

Private Non-smooth Empirical Risk Minimization and Stochastic Convex Optimization in Subquadratic Steps

Janardhan Kulkarni* Yin Tat Lee† Daogao Liu‡

Abstract

We study the differentially private Empirical Risk Minimization (ERM) and Stochastic Convex Optimization (SCO) problems for non-smooth convex functions. We get a (nearly) optimal bound on the excess empirical risk and excess population loss with subquadratic gradient complexity. More precisely, our differentially private algorithm requires $O(\frac{N^{3/2}}{d^{1/8}} + \frac{N^2}{d})$ gradient queries for optimal excess empirical risk, which is achieved with the help of subsampling and smoothing the function via convolution. This is the first subquadratic algorithm for the non-smooth case when d is super constant. As a direct application, using the iterative localization approach of Feldman et al. [FKT20], we achieve the optimal excess population loss for stochastic convex optimization problem, with $O(\min\{N^{5/4}d^{1/8}, \frac{N^{3/2}}{d^{1/8}}\})$ gradient queries. Our work makes progress towards resolving a question raised by Bassily et al. [BFGT20], giving first algorithms for private ERM and SCO with subquadratic steps.

We note that independently Asi et al. [AFKT21] gave other algorithms for private ERM and SCO with subquadratic steps.

*Algorithms Group, Microsoft Research (MSR) Redmond. jakul@microsoft.com

†yintat@uw.edu. University of Washington and Microsoft Research. Supported by NSF awards CCF-1749609, DMS-1839116, DMS-2023166, Microsoft Research Faculty Fellowship, Sloan Research Fellowship, Packard Fellowships.

‡liudaogao@gmail.com. University of Washington. Part of the work was done while visiting Shanghai Qi Zhi Institute.

Contents

1	Introduction	2
1.1	Our Contributions	5
1.2	Our Techniques	6
1.3	Concurrent and Independent Work	7
2	Preliminaries	8
2.1	Convex Optimization	8
2.2	Differential Privacy	9
3	A Meta Algorithm for DP Convex Optimization	9
4	Differentially Private ERM	10
4.1	Non-smooth Strongly-convex Functions	10
4.1.1	Smoothing function	11
4.1.2	Algorithm	12
4.1.3	Utility and Privacy	12
4.2	General Non-smooth Convex Functions	15
5	Differentially Private SCO	17
References		22
A	Proof of Theorem 3.1	25

1 Introduction

Privacy has become an important consideration for learning algorithms dealing with sensitive data. Over the past decade, differential privacy, introduced in the seminal work of [DMNS06], has established itself as the defacto notion of privacy for machine learning problems. In this paper, we revisit Empirical Risk Minimization (ERM) and Stochastic Convex Optimization (SCO) problem, which are one of the most important and simplest problems in statistics and machine learning, in differential privacy setting. In the ERM problem, we are given a family of convex functions $\{f(\cdot, x)\}_{x \in \Xi}$ over a closed convex set $\mathcal{K} \subset \mathbb{R}^d$, a data set $S = \{x_1, \dots, x_N\}$ drawn from some unknown distribution \mathcal{P} over the universe Ξ , and the objective is to

$$\text{minimize } \hat{F}(\omega) := \frac{1}{N} \sum_{x_i \in S} f(\omega, x_i) \quad \text{over } \omega \in \mathcal{K},$$

while in the SCO the objective is to

$$\text{minimize } F(\omega) := \mathbb{E}_{x \sim \mathcal{P}} f(\omega, x) \quad \text{over } \omega \in \mathcal{K},$$

Differentially private convex optimization has been studied extensively for over a decade now [CM08, RBHT09, CMS11, KST12, JT14, TTZ14, BST14, TTZ15, KJ16, WLK⁺17, FTS17, ZZMW17, WYX17, INS⁺19]. Most of the previous results are focus on DP-ERM and roughly speaking, there are three major approaches in DP-ERM: output perturbation, objective perturbation, and gradient perturbation. Output perturbation approach is based on the sensitivity method proposed by [DMNS06] and adds noise to the final output to the standard ERM problem [CM08, RBHT09, CMS11, ZZMW17]. Objective perturbation [CM08, CMS11, KST12, TTZ14] means to perturb the objective function we want to minimize. In the gradient perturbation approach, we add noise to the first order information using optimization algorithms such as Stochastic Gradient Descent (SGD). This approach was first proposed in [BST14] and was later extended by [TTZ14, WYX17], and has lead to the state-of-the-art theoretical bounds for DP-ERM. For an experimental comparison of various approaches to solving DP-ERM we refer the readers to [RBHT09, INS⁺19].

DP-ERM for smooth convex functions is well understood in the sense that we know (near) linear time algorithms that achieve optimal excess empirical risk. We refer the readers to [WYX17] for more details. However, for the more general non-smooth convex loss functions our understanding is not yet complete, which is the focus of this paper. A summary of the state-of-the-art results and our contributions for the non-smooth convex loss functions is given in Table 1 (General Convex) and Table 2 (Strongly Convex). We will discuss the concurrent work [AFKT21] separately at the end of the introduction, and the following discussion are only limited to the previous work.

[KST12] used the objective perturbation method to design a DP-algorithm with $O(\frac{GD\sqrt{d}\log(1/\delta)}{\sqrt{N}\varepsilon})$ excess empirical risk. This result was improved significantly by [BST14], who first showed a lower bound of $\Omega(\min\{GD, \frac{GD\sqrt{d}}{N\varepsilon}\})$ on the excess empirical risk for DP-ERM. Further, they gave an algorithm with excess empirical risk $O(\frac{GD\log^{\frac{3}{2}}(N/\delta)\sqrt{d\log(1/\delta)}}{N\varepsilon})$, which is sub-optimal by a factor of $\log^{\frac{3}{2}}(N/\delta)$. Their algorithm is based on a modification of SGD by adding Gaussian noise to the gradients to make it DP. The privacy analysis proceeds via amplification by sampling and the strong composition theorem. Roughly speaking, the logarithmic blowup in the excess empirical risk is due to two reasons: 1) the strong composition theorem requires that at each step one needs to add Gaussian noise with a larger variance; 2) They used sub-optimal convergence rate $O(\log T/\sqrt{T})$ for T -step SGD.

However, getting the optimal bounds with small gradient complexity for non-smooth case turns out to be a more difficult problem. This was noted by [WYX17], who raised it as an important open problem. This question was answered in [BFTT19], who gave an algorithm with almost optimal excess empirical risk. To achieve this, [BFTT19] first consider the smooth case, and give an improved privacy analysis via the Moments Accountant technique proposed by [ACG⁺16]. They extend their result to non-smooth case by applying Moreau-Yosida envelope technique (a.k.a. Moreau envelope smoothing) [Nes05] to make the function smooth. However, this technique is computationally inefficient and leads to $O(N^{4.5})$ -gradient computations for the whole algorithm. This limitation was overcome in a recent work of [BFGT20] who gave the optimal excess empirical risk guarantee with $O(N^2)$ -gradient computations. The privacy analysis of this result also used Moments Accountant method, and they used the standard online-to-batch conversion technique [CBCG04] to prove the high-probability bound on the excess empirical error of SGD, which leads to the near optimal bound in expectation. We remark that all the papers [BFTT19, BFGT20] above not only study the ERM problem, but also consider more general DP-SCO settings and uniform stability, and in some cases, results on ERM are byproducts of the more general results.

As we can see from Table 1 and Table 2, all the previously known results (except the concurrent work [AFKT21]) achieving near optimal excess empirical risk bounds require at least $O(N^2)$ -gradient computations. It is natural to ask if there are lower bounds to rule out algorithms with subquadratic gradient complexity that can match the error bounds of the above results.

	Excess Risk	Empirical Risk	Gradient Complexity
[KST12]	$\frac{GD\sqrt{d}\log(1/\delta)}{\sqrt{N}\varepsilon}$		N/A
[BST14]	$\frac{GD\log^{\frac{3}{2}}(N/\delta)\sqrt{d}\log(1/\delta)}{N\varepsilon}$		N^2
[BFTT19]	$\frac{GD\sqrt{d}\log(1/\delta)}{N\varepsilon}$		$N^{4.5}$
[BFGT20]	$\frac{GD\sqrt{d}\log(1/\delta)}{N\varepsilon}$		N^2
[AFKT21]	$\frac{GD\sqrt{d}\log(1/\delta)}{N\varepsilon}$		N^2/\sqrt{d}
Ours	$\frac{GD\sqrt{d}\log(1/\delta)}{N\varepsilon}$		$\frac{N^{3/2}}{d^{1/8}} + \frac{N^2}{d}$

Table 1: Comparisons with previous (ε, δ) -differential private algorithms when objective function is G -Lipschitz and convex over a convex set $\mathcal{K} \subset \mathbb{R}^d$ of diameter D . The results are stated asymptotically and the big O notation is hidden for simplicity. The lower bound is $\Omega(\min\{GD, \frac{GD\sqrt{d}}{N\varepsilon}\})$ [BST14].

As Table 3 and Table 4 show, a similar situation arises in Stochastic Convex Optimization (SCO), which is a closely related problem compared to ERM. In the SCO problem, we want to minimize the objective function $F(\omega) = \mathbb{E}_{x \sim \mathcal{P}}[f(\omega, x)]$ for some unknown distribution \mathcal{P} over the universe Ξ . Many results for SCO [BST14, BFTT19, BFGT20] are directly based on ERM; that is, solving the ERM and analyzing the generalization error. The first non-trivial result for general convex loss functions achieving excess population loss of $O\left(GD\left(\frac{d^{1/4}}{\sqrt{N}} + \frac{\sqrt{d}}{N\varepsilon}\right)\right)$ was given by [BST14], who showed the result by first solving the ERM problem and bounding the generalization error. They used the result on universal convergence directly, namely, bounding $\sup_{\omega \in \mathcal{K}} \mathbb{E}[F(\omega) - \hat{F}(\omega)]$. But this

	Excess Risk	Empirical Risk	Gradient Complexity
[KST12]	$\frac{G^2 d \log(1/\delta)}{\mu N^{3/2} \varepsilon^2}$		N/A
[BST14]	$\frac{G^2 \log^2(N/\delta) d \log(1/\delta)}{\mu N^2 \varepsilon^2}$		N^2
[BFTT19]	$\frac{G^2 d \log(1/\delta)}{\mu N^2 \varepsilon^2}$		$N^{4.5}$
[BFGT20]	$\frac{G^2 d \log(1/\delta)}{\mu N^2 \varepsilon^2}$		N^2
Ours	$\frac{G^2 d \log(1/\delta)}{\mu N^2 \varepsilon^2}$		$\frac{N^{3/2}}{d^{1/8}} + \frac{N^2}{d}$

Table 2: Comparisons with previous (ε, δ) -differential private algorithms when objective function is G -Lipschitz and μ -strongly convex over a convex set $\mathcal{K} \subset \mathbb{R}^d$. The results are stated asymptotically and the big O notation is hidden for simplicity. The lower bound is $\Omega(\min\{\frac{G^2}{\mu}, \frac{G^2 d}{\mu N^2 \varepsilon^2}\})$ [BST14].

method has its limitations; For example, [Fel16] showed that lower bound of universal convergence is $\Omega(\sqrt{d/N})$ for some (not necessarily convex) loss functions. Later, [BFTT19], [FKT20] and [BFGT20] obtained near optimal excess population loss with significantly better running times (gradient complexity). The privacy analysis in these papers relied on recent advances in the privacy techniques such as the Moments Accountant method [ACG⁺16], Rényi differential privacy (RDP) [Mir17] and the Privacy Amplification by Iteration [FMTT18] and other fast stochastic convex optimization algorithms such as [JNN19]. The excess population loss bound in most of these works followed by solving a (phased) convex (regularized) ERM problem and then appealing to the uniform stability property [HRS16] or the iterative localization approach [FKT20] to do the generalization error analysis.

