

Private Convex Optimization in General Norms

Sivakanth Gopi* Yin Tat Lee† Daogao Liu‡ Ruoqi Shen§ Kevin Tian¶

Abstract

We propose a new framework for differentially private optimization of convex functions which are Lipschitz in an arbitrary norm $\|\cdot\|_{\mathcal{X}}$. Our algorithms are based on a regularized exponential mechanism which samples from the density $\propto \exp(-k(F + \mu r))$ where F is the empirical loss and r is a regularizer which is strongly convex with respect to $\|\cdot\|_{\mathcal{X}}$, generalizing a recent work of [GLL22] to non-Euclidean settings. We show that this mechanism satisfies Gaussian differential privacy and solves both DP-ERM (empirical risk minimization) and DP-SCO (stochastic convex optimization) by using localization tools from convex geometry. Our framework is the first to apply to private convex optimization in general normed spaces, and directly recovers non-private SCO rates achieved by mirror descent, as the privacy parameter $\epsilon \rightarrow \infty$. As applications, for Lipschitz optimization in ℓ_p norms for all $p \in (1, 2)$, we obtain the first optimal privacy-utility tradeoffs; for $p = 1$, we improve tradeoffs obtained by the recent works [AFKT21, BGN21] by at least a logarithmic factor. Our ℓ_p norm and Schatten- p norm optimization frameworks are complemented with polynomial-time samplers whose query complexity we explicitly bound.

*Microsoft Research, sigopi@microsoft.com

†University of Washington, yintat@uw.edu

‡University of Washington, dgliu@uw.edu

§University of Washington, shenr3@cs.washington.edu

¶Microsoft Research, tiankevin@microsoft.com. Work completed while at Stanford University.

1 Introduction

The study of convex optimization in spaces where the natural geometry is non-Euclidean, beyond being a natural question of independent interest, has resulted in many successes across algorithm design. A basic example of this is the celebrated multiplicative weights, or exponentiated gradient method [AHK12], which caters to the ℓ_1 geometry and has numerous applications in learning theory and algorithms. Moreover, optimization in real vector spaces equipped with different ℓ_p norms has found use in sparse recovery [CRT06], combinatorial optimization [KLOS14, KPSW19], multi-armed bandit problems [BC12], fair resource allocation [DFO20], and more (see e.g. [AKPS19, DG21] and references therein). Furthermore, optimization in Schatten- p norm geometries (the natural generalization of ℓ_p norms to matrix spaces) has resulted in improved algorithms for matrix completion [ANW10] and outlier-robust PCA [JLT20]. In addition to ℓ_p and Schatten- p norms, the theory of non-Euclidean geometries has been very useful in settings such as linear and semidefinite programming [Nem04] and optimization on matrix spaces [AGL⁺18], amongst others.

The main result of this paper is a framework for *differentially private* convex optimization in general normed spaces under a Lipschitz parameter bound. Differential privacy [DKM⁺06, DMNS06] has been adopted as the standard privacy notion for data analysis in both theory and practice, and differentially private algorithms have been deployed in many important settings in the industry as well as the U.S. census [EPK14, Abo16, Tea17, BEM⁺17, DKY17]. Consequently, differentially private optimization is an increasingly important and fundamental primitive in modern machine learning applications [BST14, ACG⁺16]. However, despite an extensive body of theoretical work providing privacy-utility tradeoffs (and more) for optimization in the Euclidean norm geometry, e.g. [CM08, CMS11, KST12, JT14, BST14, KJ16] (and many other follow-up works), more general settings have been left relatively unexplored. This state of affairs prohibits the application of private optimization theory to problems where the natural geometry is non-Euclidean. Recent works [AFKT21, BGN21, BGM21] have investigated special cases of private convex optimization, e.g. for ℓ_p spaces or polyhedral sets, under smoothness assumptions, or under structured losses. However, the systematic study of private convex optimization in general normed spaces in the most fundamental setting of Lipschitz losses has been left open, a gap that our work addresses.

