

Gradient Perturbation is Underrated for Differentially Private Convex Optimization

Da Yu^{1*†}, Huishuai Zhang^{2*}, Wei Chen², Jian Yin¹ and Tie-Yan Liu²

¹The School of Data and Computer Science, Sun Yat-sen University.

Guangdong Key Laboratory of Big Data Analysis and Processing, Guangzhou 510006, P.R.China

²Microsoft Research Asia, Beijing, China

{yuda3@mail2, issjyin@mail}.sysu.edu.cn, {huishuai.zhang, wche, tie-yan.liu}@microsoft.com

Abstract

Gradient perturbation, widely used for differentially private optimization, injects noise at every iterative update to guarantee differential privacy. Previous work first determines the noise level that can satisfy the privacy requirement and then analyzes the utility of noisy gradient updates as in the non-private case. In contrast, we explore how the privacy noise affects the optimization property. We show that for differentially private convex optimization, the utility guarantee of differentially private (stochastic) gradient descent is determined by an *expected curvature* rather than the minimum curvature. The *expected curvature*, which represents the average curvature over the optimization path, is usually much larger than the minimum curvature. By using the *expected curvature*, we show that gradient perturbation can achieve a significantly improved utility guarantee that can theoretically justify the advantage of gradient perturbation over other perturbation methods. Finally, our extensive experiments suggest that gradient perturbation with the advanced composition method indeed outperforms other perturbation approaches by a large margin, matching our theoretical findings.

1 Introduction

Machine learning has become a powerful tool for many applications. The training process often needs access to private datasets, e.g., applications in financial and medical fields. Recent work has shown that the model learned from training data may leak unintended information of individual records [Shokri *et al.*, 2017; Hitaj *et al.*, 2017]. It is known that *Differential Privacy (DP)* [Dwork *et al.*, 2006] is a golden standard for privacy preserving data analysis, which has also been deployed into real world applications [Erlingsson *et al.*, 2014; Abowd, 2016]. It provides provable privacy guarantee by ensuring that the influence of any individual record is negligible.

We study the fundamental problem when differential privacy meets machine learning: the *differentially private empirical*

risk minimization (DP-ERM) problem [Chaudhuri *et al.*, 2011; Bassily *et al.*, 2014; Talwar *et al.*, 2015; Wu *et al.*, 2017; Zhang *et al.*, 2017; Wang *et al.*, 2017; Jayaraman *et al.*, 2018; Feldman *et al.*, 2018; Iyengar *et al.*, 2019; Wang and Gu, 2019]. DP-ERM minimizes the empirical risk while guaranteeing the output of the learning algorithm differentially private with respect to the training data. Such privacy guarantee provides strong protection against potential adversaries [Hitaj *et al.*, 2017; Rahman *et al.*, 2018]. In order to guarantee privacy, it is necessary to introduce randomness to the learning algorithm. There are usually three ways to introduce randomness according to the time of adding noise: *output perturbation*, *objective perturbation* and *gradient perturbation*.

Output perturbation [Wu *et al.*, 2017; Zhang *et al.*, 2017] first runs the learning algorithm the same as in the non-private case then adds noise to the output parameter. *Objective perturbation* [Chaudhuri *et al.*, 2011; Iyengar *et al.*, 2019] perturbs the objective (i.e., the empirical loss) then releases the (approximated) minimizer of the perturbed objective. *Gradient perturbation* [Bassily *et al.*, 2014; Abadi *et al.*, 2016; Jayaraman *et al.*, 2018] perturbs each intermediate update.

Gradient perturbation comes with several advantages over output/objective perturbations. Firstly, gradient perturbation does not require strong assumption on the objective because it only needs to bound the sensitivity of each gradient update rather than the whole learning process. Secondly, gradient perturbation can release the noisy gradient at each iteration without damaging the privacy guarantee as differential privacy is immune to *post processing* [Dwork *et al.*, 2014]. Therefore, it is a favorable choice in the distributed optimization or the federated learning setting [Konečný *et al.*, 2015; Agarwal *et al.*, 2018; Jayaraman *et al.*, 2018]. At last, based on our own experiments and the observation in [Lee and Kifer, 2018], gradient perturbation achieves much better empirical utility than output/objective perturbations for DP-ERM.

However, the existing theoretical utility guarantee for gradient perturbation is the same as or strictly inferior to that of output/objective perturbation approaches as shown in Table 1. This motivates us to ask

Can we improve the theory for gradient perturbation that can reflect its empirical advantage?

In the analysis for gradient perturbation approach, all previous works [Bassily *et al.*, 2014; Wang *et al.*, 2017; Jayaraman *et al.*, 2018] derive the utility guarantee via two steps. They first determine the noise variance that meets the

*Equal contribution.

†The work was done when this author was an intern at Microsoft Research Asia.

Perturbation	Convex	μ -S.C.
Objective [Chaudhuri <i>et al.</i> , 2011]	$\frac{\sqrt{p}}{n\epsilon}$	$\frac{p}{\mu n^2 \epsilon^2}$
Output [Zhang <i>et al.</i> , 2017]	$(\frac{\sqrt{\beta p}}{n\epsilon})^{2/3}$	$\frac{\beta p}{\mu^2 n^2 \epsilon^2}$
Gradient (SGD) [Bassily <i>et al.</i> , 2014]	$\frac{\sqrt{p} \log^{3/2}(n)}{n\epsilon}$	$\frac{p \log^2(n)}{\mu n^2 \epsilon^2}$
Gradient (GD) [Jayaraman <i>et al.</i> , 2018]	N/A	$\frac{\beta p \log^2(n)}{\mu^2 n^2 \epsilon^2}$
DP-GD (Ours)	$\frac{\sqrt{p}}{n\epsilon} \wedge \frac{\beta p \log(n)}{\nu^2 n^2 \epsilon^2}$	$\frac{\beta p \log(n)}{\nu^2 n^2 \epsilon^2}$
DP-SGD (Ours)	$\frac{\sqrt{p} \log(n)}{n\epsilon} \wedge \frac{p \log(n)}{\nu n^2 \epsilon^2}$	$\frac{p \log(n)}{\nu n^2 \epsilon^2}$