	Excess Population Loss	Gradient Complexity
[BST14]	$GD\left(\frac{d^{1/4} \log(n/\delta)}{\sqrt{N}} + \frac{d^{1/2} \log^2(n/\delta)}{N\varepsilon}\right)$	N^2
[BFTT19]	$GD\left(\frac{1}{\sqrt{N}} + \frac{\sqrt{d \log(1/\delta)}}{N\varepsilon}\right)$	$N^{4.5}$
[FKT20]	$GD\left(\frac{1}{\sqrt{N}} + \frac{\sqrt{d \log(1/\delta)}}{N\varepsilon}\right)$	$N^2 \log(1/\delta)$
[BFGT20]	$GD\left(\frac{1}{\sqrt{N}} + \frac{\sqrt{d \log(1/\delta)}}{N\varepsilon}\right)$	N^2
[AFKT21]	$GD\left(\frac{1}{\sqrt{N}} + \frac{\sqrt{d \log(1/\delta)}}{N\varepsilon}\right)$	$\min\{N^{3/2}, N^2/\sqrt{d}\}$
Ours	$GD\left(\frac{1}{\sqrt{N}} + \frac{\sqrt{d \log(1/\delta)}}{N\varepsilon}\right)$	$\min\{N^{5/4} d^{1/8}, N^{3/2}/d^{1/8}\}$

Table 3: Comparisons with previous (ε, δ) -differential private algorithms when objective function is G -Lipschitz and convex over a convex set $\mathcal{K} \subset \mathbb{R}^d$. The results are stated asymptotically and the big O notation is hidden for simplicity. The lower bound is $\Omega(GD(\frac{1}{\sqrt{N}} + \frac{\sqrt{d}}{N\varepsilon}))$ [BST14].

Despite these impressive improvements, as the Table 3 and Table 4 suggest, the previous algorithms (except the concurrent work [AFKT21]) which achieve the optimal excess population loss still require $O(N^2)$ -gradient computations. Indeed, [BFGT20] write that

	Excess Population Loss	Gradient Complexity
[BFTT19]	$\frac{G^2}{\mu} \left(\frac{1}{N} + \frac{d \log(1/\delta)}{N^2 \varepsilon^2} \right)$	$N^{4.5}$
[FKT20]	$\frac{G^2}{\mu} \left(\frac{1}{N} + \frac{d \log(1/\delta)}{N^2 \varepsilon^2} \right)$	$N^2 \log(1/\delta)$
[BFGT20]	$\frac{G^2}{\mu} \left(\frac{1}{N} + \frac{d \log(1/\delta)}{N^2 \varepsilon^2} \right)$	N^2
[AFKT21]	$\frac{G^2}{\mu} \left(\frac{1}{N} + \frac{d \log(1/\delta)}{N^2 \varepsilon^2} \right)$	$\min\{N^{3/2}, N^2/\sqrt{d}\}$
Ours	$\frac{G^2}{\mu} \left(\frac{1}{N} + \frac{d \log(1/\delta)}{N^2 \varepsilon^2} \right)$	$\min\{N^{5/4} d^{1/8}, N^{3/2}/d^{1/8}\}$

Table 4: Comparisons with previous (ε, δ) -differential private algorithms when objective function is G -Lipschitz and μ -strongly convex over a convex set $\mathcal{K} \subset \mathbb{R}^d$. The results are stated asymptotically and the big O notation is hidden for simplicity. The lower bound is $\Omega\left(\frac{G^2}{\mu} \left(\frac{1}{N} + \frac{d \log(1/\delta)}{N^2 \varepsilon^2} \right)\right)$ [BST14].

“ *Proving that quadratic running time is necessary for general non-smooth DP-SCO is a very interesting open problem...* ”

Understanding if the lower bound is the right answer to the above questions or one can design algorithms with subquadratic gradient complexity is the main motivation that spurred our work.

1.1 Our Contributions

Given the close connections between the ERM and SCO problems and the bottleneck on gradient complexity of all known algorithms, it is natural to ask if the open question raised in [BFGT20] also holds for the ERM problem. As noted earlier, the state-of-art algorithms for DP-ERM achieving optimal excess empirical risk bounds require $O(N^2)$ -gradient computations.

The main contribution of this paper is to show that we can obtain subquadratic gradient complexity bound for ERM when the dimension is super constant. In particular, for the important regime of over-parameterization ($d \geq N$), we achieve a bound of $N^{1+3/8}$. Combining our private ERM algorithm and the iterative localization approach proposed in [FKT20], we can achieve optimal excess population loss with gradient complexity $O\left(N + \min\{\sqrt{\varepsilon} N^{5/4} d^{1/8}, \frac{\varepsilon N^{3/2}}{d^{1/8} \log^{1/4}(1/\delta)}\}\right)$.

Let $\mathcal{K}_r = \{y \mid y = \omega + z, \omega \in \mathcal{K}, z \in \mathbb{R}^d, \|z\|_r \leq r\}$. We now state the main technical contributions of this paper formally.

Theorem 1.1 (DP-ERM). *Suppose $\mathcal{K} \subset \mathbb{R}^d$ is a closed convex set of diameter D and $\{f(\cdot, x)\}_{x \in \Xi}$ is a family of G -Lipschitz and convex functions over \mathcal{K}_r , where $r = \frac{D \sqrt{d \log(1/\delta)}}{\varepsilon N}$ ¹. For $\varepsilon, \delta \leq 1/2$, given any sample set S consists of N samples from Ξ and arbitrary initial point $\omega_0 \in \mathcal{K}$, we have a (ε, δ) -differentially private algorithm \mathcal{A} which takes*

$$O\left(\frac{\varepsilon N^{3/2}}{d^{1/8} \log^{1/4}(1/\delta)} + \frac{\varepsilon^2 N^2}{d \log(1/\delta)}\right)$$

gradient queries and outputs ω_T such that

$$\mathbb{E}[\hat{F}(\omega_T) - \hat{F}^*] = O\left(\frac{GD \sqrt{d \log(1/\delta)}}{\varepsilon N}\right),$$

¹We only need consider the non-trivial case when $\frac{\sqrt{d \log(1/\delta)}}{\varepsilon N} \leq 1$, or any feasible solution is good enough. This means that $r = O(D)$, which is a mild assumption.

where $D = \|\omega^* - \omega_0\|_2$, $\widehat{F}(\omega) = \frac{1}{N} \sum_{x_i \in S} f(\omega, x_i)$, $\widehat{F}^* = \min_{\omega \in \mathcal{K}} \widehat{F}(\omega)$, and the expectation is taken over the randomness of the algorithm.

Moreover, if $\{f(\cdot, x)\}_{x \in \Xi}$ is also μ -strongly convex functions over \mathcal{K}_r , we have an (ε, δ) -differentially private algorithm which takes the same bound of gradient queries and outputs ω_T such that

$$\mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}^*] = O\left(\frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2}\right).$$

As we have mentioned, combining our private ERM algorithm with the iterative localization technique, we can also give the first algorithm achieving optimal excess population loss with (strictly) sub-quadratic steps for all dimensions:

Theorem 1.2 (DP-SCO). *Suppose $\varepsilon, \delta \leq 1/2$ and sample set S consists of N samples drawn i.i.d from a distribution \mathcal{P} over Ξ . Let $\{f(\cdot, x)\}_{x \in \Xi}$ is convex and G -Lipschitz with respect to ℓ_2 norm and convex over \mathcal{K}_r , where $r = \frac{D\sqrt{d \log(1/\delta)}}{\varepsilon N}$, there is an (ε, δ) -differentially private algorithm which takes*

$$O(N + \min\{\sqrt{\varepsilon} N^{5/4} d^{1/8}, \frac{\varepsilon N^{3/2}}{d^{1/8} \log^{1/4}(1/\delta)}\})$$

gradient queries to get a solution ω_T

$$\mathbb{E}[F(\omega_T) - F(\omega^*)] = O(GD(\frac{1}{\sqrt{N}} + \frac{\sqrt{d \log(1/\delta)}}{N \varepsilon})).$$

Moreover, if $\{f(\cdot, x)\}_{x \in \Xi}$ is also μ -strongly convex over \mathcal{K}_r , we can meet the same gradient query complexity and get a solution ω_T such that:

$$\mathbb{E}[F(\omega_T) - F(\omega^*)] = O\left(\frac{G^2}{\mu} \left(\frac{d \log(1/\delta)}{\varepsilon^2 N^2} + \frac{1}{N}\right)\right).$$

Finally, we note that our results can also capture the regularized ERM and SCO, which shows up often in the previous work such as [RBHT09, KST12, WYX17, INS⁺19]. Briefly, in the regularized problem, there is one more simple (and convex) function $h(\omega)$ added to the objective function to encourage certain solutions with better structure. The objective function then takes the form $\frac{1}{N} \sum_{x_i \in S} f(\omega, x_i) + h(\omega)$. We get asymptotically same results for the regularized ERM/SCO problem with straightforward modifications.

1.2 Our Techniques

Most of the previous works [BST14, BFTT19, BFGT20] that achieve near optimal bounds for ERM and SCO are based on adaptations of SGD to make it differentially private. The information theoretic lower bound of $\Omega(1/\sqrt{T})$ for T -step SGD may be one of the important reasons why we can not get subquadratic gradient complexity for non-smooth convex ERM easily. Consider the algorithm in [BFGT20] as an example. It needs to add Gaussian noise $v \sim \mathcal{N}(0, \sigma^2 I_{d \times d})$ with $\sigma^2 = \frac{G^2 \log(1/\delta)}{\varepsilon^2}$ to each gradient. By a standard analysis of SGD, we can only show an excess empirical risk of $\Theta(\frac{D\sqrt{d\sigma^2}}{\sqrt{T}})$, which requires us to set $T = \Omega(N^2)$ to get ideal bound, thus hitting the quadratic barrier.

We deviate from the above approaches for designing private algorithms for non-smooth functions. First notice that the gradient complexity $O(\frac{\varepsilon N^{3/2}}{d^{1/8} \log^{1/4}(1/\delta)} + \frac{\varepsilon^2 N^2}{d \log(1/\delta)})$ in Theorem 1.1 is the same

for both strongly convex and general non-smooth functions; same holds for DP-SCO. This is not a coincidence; If we can achieve optimal empirical risk (population loss) for one case, then we can achieve optimal empirical risk (population loss) for another with the same privacy guarantee and gradient complexity. In fact, the Figure 1 shows the relationship among these different problems.