Our framework for private convex optimization is simple: we demonstrate strong privacy-utility tradeoffs for a *regularized exponential mechanism* when optimizing a loss over a set $\mathcal{X} \subset \mathbb{R}^d$ equipped with a norm $\|\cdot\|_{\mathcal{X}}$. More concretely, our algorithms sample from densities

$$\propto \exp(-k(F_{\mathcal{D}} + \mu r))$$

where $k, \mu > 0$ are tunable parameters, $F_{\mathcal{D}}$ is a (data-dependent) empirical risk, and r is a strongly convex regularizer in $\|\cdot\|_{\mathcal{X}}$ with bounded range over \mathcal{X} . In the analogous non-private Lipschitz convex optimization setting, most theoretical developments (namely mirror descent frameworks) have focused on precise applications where such an r is readily available [Sha12, Bub15]. In this sense, our framework directly extends existing Lipschitz convex optimization theory to the private setting (and indeed, recovers existing non-private guarantees obtained by mirror descent [NY83]).

In the remainder of the introduction, we summarize our results (Section 1.1), highlight our technical contributions (Section 1.2), and situate our paper in the context of prior work (Section 1.3).

1.1 Our results

We study both the empirical risk minimization (ERM) problem and the stochastic convex optimization (SCO) problem in this paper; the goal in the latter case is to minimize the *population risk*. We formalize this under the following assumption, which parameterizes the space we are optimizing and the (empirical and population) risks we aim to minimize.

Assumption 1. *We make the following assumptions.*

- (1) *There is a compact, convex set $\mathcal{X} \subset \mathbb{R}^d$ equipped with a norm $\|\cdot\|_{\mathcal{X}}$.*
- (2) *There is a 1-strongly convex function $r : \mathcal{X} \rightarrow \mathbb{R}$ in $\|\cdot\|_{\mathcal{X}}$, and $\Theta \geq \max_{x \in \mathcal{X}} r(x) - \min_{x \in \mathcal{X}} r(x)$.*
- (3) *There is a set Ω such that for any $s \in \Omega$, there is a convex function $f(\cdot; s) : \mathcal{X} \rightarrow \mathbb{R}$ which is G -Lipschitz in $\|\cdot\|_{\mathcal{X}}$.*

For definitions used above, see Section 2. We remark that by strong convexity, the parameter Θ scales at least as $\Omega(D^2)$, where D is the diameter of \mathcal{X} with respect to $\|\cdot\|_{\mathcal{X}}$; in many cases of interest, we may upper bound Θ by $O(D^2)$ as well up to a logarithmic factor.

Finally, throughout the paper when working under Assumption 1, $\mathcal{D} = \{s_i\}_{i \in [n]}$ denotes a dataset drawn independently from \mathcal{P} , a distribution supported on Ω , and we define $F_{\mathcal{D}} : \mathcal{X} \rightarrow \mathbb{R}$ and $F_{\text{pop}} : \mathcal{X} \rightarrow \mathbb{R}$ by

$$F_{\mathcal{D}}(x) := \frac{1}{n} \sum_{i \in [n]} f(x; s_i), \quad F_{\text{pop}}(x) := \mathbb{E}_{s \sim \mathcal{P}}[f(x; s)]. \quad (1)$$

Private ERM and SCO. We first present the following general results under Assumption 1.

Theorem 1 (Informal, see Theorems 3, 4). *Under Assumption 1 and following notation (1), drawing a sample x from the density $\nu \propto \exp(-k(F_{\mathcal{D}} + \mu r))$ for some $k, \mu > 0$ specified in Theorem 3 is (ϵ, δ) -differentially private, and produces x such that*

$$\mathbb{E}_{x \sim \nu}[F_{\mathcal{D}}(x)] - \min_{x \in \mathcal{X}} F_{\mathcal{D}}(x) \leq G\sqrt{\Theta} \cdot \frac{\sqrt{8d \log \frac{1}{2\delta}}}{n\epsilon}.$$