Table 1: Expected excess empirical risk bounds (the smaller, the better) under (ϵ, δ) -DP. Notations: n and p are the number of samples and the number of parameters, respectively, and β, μ and ν are the smooth coefficient, the strongly convex coefficient and the *expected curvature*, respectively, and $\nu \geq \mu$ (see Section 3.1). All bounds should be read as $\mathcal{O}(\cdot)$. Operator \wedge selects the smaller element. The Lipschitz constant L is assumed to be 1. We omit $\log(1/\delta)$ for simplicity.

privacy requirement and then derive the utility guarantee by using the convergence analysis the same as in the non-private case. However, the noise to guarantee privacy naturally affects the optimization procedure, which was neglected in previous analysis.

In this paper, we utilize the interaction between the privacy noise and the optimization procedure, which help us establish new utility guarantees for gradient perturbation approaches. Our contributions are as follows.

- We introduce an *expected curvature* that characterizes the optimization property accurately when there is perturbation noise at each gradient update.
- We establish new utility guarantees for DP-GD and DP-SGD for both convex and strongly convex objectives based on the *expected curvature* rather than the usual minimum curvature, with significant improvement over previous work.
- We conduct extensive experiments on real world datasets. Our experiments suggest that gradient perturbation with the advanced composition theorem outperforms other perturbation approaches, which corroborates our theoretical findings nicely.

To the best of our knowledge, this is the first utility guarantee of DP-ERM that removes the dependency on the minimum curvature. Moreover, we argue that objective/output perturbation cannot utilize this expected curvature condition to improve their utility guarantee as no noise is injected in their training process.

2 Preliminary

We introduce notations and definitions in this section. We use $D = \{d_1, \dots, d_n\}$ to denote a dataset with n records. Each $d_i = \{s_i, t_i\}$ contains a feature vector s_i and a target t_i . The objective function $F(x; D)$ is defined as $F(x; D) \triangleq \frac{1}{n} \sum_{i=1}^n f(x; d_i)$, where $f(x; d_i) : \mathbb{R}^p \rightarrow \mathbb{R}$ is the loss of model $x \in \mathbb{R}^p$ for the record d_i .

For simplicity, we use $F(x)$ to denote $F(x; D)$ and $f_i(x)$ to denote $f(x; d_i)$. We use $\|v\|$ to denote the l_2 norm of a vector v . We use $\mathcal{X}_f^* = \arg \min_{x \in \mathbb{R}^p} f(x)$ to denote the set of optimal solutions of $f(x)$. Throughout this paper, we assume \mathcal{X}_f^* non-empty.

Definition 1 (Objective properties). *For any $x, y \in \mathbb{R}^p$, a function $f : \mathbb{R}^p \rightarrow \mathbb{R}$*

- *is L -Lipschitz if $|f(x) - f(y)| \leq L \|x - y\|$.*
- *is β -smooth if $\langle \nabla f(x) - \nabla f(y), x - y \rangle \leq \beta \|x - y\|^2$.*
- *is convex if $\langle \nabla f(x) - \nabla f(y), x - y \rangle \geq 0$.*
- *is μ -strongly convex (or μ -S.C.) if $\langle \nabla f(x) - \nabla f(y), x - y \rangle \geq \mu \|x - y\|^2$.*

The strong convexity coefficient μ is the lower bound of the minimum curvature of function f over the domain.

We say that two datasets D, D' are neighboring datasets (denoted as $D \sim D'$) if D can be obtained by arbitrarily modifying one record in D' (or vice versa). In this paper we consider the (ϵ, δ) -differential privacy.

Definition 2 ((ϵ, δ) -DP [Dwork *et al.*, 2006]). *A randomized mechanism $\mathcal{M} : D \rightarrow \mathcal{R}$ guarantees (ϵ, δ) -differential privacy if for any two neighboring input datasets D, D' and for any subset of outputs $S \subseteq \mathcal{R}$ it holds that $\Pr[\mathcal{M}(D) \in S] \leq e^\epsilon \Pr[\mathcal{M}(D') \in S] + \delta$.*

The parameter δ can be viewed as the probability that original ϵ -DP [Dwork *et al.*, 2006] fails and a meaningful setting requires $\delta \ll \frac{1}{n}$. By its definition, differential privacy controls the maximum influence that any individual record can produce. Smaller ϵ, δ implies less information leak but usually leads to worse utility. One can adjust ϵ, δ to trade off between privacy and utility.

DP-ERM requires the output $x_{out} \in \mathbb{R}^p$ is differentially private with respect to the input dataset D . Let $x_* \in \mathcal{X}_F^*$ be one of the optimal solutions of $F(x)$, the utility of DP-ERM algorithm is measured by *expected excess empirical risk*: $\mathbb{E}[F(x_{out}) - F(x_*)]$, where the expectation is taken over the algorithm randomness. Following previous work, we use utility guarantee to denote the upper bound on expected risk. The algorithm with smaller utility guarantee is better.

3 Main Results

In this section, we first define the *expected curvature* and discuss its property. We then use such expected curvature to improve the utility analysis for DP-SGD and DP-GD.