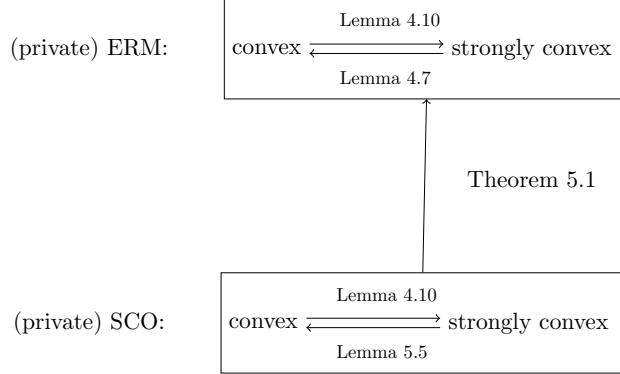


Figure 1: Reductions between ERM and SCO for general convex and strongly convex cases. As the lower bound of excess population loss is $\Omega(GD(\frac{1}{\sqrt{N}} + \frac{\sqrt{d}}{N\epsilon}))$ while the lower bound of empirical risk is $\Omega(\frac{GD\sqrt{d}}{N\epsilon})$, we do not know how to reduce from ERM to SCO.

Our result for the general convex non-smooth case is obtained by providing a reduction to the strongly convex non-smooth case. Thus, our task becomes designing better algorithms for the strongly convex non-smooth functions. Rather than using SGD, we let the objective function take convolution with a sphere kernel to make it smooth. We then use the accelerated stochastic approximation algorithm in [GL12] for solving strongly convex stochastic optimization problems. However, this is not enough, as the required noise that needs to be added to the gradients to make the algorithm private is too large to get subquadratic gradient complexity, even if we use the tighter Moments Accountant technique [ACG⁺16]. We overcome this by increasing the batch size to an appropriate value. Combining these ideas together, we show that the amount of noise we add can be reduced to achieve the optimal excess empirical loss, and we get the gradient complexity of $O(\max\{N^{3/2}/d^{1/8}, N^2/d\})$.

For SCO, we get the gradient complexity of $O(\min\{N^{5/4}d^{1/8}, N^{3/2}/d^{1/8}\})$ via a direct application of the iterative localization approach of Feldman et al [BFGT20]. The intuition behind iterative localization is using private ERM to solve regularized objective functions which have low sensitivity, iteration by iteration. Each iteration reduces the distance to an approximate minimizer by a multiplicative factor, so after logarithmic number of phases we are done.

1.3 Concurrent and Independent Work

In an independent and concurrent work, [AFKT21] give a new analysis of private regularized mirror descent to do the private ERM. Then they combine the iterative localization approach to achieve the optimal excess population loss for SCO. Their result also achieves subquadratic gradient complexity. More formally, they get $O(\log N \cdot \min(N^{3/2}\sqrt{\log d}, N^2/\sqrt{d}))$ for SCO in query complexity. We compare their gradient complexity with ours in Figure 2. Finally, we remark that the main motivation of [AFKT21] was to study SCO problem in more general ℓ_p norms as much of the literature has focussed on the ℓ_2 -norm. They also give news results in ℓ_p -bounded domain together with another concurrent work [BGN21].

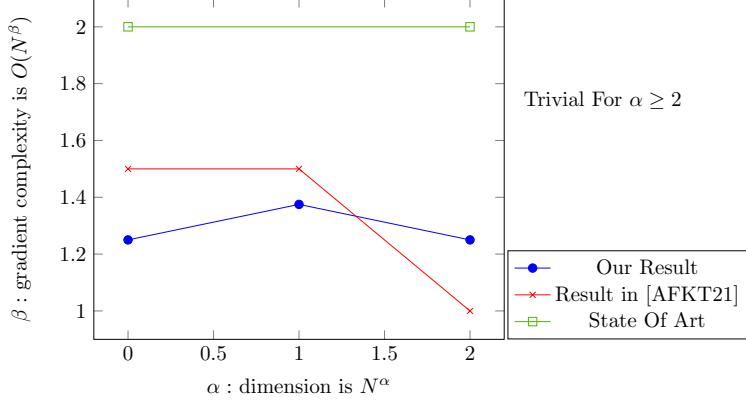


Figure 2: Comparison among our results, the recent result in [AFKT21] and the previous best one for the non-trivial regime ($d \leq N^2$). Suppose ε, δ are small constants. Our result is faster for the important case $d \leq N^{1+1/3}$.

Road map

We will give some basic definitions and theorems about convex optimization and differential privacy in Section 2. In Section 3, we give a general algorithm framework for private convex optimization. The results of DP-ERM are given in Section 4 and the results of DP-SCO are shown in Section 5. Some technical proofs are left in Appendix A.

2 Preliminaries

In this section, we briefly recall some of the main definitions we use from the convex optimization theory and differential privacy. We refer the readers to excellent books [Nes05, DR14] for more details on these topics.

2.1 Convex Optimization

Definition 2.1 (Empirical risk minimization, Stochastic Convex Optimization). Let $\mathcal{K} \subset \mathbb{R}^d$ be a closed convex set of diameter D . Given a family of convex loss functions $\{f(\omega, x)\}_{x \in \Xi}$ of ω over \mathcal{K} and a set of samples $S = \{x_1, \dots, x_n\}$ over the universe Ξ , the objective of Empirical Risk Minimization (ERM) is to minimize

$$\hat{F}(\omega) = \frac{1}{N} \sum_{x_i \in S} f(\omega, x_i).$$

The excess empirical loss with respect to a solution ω is defined by $\hat{F}(\omega) - \hat{F}^*$, where $\hat{F}^* = \min_{\omega \in \mathcal{K}} \hat{F}(\omega)$.

Stochastic Convex Optimization (SCO) wants to output a solution ω to minimize the expected loss (also referred to *population loss*) $F(\omega) - F^*$ where $F(\omega) = \mathbb{E}[x \sim \mathcal{P}] f(\omega, x)$ and $F^* = \min_{\omega \in \mathcal{K}} F(\omega)$.

Definition 2.2 (L -Lipschitz Continuity). A function $f : \mathcal{K} \rightarrow \mathbb{R}$ is L -Lipschitz continuous over the domain $\mathcal{K} \subset \mathbb{R}^d$ if the following holds for all $\omega, \omega' \in \mathcal{K}$: $|f(\omega) - f(\omega')| \leq L \|\omega - \omega'\|_2$.

Definition 2.3 (β -Smoothness). A function $f : \mathcal{K} \rightarrow \mathbb{R}$ is β -smooth over the domain $\mathcal{K} \subset \mathbb{R}^d$ if for all $\omega, \omega' \in \mathcal{K}$, $\|\nabla f(\omega) - \nabla f(\omega')\|_2 \leq \beta \|\omega - \omega'\|_2$.

Definition 2.4 (μ -Strongly convex). A differentiable function $f : \mathcal{K} \rightarrow \mathbb{R}$ is called strongly convex with parameter $\mu > 0$ if the following inequality holds for all points $\omega, \omega' \in \mathcal{K}$,

$$\langle \nabla f(\omega) - \nabla f(\omega'), \omega - \omega' \rangle \geq \mu \|\omega - \omega'\|_2^2.$$

Equivalently,

$$f(\omega') \geq f(\omega) + \nabla f(\omega)^\top (\omega' - \omega) + \frac{\mu}{2} \|\omega' - \omega\|_2^2.$$

2.2 Differential Privacy

Definition 2.5 (Differential privacy). A randomized mechanism \mathcal{M} is (ε, δ) -differentially private if for any event $\mathcal{O} \in \text{Range}(\mathcal{M})$ and for any neighboring databases that differ in a single data element, one has

$$\Pr[\mathcal{M}(S) \in \mathcal{O}] \leq \exp(\varepsilon) \Pr[\mathcal{M}(S') \in \mathcal{O}] + \delta.$$

Lemma 2.6 (Proposition 2.1 in [DR14]). *(Post-Processing) Let $\mathcal{M} : \mathbb{N}^{|\Xi|} \rightarrow R$ be a randomized algorithm that is (ε, δ) -differentially private. Let $f : R \rightarrow R'$ be an arbitrary randomized mapping. Then $f \circ \mathcal{M} : \mathbb{N}^{|\Xi|} \rightarrow R'$ is (ε, δ) -differentially private.*

Theorem 2.7 (Basic Composition). *Let $\mathcal{M}_i : \mathbb{N}^{|\Xi|} \rightarrow R_i$ be $(\varepsilon_i, \delta_i)$ -differentially private. Then if mechanism $\mathcal{M}_{[k]} : \mathbb{N}^{|\mathcal{X}|} \rightarrow \prod_{i=1}^k \mathcal{R}_i$ is defined to be $\mathcal{M}_{[k]}(x) = (\mathcal{M}_1(x), \dots, \mathcal{M}_k(x))$, then $\mathcal{M}_{[k]}$ is $(\sum_{i=1}^k \varepsilon_i, \sum_{i=1}^k \delta_i)$ -differentially private.*

3 A Meta Algorithm for DP Convex Optimization

Many convex optimization algorithms with noisy first-order information have the following simple format.

Algorithm 1: Meta Algorithm META

```

1 Input: The objective convex function  $F(\omega)$  we want to minimize, an initial point  $\omega_0$ .
2 Process: for  $phases t = 1, \dots, \mathbf{do}$ 
3   | Get the noisy gradient  $G_t \approx \nabla F(\omega_{t-1})$ ;
4   | Update the result by some sub-procedure:  $\omega_t \leftarrow \text{Sub-procedure}(\omega_{t-1}, G_t)$ ;
5 end
6 Output: Some function of  $\{\omega_i\}_{i \geq 1}$ .
```

We can use the above algorithmic framework to solve ERM privately. Specifically, we make two simple modifications to make it private. First, we compute gradients over a uniform sample of some size B . Next, we add a carefully calibrated Gaussian noise to these gradients and take average, before updating our results. This gives us a meta differentially private algorithm for convex optimization problems, and is described in Algorithm 2. The DP analysis then follows from a careful accounting of the privacy budget lost in each iteration, and the bound on excess empirical risk comes from the

property of the optimization algorithm.

Algorithm 2: Private Meta Algorithm META_{DP}

- 1 **Input:** Sample set $S = \{x_1, \dots, x_N\}$, the objective convex function $F(\omega)$ we want to minimize, the initial point ω_0 , and privacy parameter ε, δ ;
- 2 **Process:** **for** $phases t = 1, \dots, T$ **do**
- 3 Select a random sample set S_t from the uniform distribution over all subsets of S of size B ;
- 4 Let $G_t = (\sum_{x_i \in S_t} \nabla f(\omega_{t-1}, x_i) + v)/B$, where $v \sim \mathcal{N}(0, \sigma^2 I_{d \times d})$;
- 5 Update the result by some sub-procedure $\omega_t \leftarrow \text{Sub-procedure}(\omega_{t-1}, G_t)$;
- 6 **end**
- 7 **Output:** Some function of $\{\omega_i\}_{i \geq 1}$.