Moreover, drawing a sample x from the density $\nu \propto \exp(-k(F_{\mathcal{D}} + \mu r))$ for some $k, \mu > 0$ specified in Theorem 4 is (ϵ, δ) -differentially private, and produces x such that

$$\mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n, x \sim \nu}[F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} F_{\text{pop}}(x) \leq G\sqrt{\Theta} \cdot \left(\frac{\sqrt{8d \log \frac{1}{2\delta}}}{n\epsilon} + \sqrt{\frac{8}{n}} \right).$$

Minimizing the non-private population risk under the same setting as Assumption 1 is a very well-studied problem, with matching upper and lower bounds in many cases of interest, such as ℓ_p norms [NY83, ABRW12, DJW14]. The population utility achieved by our regularized exponential mechanism in Theorem 1 (namely, as $\epsilon \rightarrow \infty$) matches the rate obtained by the classical mirror descent algorithm [NY83], which to our knowledge has not been previously observed. Finally, in Appendix A we provide an analog of Theorem 1 under the stronger assumption that the sample losses $f(\cdot; s)$ are strongly convex, bypassing the need for explicit regularization. Our results in Appendix A recover the optimal rate in the Euclidean case, matching known lower bounds [BST14].

We next show how to apply the results of Theorem 1 under various instantiations of Assumption 1 to derive new rates for private convex optimization under ℓ_p and Schatten- p norm geometries.

ℓ_p and Schatten- p norms. In Corollaries 2, 3, and 4, we combine known (optimal) uniform convexity estimates for ℓ_p spaces [BCL94] with the algorithms of Theorem 3 and 4 to obtain privacy-utility tradeoffs summarized in Table 1. Interestingly, we achieve all these bounds with a single algorithmic framework, which in all cases matches or partially matches known lower bounds.

We now contextualize our results with regard to the existing literature. In the following discussion, the “privacy-dependent” loss term is the term in the SCO loss scaling with ϵ, δ , and the “privacy-independent” loss term is the SCO loss when $\epsilon \rightarrow \infty$.

ℓ_p norm	Optimality gap	
	ERM loss $F_{\mathcal{D}}$	SCO loss F_{pop}
$p \in (1, 2)$ (★)	$GD \cdot \frac{\sqrt{d \log \frac{1}{2\delta}}}{n\epsilon\sqrt{p-1}}$	$GD \cdot \left(\frac{1}{\sqrt{n(p-1)}} + \frac{\sqrt{d \log \frac{1}{2\delta}}}{n\epsilon\sqrt{p-1}} \right)$
$p = 1$ (†)	$GD \cdot \frac{\sqrt{d \log d \log \frac{1}{2\delta}}}{n\epsilon}$	$GD \cdot \left(\sqrt{\frac{\log d}{n}} + \frac{\sqrt{d \log d \log \frac{1}{2\delta}}}{n\epsilon} \right)$
$p \geq 2$ (†)	$GD \cdot \frac{d^{1-\frac{1}{p}} \sqrt{\log \frac{1}{2\delta}}}{n\epsilon}$	$GD \cdot \left(\frac{d^{\frac{1}{2}-\frac{1}{p}}}{\sqrt{n}} + \frac{d^{1-\frac{1}{p}} \sqrt{\log \frac{1}{2\delta}}}{n\epsilon} \right)$

Table 1: Privacy-utility tradeoffs for ℓ_p norm optimization under (ϵ, δ) -differential privacy obtained by: Corollary 2 ($p \in (1, 2)$), Corollary 3 ($p = 1$), and Corollary 4 ($p \geq 2$). We assume \mathcal{X} has ℓ_p diameter bounded by D and hide constants (stated in the formal results) for brevity. (★) indicates that our result matches the private ERM and SCO lower bound [BGN21, LL22]. (†) indicates that our result (as $\epsilon \rightarrow \infty$) matches the non-private SCO lower bound [ABRW12, DJW14].

In the case of constant $p \in (1, 2)$, our Corollary 2 sharpens Theorem 5 of [AFKT21] by a $\sqrt{\log d}$ factor in the privacy-dependent loss term, and is the first to match lower bounds of [BGN21, LL22]. It improves bounds by [BGN21] by at least a $\log n$ factor on both parts of the SCO loss, which further loses an $n^{\frac{1}{4}}$ factor on the privacy-dependent loss and requires additional smoothness assumptions.

