," in *ACM COLT*, 2012.
- [9] S. Zhou, J. Lafferty, and L. Wasserman, "Compressed and privacy-sensitive sparse regression," *IEEE Transactions on Information Theory*, vol. 55, no. 2, pp. 846–866, 2009.
- [10] K. Kenthapadi, A. Korolova, I. Mironov, and N. Mishra, "Privacy via the johnson-lindenstrauss transform," *Journal of Privacy and Confidentiality*, vol. 5, 2013.
- [11] S. P. Kasiviswanathan and H. Jin, "Efficient private empirical risk minimization for high-dimensional learning," in *International Conference on Machine Learning*, pp. 488–497, 2016.
- [12] J. Blocki, A. Blum, A. Datta, and O. Sheffet, "The Johnson-Lindenstrauss transform itself preserves differential privacy," in *Foundations of Computer Science (FOCS)*, pp. 410–419, IEEE, 2012.
- [13] O. Sheffet, "Differentially private ordinary least squares," in *International Conference on Machine Learning*, pp. 3105–3114, 2017.
- [14] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and differential privacy," in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pp. 51–60, Oct 2010.
- [15] A. Lapidoth and S. M. Moser, "Capacity bounds via duality with applications to multiple-antenna systems on flat-fading channels," *IEEE Transactions on Information Theory*, vol. 49, pp. 2426–2467, Oct 2003.
- [16] P.-Å. Wedin, "Perturbation theory for pseudo-inverses," *BIT Numerical Mathematics*, vol. 13, pp. 217–232, Jun 1973.
- [17] M. Rudelson and R. Vershynin, "Non-asymptotic theory of random matrices: extreme singular values," International Congress of Mathematicians, Hyderabad, India, 2010.
- [18] M. Showkatbakhsh, C. Karakus, and S. Diggavi, "Privacy-utility trade-off of linear regression under random projections and additive noise," 2018, Available on arXiv.
- [19] Y. LeCun, C. Cortes, and C. J. Burges, "The mnist database of handwritten digits,"