3.1 Expected Curvature

In the non-private setting, the convergence analysis of gradient descent for strongly convex objective explicitly relies on the strongly convex coefficient μ , i.e., the minimum curvature of the objective function over the domain, which can be extremely small for some objectives. Previous work on DP-ERM uses the same analysis as in the non-private case and therefore the resulting utility bounds rely on the minimum curvature. In our analysis, however, we study the interaction between privacy noise and the optimization procedure, and figure out that the optimization procedure is determined by certain *expected*

Dataset	Graduate Admission	Sonar	Boston Housing	Swedish Motor
Minimum	6.68×10^{-3}	3.63×10^{-4}	3.44×10^{-3}	9.07×10^{-6}
Average	0.40	0.19	0.34	0.24

Table 2: Minimum and average curvatures for linear regression with four datasets. The minimum and the average curvatures are the smallest and the average eigenvalues of the Hessian matrix, respectively. We normalize the features with standard min-max normalization for all datasets.

curvature rather than the usual minimum curvature. Definition 3 describes such *expected curvature* when the perturbation noise is Gaussian. We use $\mathbf{x}_* = \arg \min_{\mathbf{x} \in \mathcal{X}_*} \|\mathbf{x} - \mathbf{x}_1\|$ to denote the closest solution to the initial point.

Definition 3 (Expected curvature). A convex function $F : \mathbb{R}^p \rightarrow \mathbb{R}$, has expected curvature ν with respect to noise $\mathcal{N}(0, \sigma^2 \mathbf{I}_p)$ if for any $\mathbf{x} \in \mathbb{R}^p$ and $\tilde{\mathbf{x}} = \mathbf{x} - \mathbf{z}$ where $\mathbf{z} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_p)$, it holds that

$$\mathbb{E}[\langle \nabla F(\tilde{\mathbf{x}}), \tilde{\mathbf{x}} - \mathbf{x}_* \rangle] \geq \nu \mathbb{E}[\|\tilde{\mathbf{x}} - \mathbf{x}_*\|^2], \quad (1)$$

where the expectation is taken with respect to \mathbf{z} .

Remark 1. If $F(\mathbf{x})$ is μ -strongly convex, then $F(\mathbf{x})$ has ν expected curvature with $\nu \geq \mu$.

It can be verified that $\nu = \mu$ always holds because of the strongly convex definition.

In fact, ν is usually much larger than μ . Let $\mathbf{H}_{\mathbf{x}} = \nabla^2 F(\mathbf{x})$ be the Hessian matrix evaluated at \mathbf{x} . We use Taylor expansion to approximate the left hand side of Eq (1) as follows

$$\begin{aligned} \mathbb{E}[\langle \nabla F(\tilde{\mathbf{x}}), \tilde{\mathbf{x}} - \mathbf{x}_* \rangle] &\approx \mathbb{E}[\langle \nabla F(\mathbf{x}) - \mathbf{H}_{\mathbf{x}} \mathbf{z}, \mathbf{x} - \mathbf{z} - \mathbf{x}_* \rangle] \\ &= \langle \nabla F(\mathbf{x}), \mathbf{x} - \mathbf{x}_* \rangle + \mathbb{E}[\mathbf{z}^T \mathbf{H}_{\mathbf{x}} \mathbf{z}] \\ &= \langle \nabla F(\mathbf{x}), \mathbf{x} - \mathbf{x}_* \rangle + \sigma^2 \text{tr}(\mathbf{H}_{\mathbf{x}}). \end{aligned} \quad (2)$$

The approximation is reasonably accurate for smooth convex objectives. For convex objective, the Hessian matrix is positive semi-definite and $\text{tr}(\mathbf{H}_{\mathbf{x}})$ is the sum of the eigenvalues of $\mathbf{H}_{\mathbf{x}}$. The right hand side of Eq (1) can be written as

$$\nu \mathbb{E}[\|\tilde{\mathbf{x}} - \mathbf{x}_*\|^2] = \nu \mathbb{E}[\|\mathbf{x} - \mathbf{z} - \mathbf{x}_*\|^2] \quad (3)$$

$$= \nu \left(\|\mathbf{x} - \mathbf{x}_*\|^2 + p\sigma^2 \right). \quad (4)$$

Based on the above approximation, we can estimate the value of ν in Definition 3:

$$\nu \approx \frac{\text{tr}(\mathbf{H}_{\mathbf{x}})\sigma^2 + \langle \nabla F(\mathbf{x}), \mathbf{x} - \mathbf{x}_* \rangle}{p\sigma^2 + \|\mathbf{x} - \mathbf{x}_*\|^2}. \quad (5)$$

We note that $\langle \nabla F(\mathbf{x}), \mathbf{x} - \mathbf{x}_* \rangle \geq \mu \|\mathbf{x} - \mathbf{x}_*\|^2$ because of the definition of strongly convex definition¹. For bounded $\|\mathbf{x} - \mathbf{x}_*\|$ and relatively large σ^2 , this implies $\nu \approx \frac{\text{tr}(\mathbf{H}_{\mathbf{x}})}{p}$, the average curvature at \mathbf{x} . Large variance is a reasonable setting because meaningful differential privacy guarantee requires non-trivial amount of noise.

The above analysis suggests that ν can be independent of and much larger than μ . We next show that for linear regression, ν is larger than μ without approximation.

¹In fact, this is the restricted strongly convex with respect to the minimizer \mathbf{x}_* .

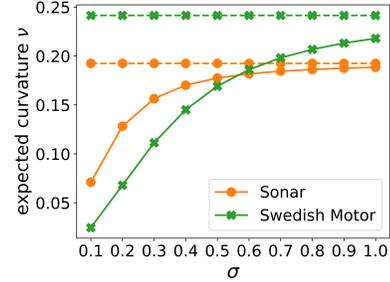


Figure 1: Expected curvatures ν with varying σ of two regression datasets. The dash line is the average curvature. The solid line is the expected curvature ν . We compute ν using Eq (5) with $\mathbf{x} = 0$.