The above framework is a sub-sampled Gaussian mechanism, for which we can use tCDP proposed in [BDRS18] to analyze its privacy guarantee. As this is a direct application of the main result in [BDRS18], we leave the proof of the following theorem in the Appendix.

Theorem 3.1. *Suppose $\{f(\cdot, x)\}_{x \in \Xi}$ is a family of G -Lipschitz and convex functions over \mathcal{K} , for $\varepsilon < c_1 B^2 T / N^2$, $B \leq N/10$ and $1/2 \geq \delta > 0$, by setting $\sigma = \frac{c_2 G B \sqrt{T \log(1/\delta)}}{\varepsilon N}$ for some constant c_1 and c_2 , META_{DP} is (ε, δ) -differential private.*

4 Differentially Private ERM

In this section, we present private algorithms achieving the optimal excess empirical loss with subquadratic gradient complexity when the dimension is super constant. We consider non-smooth strongly-convex functions first, and then show how to reduce the general non-smooth case to the strongly-convex case in the last subsection.

4.1 Non-smooth Strongly-convex Functions

We use the framework introduced in Section 3 and give a faster private algorithm. Specifically, we modify a stochastic convex optimization algorithm in [GL12] to fit into our framework. First we recall some properties of that algorithm.

Suppose $f : \mathcal{K} \rightarrow \mathbb{R}$ is a convex function, and the objective is to get

$$\Psi^* := \min_{\omega \in \mathcal{K}} \{\Psi(\omega) = f(\omega) + h(\omega)\},$$

where \mathcal{K} is a closed convex set and $h(\omega)$ is a simple convex function with known structure.

Theorem 4.1 (Proposition 9 in [GL12]). *If the following conditions are met:*

- For some $L \geq 0, M \geq 0$ and $\mu > 0$,

$$\frac{\mu}{2} \|y - \omega\|_2^2 \leq f(y) - f(\omega) - \langle g(\omega), y - \omega \rangle \leq \frac{L}{2} \|y - \omega\|_2^2 + M \|y - \omega\|_2, \quad \forall \omega, y \in \mathcal{K},$$

where $g(\omega) \in \partial f(\omega)$ and $\partial f(\omega)$ denotes the sub-differential of f at ω .

- For each call of the stochastic oracle \mathcal{G} with the input $\omega_t \in \mathcal{K}$, the stochastic oracle \mathcal{G} can output an independent vector $\mathcal{G}(\omega_t)$ such that $\mathbb{E}[\mathcal{G}(\omega_t)] \in \partial f(\omega_t)$.
- For any $t \geq 1$ and $\omega_t \in \mathcal{K}$, $\mathbb{E}[\|\mathcal{G}(\omega_t) - g(\omega_t)\|_2^2] \leq V$.

Then after T iterations, Algorithm 3 given below outputs ω_T such that

$$\mathbb{E}[\Psi(\omega_T) - \Psi^*] \leq O\left(\frac{L\|\omega_0 - \omega^*\|_2^2}{T^2} + \frac{M^2 + V}{\mu T}\right),$$

where $\omega^* = \arg \min_{\omega \in \mathcal{K}} \Psi(\omega)$ and $\Psi^* = \Psi(\omega^*)$.

Algorithm 3: Accelerated stochastic approximation (AC-SA) algorithm

```

1 Input: Initial point  $\omega_0 \in \mathcal{K}$ .
2 Initialization: Set the initial point  $\omega_0^{ag} = \omega_0$ ;
3 Set the step-size parameters  $\alpha_t = \frac{2}{t+2}$  and  $\gamma_t = \frac{4L}{t(t+1)}$ ;
4 Process:
5 for  $t = 1, \dots, T$  do
6   Let  $\omega_t^{md} = \frac{(1-\alpha_t)(\mu+\gamma_t)}{\gamma_t+(1-\alpha_t^2)\mu} \omega_{t-1}^{ag} + \frac{\alpha_t[(1-\alpha_t)\mu+\gamma_t]}{\gamma_t+(1-\alpha_t^2)\mu} \omega_{t-1}$ ;
7   Query Oracle  $\mathcal{G}_t \equiv \mathcal{G}(\omega_t^{md})$ ;
8   Let  $\omega_t = \arg \min_{\omega \in \mathcal{K}} \{\alpha_t[\langle \mathcal{G}_t, \omega \rangle + h(\omega) + \mu\|\omega_t^{md} - \omega\|_2^2] + [(1 - \alpha_t)\mu + \gamma_t]\|\omega_{t-1} - \omega\|_2^2\}$ ;
9    $\omega_t^{ag} = \alpha_t \omega_t + (1 - \alpha_t) \omega_{t-1}^{ag}$ ;
10 end
11 Return:  $\omega_T^{ag}$ .

```

4.1.1 Smoothing function

From the statement of Theorem 4.1, it is clear that the Algorithm 3 gives much better convergence rates for smooth functions. As we are considering non-smooth functions, we need an efficient way to smooth the objective function without introducing too much error. In the next few paragraphs, we show how to achieve that.

Recall that D denotes the diameter of the closed convex set $\mathcal{K} \subset \mathbb{R}^d$. Suppose $\{f(\cdot, x)\}_{x \in \Xi}$ is a family of G -Lipschitz and μ -strongly convex functions over \mathcal{K} . This implies that for any sample set S , the empirical loss function $\widehat{F}(\omega)$ we consider is G -Lipschitz and μ -strongly convex over the domain \mathcal{K} .

We do a convolution on $f(\cdot, x)$, which is denoted by $f(\cdot, x) * n_r$. The objective function after the convolution step becomes $\widehat{F}_{n_r}(\omega) = \frac{1}{N} \sum_{x_i \in S} \mathbb{E}_{y \sim n_r} f(\omega + y, x_i)$, where n_r is the uniform density on the ℓ_2 ball of radius r . By Lemma 7 and Lemma 8 in [YNS12] while the forth result on forth item was supplemented by Lemma E.2 in [DBW12], we know the claim below.

Claim 4.2. Suppose $\{f(\cdot, x)\}_{x \in \Xi}$ is convex and G -Lipschitz over $\mathcal{K} + B_2(0, r)$. For $\omega \in \mathcal{K}$, $\widehat{F}_{n_r}(\omega)$ has following properties:

- $\widehat{F}(\omega) \leq \widehat{F}_{n_r}(\omega) \leq \widehat{F}(\omega) + Gr$;
- $\widehat{F}_{n_r}(\omega)$ is G -Lipschitz;
- $\widehat{F}_{n_r}(\omega)$ is $\frac{G\sqrt{d}}{r}$ -Smooth;
- For random variables $y \sim n_r$ and x uniformly from S , one has

$$\mathbb{E}[\nabla f(\omega + y, x)] = \nabla \widehat{F}_{n_r}(\omega)$$

and

$$\mathbb{E}[\|\nabla \widehat{F}_{n_r}(\omega) - \nabla f(\omega + y, x)\|_2^2] \leq G^2.$$

Furthermore, the convolution operation preserves strong convexity, which implies the fact below.

Fact 4.3. *Let n_r be the uniform density on the ℓ_2 ball of radius r , and $f : \mathcal{K}_r \rightarrow \mathbb{R}$ be a μ -strongly convex function over \mathcal{K}_r . Then $\mathbb{E}_{y \sim n_r} f(y + \cdot)$ is μ -strongly convex over \mathcal{K} .*

4.1.2 Algorithm

Now we state the our modifications to make AC-SA private and prove its properties. Recall that $y \sim n_r$ is a d -dimension vector drawn from the uniform density on the ℓ_2 ball of radius r . We start with the description of our algorithm.

Algorithm 4: Private AC-SA

- 1 **Input:** A convex set \mathcal{K} with diameter D , a family $\{f(\cdot, x_i)\}_{i \in [N]}$ of G -Lipschitz and μ -strongly convex functions over \mathcal{K} , an initial point $\omega_0 \in \mathcal{K}$, privacy parameters ε, δ , the batch size B , and the number of steps T .
- 2 Set $r \leftarrow \frac{D}{Td^{1/4}}$ and $\sigma \leftarrow \Theta\left(\frac{GB\sqrt{T \log(1/\delta)}}{\varepsilon N}\right)$;
- 3 Run the AC-SA with the Oracle \mathcal{G} defined below;
- 4 **Return:** The output of AC-SA
- 5 **Oracle $\mathcal{G}(\omega)$:**
- 6 Select a random sample set S_t from the uniform distribution over all subsets of S of size B .
- 7 **Return:** $(\sum_{x_i \in S_t} \partial f(\omega + y, x_i) + v)/B$, where $y \sim n_r$ and $v \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_{d \times d})$.

4.1.3 Utility and Privacy

It is not hard to show that Private AC-SA (Algorithm 4) is an instance of META_{DP} (see Section 3), so we have the following guarantee directly by Theorem 3.1.

Lemma 4.4. *For $\varepsilon \leq c_1 B^2 T / N^2$, $\delta \leq 1/2$, $B \leq N/10$ and $\sigma = \frac{c_2 GB\sqrt{T \log(1/\delta)}}{\varepsilon N}$ where $c_1 \leq 1$, $c_2 \geq 1$ are constants, Private AC-SA is (ε, δ) -DP.*

Now we consider the accuracy of Private AC-SA.

Lemma 4.5. *Under the assumptions defined in Algorithm Private AC-SA, after T iterations, it outputs ω_T such that*

$$\mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}^*] = O\left(\frac{G^2/B + \sigma^2 d/B^2}{\mu T} + \frac{GDd^{1/4}}{T}\right),$$

where $\omega^* = \arg \min_{\omega \in \mathcal{K}} \widehat{F}(\omega)$, and $\widehat{F}^* = \min_{\omega} \widehat{F}(\omega)$.