In the important case of $p = 1$, of fundamental interest due to its applications in sparse recovery [CRT06] as well as online learning [Sha12, AHK12], our Corollary 3 improves the privacy-dependent loss term of [AFKT21] by a $\log d$ factor, and matches the privacy-independent loss lower bound in the SCO literature [DJW14], matching the rate of entropic mirror descent. The privacy-dependent loss term incurs an additional overhead of $\sqrt{\log d}$ compared to existing lower bounds. However, just as lower bounds on the privacy-independent loss increase as $p \rightarrow 1$, it is plausible that the upper bound obtained by Corollary 3 is optimal, which we leave as an interesting open direction.

In the $p \geq 2$ case, prior work by [BGN21] obtains a rate matched by Corollary 4. The non-private population risk term in (15) is again known to be optimal [ABRW12]. We again find it an interesting open direction to close the gap between the upper bound (13) and known lower bounds for private convex optimization when $p \geq 2$, e.g. [BGN21, LL22].

We further demonstrate in Corollary 5 that all of these results have direct analogs in the case of optimization over matrix spaces equipped with Schatten norm geometries. To the best of our knowledge, this is the first such result in the private optimization literature; we believe this showcases the generality and versatility of our approach.

Finally, we mention that all of these results are algorithmic and achieved by samplers with polynomial query complexity and runtime, following developments of [LST21, GLL22]. In all cases, by simple norm equivalence relations, the query complexity of our samplers is at most a d factor worse than the ℓ_2 case, with improvements as $p \rightarrow 2$. It is an exciting direction to develop efficient high-accuracy samplers catering to structured densities relevant to the setups considered in this paper, e.g. those whose negative log-likelihoods are strongly convex in ℓ_p norms. The design of sampling algorithms for continuous distributions has been an area of intense research activity in the machine learning community, discussed in greater detail in Section 1.3. We mention here that our hope is that our results and general optimization framework serve as additional motivation for the pursuit of efficient structured sampling algorithms working directly in non-Euclidean geometries.

1.2 Our techniques

Our results essentially build on the recent work of [GLL22], who observed that a regularized exponential mechanism achieves optimal privacy-utility tradeoffs for empirical and population risks when losses are Lipschitz in the ℓ_2 norm. Under a Euclidean specialization of Assumption 1, [GLL22] provided variants of Theorem 1 using the regularizer $r(x) = \frac{1}{2} \|x\|_2^2$, i.e. reweighting by a Gaussian.

We demonstrate several key tools used in [GLL22] have non-Euclidean extensions by using a simple, general approach based on a convex geometric tool known as *localization*. For example, the starting point of our developments is relating the privacy curves of two nearby, strongly convex densities with the privacy curve of Gaussians (see Section 2 for definitions).

Theorem 2. *Let $\mathcal{X} \subset \mathbb{R}^d$ be compact and convex, let $F, \tilde{F} : \mathcal{X} \rightarrow \mathbb{R}$ be μ -strongly convex in $\|\cdot\|_{\mathcal{X}}$, and let $P \propto \exp(-F)$ and $Q \propto \exp(-\tilde{F})$. Suppose $\tilde{F} - F$ is G -Lipschitz in $\|\cdot\|_{\mathcal{X}}$. For all $\epsilon \in \mathbb{R}_{\geq 0}$,*

$$\delta(P \parallel Q)(\epsilon) \leq \delta\left(\mathcal{N}(0, 1) \parallel \mathcal{N}\left(\frac{G}{\sqrt{\mu}}, 1\right)\right)(\epsilon).$$