Example 1. Suppose $F(\mathbf{x}) = \frac{1}{2n} \sum_i \|\mathbf{x} \mathbf{s}_i - t_i\|^2$ is the linear regression objective and $\mathbf{x} \in \mathcal{C}$, where \mathcal{C} is bounded convex set. Let β, μ be the smooth and the strongly convex coefficients, respectively, and $\beta > \mu$. Then the $\nu > \mu$ for any $\sigma > 0$.

Proof. Let $\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_n]^T \in \mathbb{R}^{n \times p}$ be the feature matrix. For linear regression, the Hessian matrix of $F(\mathbf{x})$ for any \mathbf{x} is $\mathbf{H} = \nabla^2 F(\mathbf{x}) = (1/n) \mathbf{S}^T \mathbf{S}$. The approximation in equation 2 becomes equality because higher order derivatives are 0 for linear regression objective.

Thus, we have Eq (1) hold for $\nu = \frac{\text{tr}(\mathbf{H})\sigma^2 + \langle \nabla F(\mathbf{x}), \mathbf{x} - \mathbf{x}_* \rangle}{p\sigma^2 + \|\mathbf{x} - \mathbf{x}_*\|^2}$. Because μ is the smallest eigenvalue of \mathbf{H} , we have $\text{tr}(\mathbf{H}) > p\mu$ and therefore $\nu > \mu$. \square

In practice, the Hessian matrix of linear regression could be ill-conditioned (the minimum curvature is extremely small). Table 2 shows the minimum and the average curvatures of four real world regression datasets collected from the Kaggle website². We further plot the expected curvatures with varying σ in Figure 1. We use $\mathbf{x} = 0^p$ and Eq (5) to compute the expected curvature. The expected curvature is almost as large as the average curvature when $\sigma = 1.0$.

Another example is the l_2 regularized logistic regression. The objective is strongly convex only due to the l_2 regularizer. Thus, the minimum curvature (strongly convex coefficient) is the regularization coefficient λ . [Shamir *et al.*, 2014] shows an optimal choice of λ is $\Theta(n^{-1/2})$. In practice, typical choice of λ is even smaller and could be on the order of n^{-1} . Therefore removing the dependence on minimum curvature is a significant improvement.

Another important class of objectives are the restricted strongly convex functions, which naturally arise in the under-determined systems. Although $\mu = 0$ for such case, ν is

²<https://www.kaggle.com/>

Algorithm 1 DP-GD

- 1: **Input:** running steps T ; learning rate η .
 - 2: Set $\sigma = \frac{2\sqrt{2}L\sqrt{T\log(1/\delta)}}{n\epsilon}$.
 - 3: **for** $t = 1$ **to** T **do**
 - 4: Compute $\mathbf{g}_t = \nabla F(\mathbf{x}_t)$.
 - 5: Update parameter $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta(\mathbf{g}_t + \mathbf{z}_t)$, where $\mathbf{z}_t \sim \mathcal{N}(0, \sigma^2 I_p)$.
 - 6: **end for**
-

strictly positive and is at least as large as the restricted strongly convex coefficient. Our theory can easily handle the restricted strongly convex function class.

3.2 New Utility Guarantee for DP-GD

Algorithm 1 describes the DP-GD algorithm (gradient perturbation). It has been shown that for the chosen σ , Algorithm 1 is (ϵ, δ) -DP [Jayaraman *et al.*, 2018]. We next present a new improved utility guarantee of DP-GD by using the *expected curvature*.

Theorem 1. *Suppose F is L -Lipschitz and β -smooth with ν expected curvature.*

If $\nu > 0$, letting $\eta \leq \frac{1}{\beta}$ and $T = \frac{2\log(n)}{\eta\nu}$, then we have

$$\mathbb{E}[F(\mathbf{x}_{T+1}) - F(\mathbf{x}_*)] = \mathcal{O}\left(\frac{\beta p \log(n) L^2 \log(1/\delta)}{\nu^2 n^2 \epsilon^2}\right).$$

If $\nu = 0$, letting $\eta = \frac{1}{\beta}$, $T = \frac{n\beta\epsilon}{\sqrt{p}}$ and $\bar{\mathbf{x}} = \frac{1}{T} \sum_{i=1}^T \mathbf{x}_{i+1}$, then we have

$$\mathbb{E}[F(\bar{\mathbf{x}}) - F(\mathbf{x}_*)] = \mathcal{O}\left(\frac{\sqrt{p} L^2 \log(1/\delta)}{n\epsilon}\right).$$

The expectations are taken over the algorithm randomness.

Proof. We first prove the case of $\nu > 0$. Since \mathbf{x}_t contains Gaussian perturbation noise \mathbf{z}_{t-1} , Definition 3 gives us

$$\mathbb{E}_{\mathbf{z}_{t-1}}[\langle \mathbf{x}_t - \mathbf{x}_*, \nabla F(\mathbf{x}_t) \rangle] \geq \nu \mathbb{E}_{\mathbf{z}_{t-1}}[\|\mathbf{x}_t - \mathbf{x}_*\|^2].$$

By the β -smooth condition, we have

$$\begin{aligned} \mathbb{E}_{\mathbf{z}_{t-1}}[\langle \mathbf{x}_t - \mathbf{x}_*, \nabla F(\mathbf{x}_t) \rangle] \\ \geq \frac{\nu}{2} \mathbb{E}_{\mathbf{z}_{t-1}}[\|\mathbf{x}_t - \mathbf{x}_*\|^2] + \frac{1}{2\beta} \mathbb{E}_{\mathbf{z}_{t-1}}[\|\nabla F(\mathbf{x}_t)\|^2]. \end{aligned} \quad (6)$$