Proof. By Claim 4.2, we know that \widehat{F}_{n_r} is G -Lipschitz and $\frac{G\sqrt{d}}{r}$ -smooth. Furthermore, by Fact 4.3, we know that \widehat{F}_{n_r} is μ -strongly convex. For any t th iteration, one has that $\mathbb{E}[\mathcal{G}_t] = \nabla \widehat{F}_{n_r}(\omega_t^{md})$ and $\mathbb{E}[\|\mathcal{G}_t - \nabla \widehat{F}_{n_r}(\omega_t^{md})\|_*^2] \leq G^2/B + \sigma^2 d/B^2$. Then by Theorem 4.1 with $M = 0$, $L = \frac{G\sqrt{d}}{r}$, $V = G^2/B + \sigma^2 d/B^2$, we get

$$\mathbb{E}[\widehat{F}_{n_r}(\omega_T) - \min_{\omega} \widehat{F}_{n_r}(\omega)] = O\left(\frac{G^2/B + \sigma^2 d/B^2}{\mu T} + \frac{GD^2\sqrt{d}}{T^2 r}\right).$$

Next, by the first bullet of Claim 4.2, we know that $\widehat{F}(\omega) \leq \widehat{F}_{n_r}(\omega) \leq \widehat{F}(\omega) + Gr$ for any ω . Combining these together, we get

$$\begin{aligned} & \mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}(\omega^*)] \\ &= \mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}_{n_r}(\omega_T)] + \mathbb{E}[\widehat{F}_{n_r}(\omega_T) - \min_{\omega} \widehat{F}_{n_r}(\omega)] + \min_{\omega} \widehat{F}_{n_r}(\omega) - \widehat{F}(\omega^*) \\ &\leq 2Gr + O\left(\frac{G^2/B + \sigma^2 d/B^2}{\mu T} + \frac{GD^2\sqrt{d}}{T^2r}\right). \end{aligned}$$

By setting $r = \frac{Dd^{1/4}}{T}$, we complete the proof. \square

Before stating the main result of this section, we prove two technical lemmas that can remove the dependence on the diameter term. Lemma 4.6 below is used to prove Lemma 4.7.

Lemma 4.6. *Consider a sequence x_1, x_2, \dots . Suppose $0 \leq x_1 \leq n$ and $0 \leq x_{i+1} \leq \sqrt{x_i} + 1$, then for $k \geq \lceil \log \log n \rceil$, one has that $x_k \leq 16$.*

Proof. Without loss of generality, let $x_{i+1} = \sqrt{x_i} + 1$.

We construct another sequence y_1, \dots, y_k such that $y_1 = x_1$ and $y_{i+1} = 2\sqrt{y_i}$. Then by induction, it is easy to prove that for each $i \in [k]$, $y_i \geq x_i$. So we only need to prove that $y_k \leq 16$.

Let $z_i = \log_2 y_i$, then one has $z_{i+1} = z_i/2 + 1$. Obviously, we know that $z_i = 2^{-i+1}(z_1 - 2) + 2$ and $z_k \leq 4$, which means that $x_k \leq y_k \leq 16$. \square

Recall that the lower bound of strongly convex case is $\Omega\left(\frac{G^2}{\mu} + \frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2}\right)$ while for the general case is $\Omega\left(GD + \frac{GD\sqrt{d \log(1/\delta)}}{\varepsilon N}\right)$. Therefore, we only need to think about the case when $\frac{d \log(1/\delta)}{\varepsilon^2 N^2} \leq 1$, or the bound will be trivial. The following lemma says if we can achieve sum of these two lower bounds for strongly-convex case, then we can achieve the optimal bound for the strongly-convex case, which implies we can reduce the Strongly-Convex Case to General Convex Case.

Lemma 4.7 (Reduction to General Convex Case). *Given \widehat{F} is G -Lipschitz and μ -strongly convex. Suppose for any $\varepsilon, \delta < 1/2$, we have an (ε, δ) -differentially private algorithm \mathcal{A} which takes ω_0 as the initial start point and outputs a solution ω_T such that*

$$\mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}^*] = O\left(\frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2} + \frac{GD\sqrt{d \log(1/\delta)}}{\varepsilon N}\right),$$

where $\omega^* = \arg \min_{\omega \in \mathcal{K}} \widehat{F}(\omega)$ and $D = \|\omega_0 - \omega^*\|_2$. Then by taking \mathcal{A} as sub-procedure with some modifications on parameters, we can get an (ε, δ) -differentially private solution with excess empirical loss at most

$$\mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}^*] = O\left(\frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2}\right).$$

Furthermore, if \mathcal{A} uses $g(N, \varepsilon, \delta)$ many gradients, the new algorithm uses $\sum_{i \geq 1} g(N, \varepsilon/2^i, \delta/2^i)$ many gradients.

Remark 4.8. All algorithms in this paper uses less gradients if ε and δ are smaller. So, the new algorithm uses essentially as much as the given algorithm.

Proof. Repeat the private algorithm \mathcal{A} for $k = \lceil \log \log N^3 \rceil$ times. For the i th repetition, we start from the output of the last repetition and use \mathcal{A} as a sub-procedure with privacy parameter $\varepsilon_i = \varepsilon/2^{k+1-i}$ and $\delta_i = \delta/2^{k+1-i}$. (Note that the noise is decreasing so that the last step gives the best solution). We show that the last output has excess empirical risk at most $O\left(\frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2}\right)$.

More specifically, let ω_i be the output of the i th repetition, $\Delta_i = \mathbb{E}[\hat{F}(\omega_i) - \hat{F}^*]$ and $D_i^2 = \mathbb{E}[\|\omega_i - \omega^*\|^2]$. As the objective function \hat{F} is μ -strongly convex, we know that $\frac{1}{2}\mu D_i^2 \leq \Delta_i$ for all $i \geq 0$.

By the guarantee of the algorithm, there exists some constant $c \geq 1$ such that

$$\begin{aligned}\Delta_{i+1} &= \mathbb{E}[\hat{F}(\omega_{i+1}) - \hat{F}^*] \\ &\leq c \frac{G D_i \sqrt{d \log(1/\delta_i)}}{\varepsilon_i N} + c \frac{G^2 d \log(1/\delta_i)}{\mu \varepsilon_i^2 N^2} \\ &\leq c \frac{G \sqrt{d \log(1/\delta_i)}}{\varepsilon_i N} \sqrt{\frac{2\Delta_i}{\mu}} + \frac{E_i}{c},\end{aligned}$$

where we define $E_i = 2c^2 \frac{G^2 d \log(1/\delta_i)}{\mu \varepsilon_i^2 N^2}$.

As $E_i/E_{i+1} = \frac{\varepsilon_{i+1}^2 \log(1/\delta_i)}{\varepsilon_i^2 \log(1/\delta_{i+1})} \leq 8$, we can rearrange the above function and get

$$\begin{aligned}\frac{\Delta_{i+1}}{64E_{i+1}} &\leq \frac{\sqrt{\Delta_i E_i} + \frac{E_i}{c}}{64E_{i+1}} \\ &\leq \frac{E_i}{64E_{i+1}} \left(\sqrt{\frac{\Delta_i}{E_i}} + \frac{4}{c} \right) \\ &\leq \sqrt{\frac{\Delta_i}{64E_i}} + 1.\end{aligned}$$

By strong convexity one has that $\Delta_1 \leq G^2/\mu$, and $E_1 = \Omega(G^2 \log^3 N / (\mu N^2)) = \Omega(G^2 / (\mu N^3))$ by the definition, so $\Delta_1/E_1 \leq N^3$. Then by Lemma 4.6, after $k = \lceil \log \log N^3 \rceil$ repetitions, we get $\frac{\Delta_k}{64E_k} \leq 16$. This further implies that there is a solution with expected error

$$\mathbb{E}[\hat{F}(\omega_k) - \hat{F}^*] = O\left(\frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2}\right).$$

The privacy guarantee comes directly from the basic composition theorem (See Theorem 2.7). \square

We did not optimize constants in the calculations above. Now we are ready to state the main result for the strongly-convex case.

Theorem 4.9 (Strongly Convex Case for Theorem 1.1). *Suppose $\mathcal{K} \subset \mathbb{R}^d$ is a closed convex set of diameter D and $\{f(\cdot, x)\}_{x \in \Xi}$ is a family of G -Lipschitz and μ -strongly convex functions over \mathcal{K}_r where $r = \frac{D \sqrt{d \log(1/\delta)}}{\varepsilon N}$. For $\varepsilon, \delta \leq 1/2$, given any sample set S consists of N samples from Ξ and arbitrary initial point $\omega_0 \in \mathcal{K}$, we have an (ε, δ) -differentially private algorithm \mathcal{A} which takes $O\left(\frac{\varepsilon N^{\frac{3}{2}}}{d^{1/8} \log^{1/4}(1/\delta)} + \frac{\varepsilon^2 N^2}{d \log(1/\delta)}\right)$ gradient queries and outputs ω_T such that*

$$\mathbb{E}[\hat{F}(\omega_T) - \hat{F}^*] = O\left(\frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2}\right),$$

where $\hat{F}(\omega) = \frac{1}{N} \sum_{x_i \in S} f(\omega, x_i)$, $\hat{F}^* = \min_{\omega \in \mathcal{K}} \hat{F}(\omega)$, and the expectation is taken over the randomness of the algorithm itself.

Proof. By Lemma 4.5, the output ω of Private AC-SA satisfies

$$\mathbb{E}[\widehat{F}(\omega) - \widehat{F}^*] = O\left(\frac{\frac{G^2}{B} + \frac{\sigma^2 d}{B^2}}{\mu T} + \frac{GDd^{1/4}}{T}\right).$$

By setting $\sigma = \frac{c_2 GB \sqrt{T \log(1/\delta)}}{\varepsilon N}$ and $T = \lceil \frac{100\varepsilon N}{c_1 d^{1/4} \sqrt{\log(1/\delta)}} \rceil$ (c_1, c_2 are defined in Lemma 4.4), one has

$$\begin{aligned}\mathbb{E}[\widehat{F}(\omega) - \widehat{F}^*] &= O\left(\frac{G^2}{\mu BT} + \frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2} + \frac{GDd^{1/4}}{T}\right) \\ &= O\left(\frac{G^2}{\mu BT} + \frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2} + \frac{GD \sqrt{d \log(1/\delta)}}{\varepsilon N}\right)\end{aligned}$$

To ensure that Private AC-SA is (ε, δ) -DP, we set $B = \lceil \sqrt{\frac{\varepsilon N^2}{c_1 T}} + \frac{\varepsilon^2 N^2}{d \log(1/\delta) T} \rceil$. By our choice of T , we have $B \leq N/10$ and $\varepsilon \leq c_1 B^2 T / N^2$. Hence, we can apply Lemma 4.4 to conclude the guarantee of (ε, δ) differential privacy.

Furthermore, we get a solution ω such that

$$\mathbb{E}[\widehat{F}(\omega) - \widehat{F}^*] = O\left(\frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2} + \frac{GD \sqrt{d \log(1/\delta)}}{\varepsilon N}\right).$$

As for the total gradient complexity of our algorithm, we are under the assumption that $\frac{d \log(1/\delta)}{\varepsilon^2 N^2} \leq 1$, which means that $\frac{\varepsilon N}{d^{1/4} \sqrt{\log(1/\delta)}} \geq d^{1/4}$, and $T = \lceil \frac{100\varepsilon N}{c_1 d^{1/4} \sqrt{\log(1/\delta)}} \rceil = \Theta(\frac{\varepsilon N}{d^{1/4} \sqrt{\log(1/\delta)}})$.