An analog of Theorem 2 when $\|\cdot\|_{\mathcal{X}}$ is the Euclidean norm was proven as Theorem 4.1 of [GLL22]. Moreover, the analog of Theorem 1 in [GLL22] follows from combining Theorem 4.1 of that work, and their Theorem 6.10, a reduction from the SCO problem to the ERM problem (containing a generalization error bound). These proofs in [GLL22] rely on powerful inequalities from probability theory, which were initially studied in the Gaussian (Euclidean norm regularizer) setting. For example, Theorem 4.1 applied the *Gaussian isoperimetric inequality* of [ST74, Bor75a] (see also Theorem 1.1, [Led99]), which states that strongly logconcave distributions in the Euclidean norm have expansion quality at least as good as a corresponding Gaussian. Moreover, the generalization error bound in Theorem 6.10 was proven based on a Euclidean norm *log-Sobolev inequality* and *transportation inequality*, relating Wasserstein distances, KL divergences, and Lipschitz bounds on negative log-densities. Fortunately, it turns out that all of these inequalities have non-Euclidean generalizations (possibly losing constant factors). For example, a non-Euclidean log-Sobolev inequality was shown by Proposition 3.1 of [BL00], and a non-Euclidean transport inequality sufficient for our purposes is proved as Proposition 1 of [CE17]. Finally, variants of the Gaussian isoperimetric inequality in general norms are given by [MS08, Kol11]. Plugging in these tools into the proofs of [GLL22] allows us to recover Theorems 1 and 2, as well as our applications.

In this work, we take a different (and in our opinion, simpler) strategy to proving the probability-theoretic inequalities required by Theorems 1, 2, yielding an alternative to the proofs in [GLL22] which we believe may be of independent intellectual interest to the privacy community. In particular, our technical insight is the simple observation that several of the definitions in differential privacy are naturally cast in the language of localization [KLS95, FG04], which characterizes extremal logconcave densities subject to linear constraints (see our proof of Lemma 2 for an example of this). This observation allows us to reduce the proofs of key technical tools used in Theorems 1 and 2 to proving these tools in one dimension, where *all norms are equivalent* up to constant factor rescaling.¹ After deriving several extensions of basic localization arguments in Section 3.1, we follow this reduction approach to give a more unified proof to Theorems 1 and 2. To our knowledge, this is the first direct application of localization techniques in differential privacy.

The interplay between the privacy and probability theory communities is an increasingly active area of exploration [DRS21, GLL22, GTU22] (discussed in more detail in Section 1.3). We are

¹The one-dimensional case can then typically be handled by more straightforward “combinatorial” arguments, see e.g. Section 2.b of [LS93] or Appendix B.3 of [CDWY20] for examples.

hence optimistic that localization-based proof strategies will have further applications in the privacy literature, especially in situations (beyond this paper) where probability theoretic tools used in the Euclidean case do not have non-Euclidean variants in full generality. In such settings, it may be a valuable endeavor to see if necessary inequalities may be directly recast in the language of localization.

1.3 Prior work

Private optimization in Euclidean norm. Many prior works on private convex optimization have focused on variants of the ERM and SCO problems studied in this work, under ℓ_2 Lipschitz losses and ℓ_2 bounded domains, such as [CMS11, KST12, BST14, BFTT19, BFGT20]. The optimal information-theoretic rate for these private optimization problems was given by [BST14], which was matched algorithmically up to constant factors by [BFTT19, BFGT20].

From an algorithmic perspective, a topic of recent interest in the Euclidean case is the problem of attaining optimal privacy-utility tradeoffs in *nearly-linear time*, namely, with $\approx n$ gradient queries [FKT20, AFKT21, KLL21]. Under additional smoothness assumptions, this goal was achieved by [FKT20]; however, achieving near-optimal gradient oracle query rates in the general Lipschitz case remains open. We note that under *value oracle* access, a near-optimal bound was recently achieved by [GLL22]. This paper primarily focuses on the information-theoretic problem of achieving optimal privacy-utility tradeoffs for a given dataset size. However, we believe the corresponding problem of designing algorithms with near-optimal query complexities and runtimes (under value or gradient oracle access) is also an important open direction in the case of general norm geometries.