Let $r_t = \|\mathbf{x}_t - \mathbf{x}_*\|$ be the estimation error at step t . The update rule gives us

$$r_{t+1}^2 = \|\mathbf{x}_t - \eta \nabla F(\mathbf{x}_t) - \eta \mathbf{z}_t - \mathbf{x}_*\|^2. \quad (7)$$

Taking expectation of Eq (7) with respect to $\mathbf{z}_t, \mathbf{z}_{t-1}$ and using Eq (6), we have

$$\begin{aligned} \mathbb{E}_{\mathbf{z}_t, \mathbf{z}_{t-1}}[r_{t+1}^2] &\leq (1 - \eta\nu) \mathbb{E}_{\mathbf{z}_{t-1}}[r_t^2] \\ &+ \left(\eta^2 - \frac{\eta}{\beta}\right) \mathbb{E}_{\mathbf{z}_{t-1}}[\|\nabla F(\mathbf{x}_t)\|^2] + p\eta^2 \sigma^2. \end{aligned} \quad (8)$$

Choose $\eta \leq \frac{1}{\beta}$. Applying Eq (8) and taking expectation with

respect to $\mathbf{z}_t, \mathbf{z}_{t-1}, \dots, \mathbf{z}_1$ iteratively yields

$$\mathbb{E}[r_{t+1}^2] \leq (1 - \eta\nu)^t r_1^2 + p\eta^2 \sum_{i=1}^t (1 - \eta\nu)^{t-i} \sigma_i^2. \quad (9)$$

Then plug $t = 2\log(n)/\eta\nu$, $\sigma = L\sqrt{T\log(1/\delta)}/n\epsilon$ in Eq (9). We obtain the desired bound by using the smoothness condition to convert Eq (9) to the objective function value bound.

We now prove the case of $\nu = 0$. The β -smooth condition gives us

$$F(\mathbf{x}_{t+1}) \leq F(\mathbf{x}_t) + \langle \nabla F(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle + \frac{\beta}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2.$$

Set $\eta = \frac{1}{\beta}$ and take expectation with respect to \mathbf{z}_t .

$$\mathbb{E}_{\mathbf{z}_t}[F(\mathbf{x}_{t+1})] = F(\mathbf{x}_t) - \frac{1}{2} \|\nabla F(\mathbf{x}_t)\|^2 / \beta + \frac{1}{2} p\sigma^2 / \beta.$$

Subtract $F(\mathbf{x}_*)$ from both sides and use convexity.

$$\begin{aligned} \mathbb{E}_{\mathbf{z}_t}[F(\mathbf{x}_{t+1}) - F(\mathbf{x}_*)] \\ \leq \langle \nabla F(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}_* \rangle - \frac{1}{2\beta} \|\nabla F(\mathbf{x}_t)\|^2 + \frac{1}{2\beta} p\sigma^2. \end{aligned} \quad (10)$$

Plug in $\nabla F(\mathbf{x}_t) = \beta(\mathbf{x}_t - \mathbf{x}_{t+1}) - \mathbf{z}_t$ and rearrange.

$$\begin{aligned} \mathbb{E}_{\mathbf{z}_t}[F(\mathbf{x}_{t+1}) - F(\mathbf{x}_*)] \\ \leq \frac{\beta}{2} (\|\mathbf{x}_t - \mathbf{x}_*\|^2 - \|\mathbf{x}_{t+1} - \mathbf{x}_*\|^2) + \frac{1}{\beta} p\sigma^2. \end{aligned} \quad (11)$$

Sum over $t = 1, \dots, T$, take expectation with respect to $\mathbf{z}_1, \dots, \mathbf{z}_T$ and use the convexity.

$$\mathbb{E}[F(\bar{\mathbf{x}}) - F(\mathbf{x}_*)] \leq \frac{\beta}{2T} \|\mathbf{x}_1 - \mathbf{x}_*\|^2 + \frac{1}{\beta} p\sigma^2. \quad (12)$$

Plugging in T and σ yields the desired bound. \square

Remark 2. *Theorem 1 only depends on the expected curvature ν over the training path.*

Theorem 1 significantly improves the original analysis of DP-GD because of our argument in Section 3.1. We note that if $\nu = 0$, then the curvatures are flatten in all directions. One example is the dot product function, which is used by [Bassily *et al.*, 2014] to derive their utility lower bound. Such simple function may not be commonly used as loss function in practice.

3.3 New Utility Guarantee for DP-SGD

Stochastic gradient descent has become one of the most popular optimization methods because of the cheap one-iteration cost. In this section we describe the DP-SGD (Algorithm 2) and show that *expected curvature* can also improve the utility analysis for DP-SGD. We note that $\nabla f(\mathbf{x})$ represents an element from the subgradient set evaluated at \mathbf{x} when the objective is not smooth. Before stating our theorems, we introduce the *moments accountant* technique (Lemma 1) that is essential to establish privacy guarantee.

Lemma 1 ([Abadi *et al.*, 2016]). *There exist constants c_1 and c_2 so that given running steps T , for any $\epsilon < c_1 T/n^2$,*

Algorithm 2 DP-SGD

- 1: **Input:** running steps T ; learning rate η_t .
 - 2: **for** $t = 1$ **to** T **do**
 - 3: Sample i_t from $[n]$ uniformly.
 - 4: Compute $\mathbf{g}_t = \nabla f_{i_t}(\mathbf{x}_t)$.
 - 5: Update parameter $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t (\mathbf{g}_t + \mathbf{z}_t)$, where $\mathbf{z}_t \sim \mathcal{N}(0, L^2 \sigma^2 I_p)$.
 - 6: **end for**
-

Algorithm 2 is (ϵ, δ) -differentially private for any $\delta > 0$ if we choose $\sigma \geq c_2 \frac{\sqrt{T \log(1/\delta)}}{n\epsilon}$.