As for the batch size, we know $\sqrt{\frac{\varepsilon N^2}{T}} + \frac{\varepsilon^2 N^2}{d \log(1/\delta) T} = \omega(1)$ and thus $B = \lceil \sqrt{\frac{\varepsilon N^2}{T}} + \frac{\varepsilon^2 N^2}{d \log(1/\delta) T} \rceil = \Theta(\sqrt{\frac{\varepsilon N^2}{T}} + \frac{\varepsilon^2 N^2}{d \log(1/\delta) T})$, from which we get the gradient complexity is

$$BT = \Theta\left(\frac{\varepsilon N^{\frac{3}{2}}}{d^{1/8} \log^{1/4}(1/\delta)} + \frac{\varepsilon^2 N^2}{d \log(1/\delta)}\right).$$

By Lemma 4.7 we can adjust Private AC-SA and get a final solution ω_T such that

$$\mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}^*] = O\left(\frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2}\right),$$

with gradient complexity $\Theta(\sum_{i=1}^{\log \log N^3} \frac{(\varepsilon/2^i) N^{3/2}}{d^{1/8} \log^{1/4}(2^i/\delta)} + \frac{(\varepsilon/2^i)^2 N^2}{d \log(2^i/\delta)}) = \Theta(\frac{\varepsilon N^{\frac{3}{2}}}{d^{1/8} \log^{1/4}(1/\delta)} + \frac{\varepsilon^2 N^2}{d \log(1/\delta)})$, which completes the proof. \square

4.2 General Non-smooth Convex Functions

In the general non-smooth case, we only assume that the family of functions $\{f(\cdot, x)\}_{x \in \Xi}$ is G -Lipschitz and convex over \mathcal{K} . We now give a reduction from this case to the strongly-convex case, which completes our second main result.

Lemma 4.10. *Suppose $\mathcal{K} \subset \mathbb{R}^d$ is a convex set of diameter D and let $\{f(\cdot, x)\}_{x \in \Xi}$ be a family of convex functions over \mathcal{K} , which are G -Lipschitz and μ -strongly convex. Given any sample set S*

consists of N samples from Ξ and other necessary inputs, suppose we have a (ε, δ) -DP algorithm \mathcal{A} which can output a solution ω_T such that

$$\mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}^*] = O\left(\frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2}\right),$$

where $\widehat{F}^* = \min_{\omega \in \mathcal{K}} \widehat{F}(\omega)$.

Then when $\{h(\cdot, x)\}_{x \in \Xi}$ is only G -Lipschitz and convex with necessary inputs, for any sample set S of size N , we also have a (ε, δ) -DP algorithm \mathcal{A}' which can get a solution ω_T such that

$$\mathbb{E}[\widehat{H}(\omega_T) - \widehat{H}^*] = O\left(\frac{GD \sqrt{d \log(1/\delta)}}{\varepsilon N}\right).$$

where $\widehat{H}(\omega) = \frac{1}{N} \sum_{x_i \in S} h(\omega, x_i)$, $\widehat{H}^* = \min_{\omega \in \mathcal{K}} H(\omega)$. The gradient complexity and privacy guarantee of \mathcal{A} and \mathcal{A}' are the same.

Moreover, the reduction also holds in the context of SCO.

Proof. We only consider this lemma in the context of ERM, as we can use the nearly the same argument for SCO.

The proof of this reduction is rather simple: After getting $\{h(\cdot, x_i)\}_{x_i \in S}$, we only need to consider $h_u(\omega, x) = h(\omega, x) + u\|\omega\|^2$. Then $h_u(\cdot, x)$ is u -strongly convex and $O(G + uD)$ -Lipschitz for any x with $\|x\|_2 \leq 2D$.

For the case $uD \leq G$, we run \mathcal{A} on $\{h_u(\cdot, x_i)\}_{x_i \in S}$ to get a solution ω_T with loss

$$\mathbb{E}[H_u(\omega_T) - H_u^*] = O\left(\frac{G^2 d \log(1/\delta)}{u \varepsilon^2 N^2}\right),$$

where $H_u(\omega) = \frac{1}{N} \sum_{x_i \in S} h(\omega, x_i) + u\|\omega\|^2$ and $H_u^* = \min_{\omega \in \mathcal{K}} H_u(\omega)$. Now by setting $u = \Theta\left(\frac{G \sqrt{d \log(1/\delta)}}{D \varepsilon N}\right)$, one has

$$\begin{aligned} \mathbb{E}[\widehat{H}(\omega_T) - \widehat{H}^*] &= O\left(\frac{G^2 d \log(1/\delta)}{u \varepsilon^2 N^2} + uD^2\right) \\ &= O\left(\frac{GD \sqrt{d \log(1/\delta)}}{\varepsilon N}\right). \end{aligned}$$

For the case $uD \geq G$, we have $\frac{GD \sqrt{d \log(1/\delta)}}{\varepsilon N} \geq GD$ and hence we can simply output the initial point ω_0 as the solution with a loss no more than GD . \square

The above reduction completes the main result of this subsection.

Theorem 4.11 (General Convex Case for Theorem 1.1). *Suppose $\mathcal{K} \subset \mathbb{R}^d$ is a convex set of diameter D and $\{f(\cdot, x)\}_{x \in \Xi}$ is a family of G -Lipschitz and convex functions over \mathcal{K}_r where $r = \frac{D \sqrt{d \log(1/\delta)}}{\varepsilon N}$. For $\varepsilon, \delta \leq 1/2$, given any sample set S consists of N samples from Ξ and arbitrary initial point $\omega_0 \in \mathcal{K}$, we have a (ε, δ) -differentially private algorithm \mathcal{A} which takes $O\left(\frac{\varepsilon N^{\frac{3}{2}}}{d^{1/8} \log^{1/4}(1/\delta)} + \frac{\varepsilon^2 N^2}{d \log(1/\delta)}\right)$ gradient queries and outputs ω_T such that*

$$\mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}^*] = O\left(\frac{GD \sqrt{d \log(1/\delta)}}{\varepsilon N}\right),$$

where $\widehat{F}^* = \min_{\omega \in \mathcal{K}} \widehat{F}(\omega)$, and the expectation is taken over the randomness of the algorithm.

5 Differentially Private SCO

In this section we study SCO. We can use the iterative localization technique in [FKT20] to reduce the SCO problem to an ERM problem. More specifically, if we can solve private ERM and get a (nearly) optimal empirical loss, then we can solve private SCO with (nearly) optimal excess population loss with the following algorithm framework (Algorithm 5). See Theorem 5.1 for the corresponding formal statement.

Algorithm 5: Iterative Localized Algorithm Framework \mathcal{A}'

```

1 Input: A family of  $G$ -Lipschitz and  $\mu$ -strongly convex function  $f : \mathcal{K} \times \Xi \rightarrow \mathbb{R}$ , initial
  point  $\omega_0 \in \mathcal{K}$  and privacy parameter  $\varepsilon, \delta$ .
2 Process: Set  $k = \lceil \log N \rceil$ ;
3 for  $i = 1, \dots, k$  do
4   Set  $\varepsilon_i = \varepsilon/2^i, N_i = N/2^i, \eta_i = \eta/2^{5i}$ ;
5   Apply  $(\varepsilon_i, \delta_i)$ -DP ERM algorithm  $\mathcal{A}_{\varepsilon_i, \delta_i}$  over  $\mathcal{K}_i = \{\omega \in \mathcal{K} : \|\omega - \omega_{i-1}\|_2 \leq 2G\eta_i N_i\}$ 
    with the function  $\widehat{F}_i(\omega) = \frac{1}{N_i} \sum_{j \in S_i} f(\omega, x_j) + \frac{1}{\eta_i N_i} \|\omega - \omega_{i-1}\|^2$  where  $S_i$  consists of  $N_i$ 
    samples with replacement from  $\mathcal{P}$ ;
6   Let  $\omega_i$  be the output of the ERM algorithm;
7 end
8 Return: The final iterate  $\omega_k$ ;

```

Theorem 5.1. Suppose we have an algorithm \mathcal{A} which can solve ERM under strongly convex case and gets a solution with excess empirical loss $O(\frac{G^2}{\mu N})$ by using $g(N)$ many gradients, then we have an algorithm \mathcal{A}' which can solve SCO under general case and gets a solution with excess population loss $O(\frac{GD}{\sqrt{N}})$ by using $\sum_{i=1}^{\lceil \log N \rceil} g(N/2^i)$ many gradients.

Moreover, for $\varepsilon, \delta \leq 1/2$, if $\mathcal{A}_{\varepsilon, \delta}$ is (ε, δ) -differentially private with excess empirical loss $O(\frac{G^2}{\mu}(\frac{1}{N} + \frac{d \log(1/\delta)}{\varepsilon^2 N^2}))$ under the strongly convex case by using $g(N, \varepsilon, \delta)$ many gradients, then we can get (ε, δ) -differentially private \mathcal{A}' with excess population loss $O(GD(\frac{1}{\sqrt{N}} + \frac{\sqrt{d \log(1/\delta)}}{\varepsilon N}))$ by querying gradients at most $\sum_{i=1}^{\lceil \log N \rceil} g(N/2^i, \varepsilon/2^i, \delta/2^i)$ times.

We only prove the bound with privacy guarantee, as the (non-private) bound can be proved with similar argument. Two technical lemmas will be proved at first, after which we will complete the proof.

Lemma 5.2. Let $\widehat{\omega}_i = \arg \min_{\omega \in \mathcal{K}} \widehat{F}_i(\omega)$, then

$$\mathbb{E}[\|\omega_i - \widehat{\omega}_i\|_2^2] \leq O\left(\frac{G^2 \eta_i^2 d \log(1/\delta_i)}{\varepsilon_i^2} + G^2 \eta_i^2 N_i\right).$$

Proof. At first, we prove that $\widehat{\omega}_i \in \mathcal{K}_i$. The definition of $\widehat{\omega}_i$ implies that

$$\frac{1}{N_i} \sum_{j=1}^{N_i} f(\widehat{\omega}_i, x_j) + \frac{1}{\eta_i N_i} \|\widehat{\omega}_i - \omega_{i-1}\|_2^2 \leq \frac{1}{N_i} \sum_{j=1}^{N_i} f(\omega_{i-1}, x_j).$$

Then we know that

$$\frac{1}{\eta_i N_i} \|\widehat{\omega}_i - \omega_{i-1}\|_2^2 \leq G \|\widehat{\omega}_i - \omega_{i-1}\|_2,$$

which implies $\widehat{\omega}_i \in \mathcal{K}_i$.

Next, note that \widehat{F}_i is $\lambda_i = \frac{1}{\eta_i N_i}$ -strongly convex, by the guarantee of our ERM algorithm, we know that

$$\begin{aligned} \frac{\lambda_i}{2} \mathbb{E}[\|\widehat{\omega}_i - \omega_i\|_2^2] &\leq \mathbb{E}[\widehat{F}_i(\widehat{\omega}_i) - \widehat{F}_i(\omega_i)] \\ &\leq O\left(\frac{G^2 d \log(1/\delta_i)}{\lambda_i \varepsilon_i^2 N_i^2} + \frac{G^2}{\lambda_i N_i}\right) \\ &= O\left(\frac{G^2 \eta_i d \log(1/\delta_i)}{\varepsilon_i^2 N_i} + G^2 \eta_i\right), \end{aligned}$$

which implies

$$\mathbb{E}[\|\widehat{\omega}_i - \omega_i\|_2^2] \leq O\left(\frac{G^2 \eta_i^2 d \log(1/\delta_i)}{\varepsilon_i^2} + G^2 \eta_i^2 N_i\right).$$

□

Lemma 5.3. *For any $y \in \mathcal{K}$, we know that*

$$\mathbb{E}[F(\widehat{\omega}_i) - F(y)] \leq \frac{\mathbb{E}[\|\omega_{i-1} - y\|_2^2]}{\eta_i N_i} + O(G^2 \eta_i).$$

Proof. Let $r(\omega, x) = f(\omega, x) + \frac{1}{\eta_i N_i} \|\omega - \omega_{i-1}\|_2^2$, $R(\omega) = \mathbb{E}_{x \sim \mathcal{P}} r(\omega, x)$ and $y^* = \arg \min_{\omega \in \mathcal{K}} R(\omega)$. By Theorem 6 in [SSSSS09], one has that

$$\begin{aligned} \mathbb{E}[R(\widehat{\omega}_i) - R(y)] &= \mathbb{E}[F(\widehat{\omega}_i) + \frac{1}{\eta_i N_i} \|\widehat{\omega}_i - \omega_{i-1}\|_2^2 - F(y) - \frac{1}{\eta_i N_i} \|y - \omega_{i-1}\|_2^2] \\ &\leq \mathbb{E}[R(\widehat{\omega}_i) - R(y^*)] \\ &\leq O(G^2 \eta_i), \end{aligned}$$

which implies that

$$\begin{aligned} \mathbb{E}[F(\widehat{\omega}_i) - F(y)] &\leq O(G^2 \eta_i) - \frac{1}{\eta_i N_i} \mathbb{E}[\|\widehat{\omega}_i - \omega_{i-1}\|_2^2] + \frac{1}{\eta_i N_i} \mathbb{E}[\|y - \omega_{i-1}\|_2^2] \\ &\leq O(G^2 \eta_i) + \frac{1}{\eta_i N_i} \mathbb{E}[\|y - \omega_{i-1}\|_2^2]. \end{aligned}$$

□

Having these two lemmas, we can begin the proof.

Proof of Theorem 5.1. The privacy guarantee comes directly from the basic composition theorem (See Theorem 2.7).

Let $S_i = \{x_j\}_{N-N/2^{i-1} \leq j \leq N-N/2^i}$. Let $N_i = N/2^i$, $\varepsilon_i = \varepsilon/2^i$, $\delta_i = \delta/2^i$ and $\eta_i = \eta/2^{5i}$ where η will be defined soon. For $i \in [k]$, let $\widehat{F}_i(\omega) = \sum_{x_j \in S_i} f(\omega, x_j) + \frac{1}{\eta_i N_i} \|\omega - \omega_{i-1}\|_2^2$.

Let $\widehat{\omega}_0 = \omega^*$, we have

$$\mathbb{E}[F(\omega_k)] - F(\omega^*) = \sum_{i=1}^k \mathbb{E}[F(\widehat{\omega}_i) - F(\widehat{\omega}_{i-1})] + \mathbb{E}[F(\omega_k) - F(\widehat{\omega}_k)].$$

First, Lemma 5.2 implies that

$$\begin{aligned}
\mathbb{E}[F(\omega_k) - F(\widehat{\omega}_k)] &\leq O(G\sqrt{\mathbb{E}[\|\omega_k - \widehat{\omega}_k\|_2^2]}) \\
&\leq O\left(\frac{G^2\eta_k\sqrt{d\log(1/\delta_k)}}{\varepsilon_k} + G^2\eta_k\sqrt{N_k}\right) \\
&= O\left(\frac{G^2\eta\sqrt{d\log(N/\delta)}}{\varepsilon N^3} + \frac{G^2\eta}{N^4}\right),
\end{aligned}$$

which is negligible.

Then one has

$$\begin{aligned}
\sum_{i=1}^k \mathbb{E}[F(\widehat{\omega}_i) - F(\widehat{\omega}_{i-1})] &\leq \sum_{i=1}^k \frac{\mathbb{E}[\|\widehat{\omega}_{i-1} - \omega_{i-1}\|_2^2]}{\eta_i N_i} + O(G^2\eta_i) \\
&\leq O\left(\frac{D^2}{\eta N} + \eta G^2 + \sum_{i=2}^k \left(\frac{G^2\eta_i d \log(1/\delta_i)}{\varepsilon_i^2 N_i} + G^2\eta_i\right)\right) \\
&\leq O\left(\frac{D^2}{\eta N} + \eta G^2 + \frac{G^2\eta d \log(1/\delta)}{\varepsilon^2 N}\right).
\end{aligned}$$

By setting $\eta = \frac{D}{G} \cdot \min\left\{\frac{1}{\sqrt{N}}, \frac{\varepsilon}{\sqrt{d\log(1/\delta)}}\right\}$, we get the excess population loss:

$$\mathbb{E}[F(\widehat{\omega}_k) - F(\omega^*)] = O(GD\left(\frac{1}{\sqrt{N}} + \frac{\sqrt{d\log(1/\delta)}}{N\varepsilon}\right)).$$

As for the gradient complexity, as we use $g(N_i, \varepsilon_i, \delta_i)$ queries of gradients in i -th iteration, the total gradient complexity is $\sum_{i=1}^k g(N_i, \varepsilon_i, \delta_i)$ as claimed. \square

Note that Theorem 5.1 allows the ERM algorithm has an extra $G^2/(\mu N)$ loss. This allows us to design a faster ERM algorithm compared Theorem 4.9 by choosing a different set of parameters.

Lemma 5.4. *Under the assumption defined in Algorithm Private AC-SA, with*

$$O\left(N + \min\{\sqrt{\varepsilon}N^{5/4}d^{1/8}, \frac{\varepsilon N^{3/2}}{d^{1/8}\log^{1/4}(1/\delta)}\}\right)$$

gradient complexity, one can get a solution ω_T such that

$$\mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}^*] = O\left(\frac{G^2}{\mu}\left(\frac{1}{N} + \frac{d\log(1/\delta)}{\varepsilon^2 N^2}\right)\right).$$

Proof. By Lemma 4.5, one has

$$\mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}^*] = O\left(\frac{G^2/B + \sigma^2 d/B^2}{\mu T} + \frac{GDd^{1/4}}{T}\right).$$

Again, setting $\sigma = \frac{c_2 GB \sqrt{T \log(1/\delta)}}{\varepsilon N}$ one has

$$\mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}^*] = O\left(\frac{G^2}{\mu BT} + \frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2} + \frac{GDd^{1/4}}{T}\right).$$

Taking $T = 400 \lceil \min\{N^{1/2}d^{1/4}, \frac{N\varepsilon}{d^{1/4}\sqrt{\log(1/\delta)}}\} \rceil$ and using $BT \geq N$ (which we will ensure), we have

$$\begin{aligned}\mathbb{E}[\widehat{F}(\omega_T) - \widehat{F}^*] &= O\left(\frac{G^2}{\mu N} + \frac{G^2 d \log(1/\delta)}{\mu \varepsilon^2 N^2} + \frac{GD}{\sqrt{N}} + \frac{GD\sqrt{d \log(1/\delta)}}{\mu \varepsilon N}\right) \\ &= O\left(\frac{G^2}{\mu} \zeta + GD\sqrt{\zeta}\right).\end{aligned}$$

where $\zeta = \frac{1}{N} + \frac{d \log(1/\delta)}{\varepsilon^2 N^2}$.

To ensure that Private AC-SA is (ε, δ) -DP, we set $B = \lceil N/T + N\sqrt{\varepsilon/T} \rceil = \Theta(N/T + N\sqrt{\varepsilon/T})$. By our choice of T , we have $B \leq N/10$ and $\varepsilon \leq c_1 B^2 T / N^2$. Hence, we can apply Lemma 4.4 to conclude (ε, δ) -DP.

Hence, we have a (ε, δ) -DP for ERM with loss $O\left(\frac{G^2}{\mu} \zeta + GD\sqrt{\zeta}\right)$ with $\zeta = \frac{1}{N} + \frac{d \log(1/\delta)}{\varepsilon^2 N^2}$. Note however that Theorem 5.1 requires us to have a DP-ERM algorithm with loss $O(\frac{G^2}{\mu} \zeta)$, namely, we have the extra term $O(GD\sqrt{\zeta})$. To remove this term, we follow the reduction in Lemma 4.7.

We note that the exact same proof as Lemma 4.7 shows that for any $\zeta > 0$, if we can solve strongly ERM with loss $O(G^2 \zeta^2 / \mu + GD\zeta)$, then we can solve strongly ERM with loss $O(G^2 \zeta^2 / \mu)$ by using the same number of gradient. This completes the proof. \square

Before stating our result on SCO, we need the following variant of Lemma 4.7. The proof is essentially the same, we state it for future reference.

Lemma 5.5 (Reduction to General Convex Case). *Given F is G -Lipschitz and μ -strongly convex. Suppose for any $\varepsilon, \delta < 1/2$, we have an (ε, δ) -differentially private algorithm \mathcal{A} which takes ω_0 as the initial start point and N samples i.i.d drawn from some distribution \mathcal{P} , and outputs a solution ω_T such that*

$$\mathbb{E}[F(\omega_T) - F^*] = O\left(\frac{G^2}{\mu}\left(\frac{1}{N} + \frac{d \log(1/\delta)}{N^2 \varepsilon^2}\right) + GD\left(\frac{1}{\sqrt{N}} + \frac{\sqrt{d \log(1/\delta)}}{N \varepsilon}\right)\right),$$

where $\omega^* = \arg \min_{\omega \in \mathcal{K}} F(\omega)$ and $D = \|\omega_0 - \omega^*\|$. Then by taking \mathcal{A} as sub-procedure with some modifications on parameters, we can get an (ε, δ) -differentially private solution with excess population loss at most

$$\mathbb{E}[F(\omega_T) - F^*] = O\left(\frac{G^2}{\mu}\left(\frac{1}{N} + \frac{d \log(1/\delta)}{N^2 \varepsilon^2}\right)\right).$$

Furthermore, if \mathcal{A} uses $g(N, \varepsilon, \delta)$ many gradients, the new algorithm uses $\sum_{i \geq 1} g(N/2^i, \varepsilon/2^i, \delta/2^i)$ many gradients.

Proof. The only difference to Lemma 4.7 is that this algorithm takes $N/2^{k+1-i}$ samples instead of N samples in the i -th step for $k = \lceil \log \log N^3 \rceil$, so it may have less gradient complexity. The rest of the proof is identical. \square

Now, we can get the result for general convex case by Theorem 5.1 and Lemma 5.4, then extend it to strongly convex case by Lemma 5.5.