We next establish the utility guarantee of Algorithm 2 in Theorem 2.

Theorem 2. Suppose F is L -Lipschitz with ν expected curvature. Choose σ as stated in Lemma 1.

If $\nu > 0$, letting $\eta_t = \frac{1}{t\nu}$ and $T = n^2 \epsilon^2$, then we have

$$\mathbb{E}[F(\mathbf{x}_T) - F(\mathbf{x}_*)] = \mathcal{O}\left(\frac{pL^2 \log(n) \log(1/\delta)}{n^2 \epsilon^2 \nu}\right).$$

If $\nu = 0$, letting $G = L\sqrt{1 + p\sigma^2}$, $\eta_t = \frac{D}{G\sqrt{t}}$ and $T = n^2 \epsilon^2$, then we have

$$\mathbb{E}[F(\mathbf{x}_T) - F(\mathbf{x}_*)] = \mathcal{O}\left(\frac{\sqrt{p \log(1/\delta)} L \log(n)}{n\epsilon}\right).$$

Proof. For the case of $\nu > 0$, we first show that $r_t := \|\mathbf{x}_t - \mathbf{x}_*\|$ satisfies

$$\mathbb{E}[r_t^2] \leq \frac{2L^2(1 + p\sigma^2)}{t\nu^2}, \quad (13)$$

which depends on ν rather than μ . Using the update rule to express out r_{t+1} and taking expectation with respect to \mathbf{z}_t and i_t , we have

$$\mathbb{E}_{\mathbf{z}_t, i_t}[r_{t+1}^2] \leq r_t^2 - 2\eta_t \langle \mathbf{x}_t - \mathbf{x}_*, \nabla F(\mathbf{x}_t) \rangle + \eta_t^2 L^2 + p\eta_t^2 L^2 \sigma^2. \quad (14)$$

Taking expectation to \mathbf{z}_{t-1} and using Definition 3, we have

$$\mathbb{E}[r_{t+1}^2] \leq (1 - 2\nu\eta_t) \mathbb{E}[r_t^2] + \eta_t^2 L^2 (1 + p\sigma^2). \quad (15)$$

Now we use induction to conduct the proof. Plugging $\eta_t = \frac{1}{t\nu}$ into Eq (15), we have Eq (13) hold for $t = 2$. Assuming that $\mathbb{E}[r_t^2] \leq \frac{2L^2(1+p\sigma^2)}{t\nu^2}$ holds for $t > 2$, then

$$\begin{aligned} \mathbb{E}[r_{t+1}^2] &\leq \left(1 - \frac{2}{t}\right) \mathbb{E}[r_t^2] + \frac{L^2(1 + p\sigma^2)}{\nu^2 t^2} \\ &\leq \left(\frac{2}{t} - \frac{3}{t^2}\right) \frac{L^2(1 + p\sigma^2)}{\nu^2} \leq \frac{2L^2(1 + p\sigma^2)}{(t+1)\nu^2}. \end{aligned} \quad (16)$$

We next use the results in [Shamir and Zhang, 2013] to derive the final utility bound. Shamir and Zhang show how to convert the convergence in r_t into the convergence in objective function value without any smooth assumption. The utility guarantee for the case of $\nu > 0$ can be derived by plugging

Eq (13) and $T = n^2 \epsilon^2$ into the Theorem 1 of [Shamir and Zhang, 2013].

For the case of $\nu = 0$, the proof is similar to the $\mu = 0$ case in [Bassily *et al.*, 2014]. We improve the bound by a $\log(n)$ factor because we use the advanced *moments accountant* technique. □

Theorem 2 shows the utility guarantee of DP-SGD depends on ν rather than μ . We set $T = \Theta(n^2)$ following [Bassily *et al.*, 2014]. We note that $T = \Theta(n^2)$ is necessary even for non-private SGD to reach $1/n^2$ precision. If we choose $T = \frac{n\epsilon}{\sqrt{p}}$ in the proof of Theorem 2, then the utility bound for $\nu > 0$ is

$$\mathbb{E}[F(\mathbf{x}_T) - F(\mathbf{x}_*)] = \mathcal{O}\left(\frac{\sqrt{p}L^2 \log(n)}{n\epsilon\nu}\right).$$

In contrast, the analysis of [Bassily *et al.*, 2014] yields $\mathbb{E}[F(\mathbf{x}_T) - F(\mathbf{x}_*)] = \mathcal{O}\left(\frac{\sqrt{p}L^2 \log^2(n)}{n\epsilon\mu}\right)$ if setting $T = \frac{n\epsilon}{\sqrt{p}}$, which still depends on the minimum curvature.

3.4 Discussion

In this section, we briefly discuss two other perturbation approaches and compare them to the gradient perturbation approach.

Output perturbation [Wu *et al.*, 2017; Zhang *et al.*, 2017] first runs the learning algorithm the same as in the non-private case and then adds data-independent noise to the output parameter according to sensitivity of the whole learning procedure. The magnitude of perturbation noise is propositional to the maximum influence one record can cause on the learned model over the whole learning procedure. The sensitivity is determined by the Lipschitz coefficient (to control the maximum expansion of two different records), and the smooth and the strongly convex coefficients (to measure the contraction effect of gradient update of the “good” records).

Objective perturbation [Chaudhuri *et al.*, 2011; Kifer *et al.*, 2012; Iyengar *et al.*, 2019] perturbs the objective (i.e., the empirical loss) with a random linear term and then releases the minimizer of the perturbed objective. The strong convexity is necessary to ensure the minimizer is unique³. The sensitivity of objective perturbation is the maximum change of the minimizer that one record can produce. [Chaudhuri *et al.*, 2011] uses the largest and the smallest eigenvalue (i.e. the smooth and strongly convex coefficient) of the objective’s Hessian matrix to upper bound such change.