Theorem 5.6 (DP-SCO, Theorem 1.2 restated). Suppose $\varepsilon, \delta \leq 1/2$. Let $\{f(\cdot, x)\}_{x \in \Xi}$ is convex and G -Lipschitz with respect to ℓ_2 norm and convex over \mathcal{K}_r , where $r = \frac{D\sqrt{d\log(1/\delta)}}{\varepsilon N}$, there is an (ε, δ) -differentially private algorithm which takes

$$O(N + \min\{\sqrt{\varepsilon}N^{5/4}d^{1/8}, \frac{\varepsilon N^{3/2}}{d^{1/8}\log^{1/4}(1/\delta)}\})$$

gradient queries to get a solution ω_T

$$\mathbb{E}[F(\omega_T) - F(\omega^*)] = O(GD(\frac{1}{\sqrt{N}} + \frac{\sqrt{d\log(1/\delta)}}{N\varepsilon})).$$

Moreover, if $\{f(\cdot, x)\}_{x \in \Xi}$ is also μ -strongly convex over \mathcal{K}_r , we can use the same gradient complexity and get a solution ω such that:

$$\mathbb{E}[F(\omega_T) - F(\omega^*)] = O\left(\frac{G^2}{\mu}(\frac{d\log(1/\delta)}{\varepsilon^2 N^2} + \frac{1}{N})\right).$$

References

- [ACG⁺16] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318, 2016.
- [AFKT21] Hilal Asi, Vitaly Feldman, Tomer Koren, and Kunal Talwar. Private stochastic convex optimization: Optimal rates in ℓ_1 geometry. *arXiv preprint arXiv:2103.01516*, 2021.
- [BDRS18] Mark Bun, Cynthia Dwork, Guy N Rothblum, and Thomas Steinke. Composable and versatile privacy via truncated cdp. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 74–86, 2018.
- [BFGT20] Raef Bassily, Vitaly Feldman, Cristóbal Guzmán, and Kunal Talwar. Stability of stochastic gradient descent on nonsmooth convex losses. *arXiv preprint arXiv:2006.06914*, 2020.
- [BFTT19] Raef Bassily, Vitaly Feldman, Kunal Talwar, and Abhradeep Guha Thakurta. Private stochastic convex optimization with optimal rates. In *Advances in Neural Information Processing Systems*, pages 11282–11291, 2019.
- [BGN21] Raef Bassily, Cristóbal Guzmán, and Anupama Nandi. Non-euclidean differentially private stochastic convex optimization. *arXiv preprint arXiv:2103.01278*, 2021.
- [BST14] Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 464–473. IEEE, 2014.
- [CBCG04] Nicolo Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.
- [CM08] Kamalika Chaudhuri and Claire Monteleoni. Privacy-preserving logistic regression. In *NIPS*, volume 8, pages 289–296. Citeseer, 2008.
- [CMS11] Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(3), 2011.
- [DBW12] John C Duchi, Peter L Bartlett, and Martin J Wainwright. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.
- [DR14] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- [Fel16] Vitaly Feldman. Generalization of erm in stochastic convex optimization: The dimension strikes back. *arXiv preprint arXiv:1608.04414*, 2016.

[FKT20] Vitaly Feldman, Tomer Koren, and Kunal Talwar. Private stochastic convex optimization: optimal rates in linear time. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 439–449, 2020.

[FMTT18] Vitaly Feldman, Ilya Mironov, Kunal Talwar, and Abhradeep Thakurta. Privacy amplification by iteration. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 521–532. IEEE, 2018.

[FTS17] Kazuto Fukuchi, Quang Khai Tran, and Jun Sakuma. Differentially private empirical risk minimization with input perturbation. In *International Conference on Discovery Science*, pages 82–90. Springer, 2017.

[GL12] Saeed Ghadimi and Guanghui Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization i: A generic algorithmic framework. *SIAM Journal on Optimization*, 22(4):1469–1492, 2012.

[HRS16] Moritz Hardt, Ben Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine Learning*, pages 1225–1234. PMLR, 2016.

[INS⁺19] Roger Iyengar, Joseph P Near, Dawn Song, Om Thakkar, Abhradeep Thakurta, and Lun Wang. Towards practical differentially private convex optimization. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 299–316. IEEE, 2019.

[JNN19] Prateek Jain, Dheeraj Nagaraj, and Praneeth Netrapalli. Making the last iterate of sgd information theoretically optimal. In *Conference on Learning Theory*, pages 1752–1755. PMLR, 2019.

[JT14] Prateek Jain and Abhradeep Guha Thakurta. (near) dimension independent risk bounds for differentially private learning. In *International Conference on Machine Learning*, pages 476–484. PMLR, 2014.

[KJ16] Shiva Prasad Kasiviswanathan and Hongxia Jin. Efficient private empirical risk minimization for high-dimensional learning. In *International Conference on Machine Learning*, pages 488–497. PMLR, 2016.

[KST12] Daniel Kifer, Adam Smith, and Abhradeep Thakurta. Private convex empirical risk minimization and high-dimensional regression. In *Conference on Learning Theory*, pages 25–1. JMLR Workshop and Conference Proceedings, 2012.

[Mir17] Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.

[Nes05] Yu Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.

[RBHT09] Benjamin IP Rubinstein, Peter L Bartlett, Ling Huang, and Nina Taft. Learning in a large function space: Privacy-preserving mechanisms for svm learning. *arXiv preprint arXiv:0911.5708*, 2009.

[Rén61] Alfréd Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. The Regents of the University of California, 1961.

[SSSS09] Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Stochastic convex optimization. In *COLT*, 2009.

[TZ14] Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Private empirical risk minimization beyond the worst case: The effect of the constraint set geometry. *arXiv preprint arXiv:1411.5417*, 2014.

[TZ15] Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Nearly-optimal private lasso. In *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 2*, pages 3025–3033, 2015.

[WLK⁺17] 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 *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 1307–1322, 2017.

[WYX17] Di Wang, Minwei Ye, and Jinhui Xu. Differentially private empirical risk minimization revisited: Faster and more general. In *Advances in Neural Information Processing Systems*, pages 2722–2731, 2017.

[YNS12] Farzad Yousefian, Angelia Nedić, and Uday V Shanbhag. On stochastic gradient and subgradient methods with adaptive steplength sequences. *Automatica*, 48(1):56–67, 2012.

[ZZMW17] Jiaqi Zhang, Kai Zheng, Wenlong Mou, and Liwei Wang. Efficient private erm for smooth objectives. *arXiv preprint arXiv:1703.09947*, 2017.

A Proof of Theorem 3.1

As mentioned before, we can use the result in [BDRS18] to give a formal proof of our result. Before we start, let us define something necessary.

Definition A.1 (Truncated CDP). Let $\rho > 0$ and $\omega > 1$. A randomized algorithm $\mathcal{M} : \mathbb{N}^{|\Xi|} \rightarrow R$ satisfies ω -truncated ρ -concentrated differential privacy (or (ρ, ω) -tCDP) if for all neighboring S, S' that differ in a single entry,

$$\forall \alpha \in (1, \omega), D_\alpha(\mathcal{M}(S) \| \mathcal{M}(S')) \leq \rho\alpha,$$

where $D_\alpha(\cdot \| \cdot)$ denotes the Rényi divergence [Rén61] of order α (in nats, rather than bits).

Similar to classic differential privacy, tCDP also enjoys a property of composition:

Lemma A.2 (Composition of tCDP). *Let $\mathcal{M}_1 : \mathbb{N}^{|\Xi|} \rightarrow R_1$ satisfy (ρ, ω) -tCDP and let $\mathcal{M}_2 : \mathbb{N}^{|\Xi|} \times R_1 \rightarrow R_2$ satisfy (ρ', ω') -tCDP for all $y \in R_1$. Define $\mathcal{M} : \mathbb{N}^{|\Xi|} \rightarrow R_3$ by $\mathcal{M}(S) = \mathcal{M}_2(S, \mathcal{M}_1(S))$. Then \mathcal{M} satisfies $(\rho + \rho', \min\{\omega, \omega'\})$ -tCDP.*

Now we state the main result of [BDRS18]:

Theorem A.3 (Privacy Amplification By Subsampling). *Let $\rho, s \in (0, 0.1]$ and $B, N \in \mathbb{N}$ with $q = B/N$ and $\log(1/q) \geq 3\rho(2 + \log_2(1/\rho))$. Let $\mathcal{M} : \mathbb{N}^{|\Xi|} \rightarrow R$ satisfy (ρ, ω') -tCDP for $\omega' \geq \frac{\log(1/q)}{2\rho} \geq 3$. Define the mechanism $\mathcal{M}_q : \mathbb{N}^{|\Xi|} \rightarrow R$ by $\mathcal{M}_q(S) = \mathcal{M}(S_q)$ where $S_q \in \mathbb{N}^{|\Xi|}$ is the restriction of S to the entries specified by a uniformly random subset of size B .*

The algorithm \mathcal{M}_q satisfies $(13q^2\rho, \omega)$ -tCDP for

$$\omega = \frac{\log(1/q)}{4\rho}.$$

This theorem can apply to our algorithm META_{DP} directly, as we are using subsampling without replacement. More specifically, we are using subsampling Gaussian Mechanism, and for Gaussian Mechanism we have the following fact:

Fact A.4. *Let $P = \mathcal{N}(1, 1/2\rho)$ and $Q = \mathcal{N}(0, 1/2\rho)$. Then $D_\alpha(P \| Q) = \rho\alpha$ for all $\alpha \in (1, \infty)$. In other word, the Gaussian Mechanism with sensitive 1 satisfies (ρ, ∞) -tCDP.*

Now we can start our proof.

Proof of Theorem 3.1. For the t -th phase of META_{DP} , let $\mathcal{M}(S) = \sum_{x \in S} \nabla f(\omega_{t-1}, x) + v$ where $v \sim \mathcal{N}(0, \sigma^2 I_{d \times d})$. As we are considering G -Lipschitz function f , then we know that $\|\nabla(f)\|_2 \leq G$, which means that \mathcal{M} is (ρ, ∞) -tCDP where $\rho = G^2/(2\sigma^2)$.

Assume our parameters satisfy the precondition of Theorem A.3 first, then we know that the t -th phase of META_{DP} is $(13q^2\rho, 1/\rho)$ -tCDP. By the composition property (Lemma A.2), we know that META_{DP} is $(13Tq^2\rho, 1/\rho)$ -tCDP.

When $Tq^2\rho \cdot \frac{\log(1/\delta)}{\varepsilon} \leq O(\varepsilon)$ and $\frac{\log(1/\delta)}{\varepsilon} \leq O(\frac{1}{\rho})$, we know that META_{DP} is (ε, δ) -differentially private [BDRS18].

By setting $\sigma = \frac{c_2 GB \sqrt{T \log(1/\delta)}}{\varepsilon N}$, we have that $\rho = \frac{\varepsilon^2 N^2}{2c_2^2 B^2 T \log(1/\delta)}$. Together with the assumption $\varepsilon \leq c_1 B^2 T / N^2$, we have both $Tq^2\rho \cdot \frac{\log(1/\delta)}{\varepsilon} \leq O(\varepsilon)$ and $\frac{\log(1/\delta)}{\varepsilon} \leq O(\frac{1}{\rho})$ as claimed. This completes the proof. \square