In comparison, *Gradient perturbation* is more flexible than the output/objective perturbations. For gradient perturbation, the sensitivity is only determined by the Lipschitz coefficient which is easy to obtain by using the gradient clipping technique. More critically, we always have $\nu = \mu$ when the gradient update does not involve perturbation noise. Therefore, gradient perturbation is the only method can leverage such effect of noise among three existing perturbation methods.

³It adds L_2 regularization to obtain strong convexity if the original objective is not strongly convex.

dataset	Adult	MNIST	Real-sim	RCV1
# records	45220	65000	72309	50000
# features	104	784	20958	47236
# classes	2	10	2	2

Table 3: Detailed description of four real world datasets.

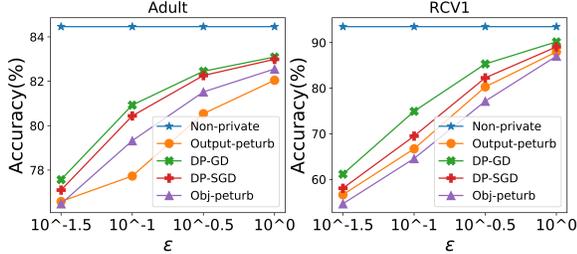


Figure 2: Algorithm validation accuracy (in %) with varying ϵ . NP represents non-private baseline.

4 Experiment

In this section, we evaluate the performance of DP-GD and DP-SGD on multiple real world datasets. Objective functions are *logistic regression* and *softmax regression* for binary and multi-class datasets, respectively.

Datasets. We present the results of four benchmark datasets in [Iyengar *et al.*, 2019], including one multi-class dataset (MNIST) and two with high dimensional features (Real-sim, RCV1)⁴. Detailed description of datasets can be found in Table 3. We use the same pre-processing as in [Iyengar *et al.*, 2017]. We use 80% data for training and the rest for testing, the same as [Iyengar *et al.*, 2019].

Implementation details. We track *Rényi differential privacy (RDP)* [Mironov, 2017] and convert it to (ϵ, δ) -DP. Running step T is chosen from $\{50, 200, 800\}$ for both DP-GD and DP-SGD. The standard deviation of the added noise σ is set to be the smallest value such that the privacy budget is allowable to run desired steps. We ensure each loss function is Lipschitz by clipping individual gradient. The method in [Goodfellow, 2015] allows us to clip individual gradient efficiently. Clipping threshold is set as 1 (0.5 for high dimensional datasets because of the sparse gradient). Privacy parameter δ is set as $\frac{1}{n^2}$. The l_2 regularization coefficient is set as 1×10^{-4} . For DP-GD, learning rate is chosen from $\{0.1, 1.0, 5.0\}$ ($\{0.2, 2.0, 10.0\}$ for high dimensional datasets). For DP-SGD, we use moments accountant to track the privacy loss and the sampling ratio is set as 0.1 (roughly the mini-batch size is $0.1 \times$ dataset size). The learning rate of DP-SGD is twice as large as DP-GD and it is divided by 2 at the middle of training. All reported numbers are averaged over 20 runs.

Baseline algorithms. The baseline algorithms include state-of-the-art objective and output perturbation algorithms. For the objective perturbation, we use *Approximate Minima Perturbation (AMP)* [Iyengar *et al.*, 2019]. For the output perturbation, we use the algorithm in [Wu *et al.*, 2017] (Output

⁴We conduct experiments on all 7 datasets in [Iyengar *et al.*, 2019] and present 4 of them due to limited space.

	Adult	MNIST	Real-sim	RCV1
Non-private	84.8	91.9	93.3	93.5
AMP	79.3	71.9	73.1	64.5
Out-SGD	77.4	69.4	73.2	66.7
DP-SGD	80.4	87.5	73.8	70.4
DP-GD	80.9	88.6	76.1	74.9

Table 4: Algorithm validation accuracy (in %) on various kinds of real world datasets. Privacy parameter ϵ is 0.1 for binary dataset and 1 for multi-classes datasets.

perturbation SGD). We adopt the implementation and hyperparameters in [Iyengar *et al.*, 2019] for both algorithms. For the multi-class classification tasks, [Wu *et al.*, 2017] and [Iyengar *et al.*, 2019] divide the privacy budget evenly and train multiple binary classifiers because their algorithms need to compute smooth coefficient before training and therefore are not directly applicable to softmax regression.

Experiment results. The validation accuracy results for all evaluated algorithms with $\epsilon = 0.1$ (1.0 for multi-class datasets) are presented in Table 4. We also plot the accuracy results with varying ϵ in Figure 2. These results confirm our theory in Section 3: gradient perturbation achieves better performance than other perturbation methods as it leverages the average curvature ν . We note that in [Iyengar *et al.*, 2019] the performance of DP-SGD is not as good as in our experiment. Our setup has two differences from theirs. First, they set the noise variance to be the upper bound (Algorithm 2 in their Appendix C) in Lemma 1. In contrast, we use smaller noise by carefully estimating the exact privacy loss through numerical integration as in [Abadi *et al.*, 2016]. Second, we decay the learning rate of DP-SGD during training while [Iyengar *et al.*, 2019] use constant learning rate. Learning rate decay is important for both theoretical analysis and empirical performance of DP-SGD. The above differences lead to a significant improvement on the performance of DP-SGD.

5 Conclusion

In this paper, we show the privacy noise helps optimization analysis, which is used to improve the utility guarantee of both DP-GD and DP-SGD. Experiments on real world datasets corroborate our theoretical findings nicely. In the future, it is interesting to consider utilizing the expected curvature to improve the utility guarantee of other gradient perturbation based algorithms.

Acknowledgements

The authors want to thank the anonymous reviewers for their constructive comments. This work is supported by the National Natural Science Foundation of China (U1711262, U1611264, U1711261, U1811261, U1811264), National Key R&D Program of China (2018YFB1004404), Guangdong Basic and Applied Basic Research Foundation (2019B1515130001), Key R&D Program of Guangdong Province (2018B010107005). Jian Yin and Huishuai Zhang are corresponding authors.

References

- [Abadi *et al.*, 2016] Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *ACM SIGSAC Conference on Computer and Communications Security*, 2016.
- [Abowd, 2016] John M. Abowd. The challenge of scientific reproducibility and privacy protection for statistical agencies. *Census Scientific Advisory Committee*, 2016.
- [Agarwal *et al.*, 2018] Naman Agarwal, Ananda Theertha Suresh, Felix Yu, Sanjiv Kumar, and Brendan McMahan. cpsgd: Communication-efficient and differentially-private distributed SGD. In *Advances in Neural Information Processing Systems*, 2018.
- [Bassily *et al.*, 2014] Raef Bassily, Adam Smith, and Abhradeep Thakurta. Differentially private empirical risk minimization: Efficient algorithms and tight error bounds. *Annual Symposium on Foundations of Computer Science*, 2014.
- [Chaudhuri *et al.*, 2011] Kamalika Chaudhuri, Claire Monteleoni, and Anand D. Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 2011.
- [Dwork *et al.*, 2006] Cynthia Dwork, Krishnamurthy Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2006.
- [Dwork *et al.*, 2014] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 2014.
- [Erlingsson *et al.*, 2014] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, 2014.
- [Feldman *et al.*, 2018] Vitaly Feldman, Ilya Mironov, Kunal Talwar, and Abhradeep Thakurta. Privacy amplification by iteration. In *IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 2018.
- [Goodfellow, 2015] Ian Goodfellow. Efficient per-example gradient computations. *arXiv:1510.01799*, 2015.
- [Hitaj *et al.*, 2017] Briland Hitaj, Giuseppe Ateniese, and Fernando Pérez-Cruz. Deep models under the gan: information leakage from collaborative deep learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017.
- [Iyengar *et al.*, 2017] Roger Iyengar, Joseph P. Near, Dawn Song, Om Thakkar, Abhradeep Thakurta, and Lun Wang. Differentially private convex optimization benchmark, 2017.
- [Iyengar *et al.*, 2019] Roger Iyengar, Joseph P. Near, Dawn Song, Om Thakkar, Abhradeep Thakurta, and Lun Wang. Towards practical differentially private convex optimization. In *IEEE Symposium on Security and Privacy*, 2019.
- [Jayaraman *et al.*, 2018] Bargav Jayaraman, Lingxiao Wang, David Evans, and Quanquan Gu. Distributed learning without distress: Privacy-preserving empirical risk minimization. In *Advances in Neural Information Processing Systems*, 2018.
- [Kifer *et al.*, 2012] Daniel Kifer, Adam Smith, and Abhradeep Thakurta. Private convex empirical risk minimization and high-dimensional regression. In *Conference on Learning Theory*, 2012.
- [Konečný *et al.*, 2015] Jakub Konečný, Brendan McMahan, and Daniel Ramage. Federated optimization: Distributed optimization beyond the datacenter. *arXiv preprint arXiv:1511.03575*, 2015.
- [Lee and Kifer, 2018] Jaewoo Lee and Daniel Kifer. Concentrated differentially private gradient descent with adaptive per-iteration privacy budget. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018.
- [Mironov, 2017] Ilya Mironov. Rényi differential privacy. In *IEEE Computer Security Foundations Symposium*, 2017.
- [Rahman *et al.*, 2018] Md Atiqur Rahman, Tanzila Rahman, Robert Laganiere, Noman Mohammed, and Yang Wang. Membership inference attack against differentially private deep learning model. *Transactions on Data Privacy*, 2018.
- [Shamir and Zhang, 2013] Ohad Shamir and Tong Zhang. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *International Conference on Machine Learning*, 2013.
- [Shamir *et al.*, 2014] Ohad Shamir, Nati Srebro, and Tong Zhang. Communication-efficient distributed optimization using an approximate newton-type method. In *International conference on machine learning*, 2014.
- [Shokri *et al.*, 2017] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *IEEE Symposium on Security and Privacy (SP)*, 2017.
- [Talwar *et al.*, 2015] Kunal Talwar, Abhradeep Guha Thakurta, and Li Zhang. Nearly optimal private lasso. In *Advances in Neural Information Processing Systems*, 2015.
- [Wang and Gu, 2019] Lingxiao Wang and Quanquan Gu. Differentially private iterative gradient hard thresholding for sparse learning. In *International Joint Conference on Artificial Intelligence*, 2019.
- [Wang *et al.*, 2017] Di Wang, Minwei Ye, and Jinhui Xu. Differentially private empirical risk minimization revisited: Faster and more general. In *Advances in Neural Information Processing Systems*, 2017.
- [Wu *et al.*, 2017] Xi Wu, Fengan Li, Arun Kumar, Kamalika Chaudhuri, Somesh Jha, and Jeffrey Naughton. Bolt-on differential privacy for scalable stochastic gradient descent-based analytics. In *ACM International Conference on Management of Data*, 2017.
- [Zhang *et al.*, 2017] Jiaqi Zhang, Kai Zheng, Wenlong Mou, and Liwei Wang. Efficient private erm for smooth objectives. In *International Joint Conference on Artificial Intelligence*, 2017.