Private optimization in non-Euclidean norms. The study of convex optimization in non-Euclidean geometries was recently initiated by [AFKT21, BGN21], who focused primarily on developing algorithms under ℓ_p regularity assumptions and bounded domains. In follow-up work, [BGM21] gave improved guarantees for the family of generalized linear losses. We discuss the rates we achieve for ℓ_p norm geometries compared to [AFKT21, BGN21] in Section 1.1; in short, we improve prior results by logarithmic factors in the case $p \in [1, 2)$, and match them when $p \geq 2$. Independently from our work, [HLL⁺22] designed an algorithm for private optimization in ℓ_p geometries improving upon [BGN21] in gradient query complexity (matching their privacy-utility tradeoffs); both [BGN21, HLL⁺22] require further smoothness assumptions on the loss functions.

One of the main motivations for this work was to develop a general theory for private convex optimization under non-Euclidean geometries, beyond ℓ_p setups. In particular, [BGN21] designed a *generalized Gaussian mechanism* for the case $p \in [1, 2)$, where gradients were perturbed by a noise distribution catering to the ℓ_p geometry. However, how to design a corresponding mechanism for more general norms may be less clear. The algorithm of [AFKT21] in the non-smooth case was based on a (Euclidean norm) Gaussian mechanism; again, this strategy is potentially more specialized to ℓ_p geometries. Beyond giving a general algorithmic framework for non-Euclidean convex optimization based on structured logconcave sampling, we hope that the information-theoretic properties we show regarding regularized exponential mechanisms (e.g. Theorem 2) may find use in designing “generalized Gaussian mechanisms” beyond ℓ_p norms.

Connections between privacy and sampling. Our work extends a line of work exploring privacy-utility tradeoffs for the exponential mechanism, a general strategy for designing private algorithms introduced by [MT07] (see additional discussion in [GLL22]). For example, the regularized exponential mechanisms we design are similar in spirit to the exponential mechanism “in the \mathcal{X} norm²” designed by [HT10, BDKT12]. Moreover, our work continues a recent interface between the sampling and privacy literature, where (continuous and discrete-time) sampling al-

²That is, the norm induced by the convex body \mathcal{X} , not to be confused with the $\|\cdot\|_{\mathcal{X}}$ of Assumption 1.

gorithms are shown to efficiently obtain strong privacy-utility tradeoffs for optimization problems [GLL22, GTU22]. This work further develops this interface, motivating the design of efficient samplers for densities satisfying non-Euclidean regularity assumptions.

The design of sampling algorithms under general geometries (e.g. “mirrored Langevin algorithms”) has been a topic of great recent interest, independently from applications in private optimization. Obtaining mixing guarantees under regularity assumptions naturally found in applications is a notoriously challenging problem in the recent algorithmic sampling literature [HKRC18, ZPFP20, AC21, Jia21, LTVW22]. For example, it has been observed both theoretically and empirically that without (potentially restrictive) relationships between regularity parameters, natural discretizations of the mirrored Langevin dynamics may not even have vanishing bias [ZPFP20, Jia21, LTVW22]. Recently, [LST21] gave an alternative strategy (to discretizing Langevin dynamics) for designing sampling algorithms in the Euclidean case, used in [GLL22] to obtain private algorithms for ℓ_2 -structured ERM and SCO problems (see Proposition 8). Our work suggests a natural non-Euclidean generalization of these sampling problems, which is useful to study from an algorithmic perspective. We are optimistic that a non-Euclidean variant of [LST21] may shed light on these mysteries and yield new efficient private algorithms. More generally (beyond the particular [LST21] framework), we state the direction of designing efficient samplers for densities of the form $\exp(-F_D - \mu r)$ satisfying Assumption 1 as an important open research endeavor with implications for both sampling and private optimization, the latter of which this paper demonstrates.

2 Preliminaries

General notation. Throughout, \tilde{O} hides logarithmic factors in problem parameters when clear from the context. For $n \in \mathbb{N}$, $[n]$ refers to the naturals $1 \leq i \leq n$. We use \mathcal{X} to denote a compact convex subset of \mathbb{R}^d . The standard (ℓ_2) Euclidean norm is denoted $\|\cdot\|_2$. We will be concerned with optimizing functions $f : \mathcal{X} \rightarrow \mathbb{R}$, and $\|\cdot\|_{\mathcal{X}}$ will refer to a norm on \mathcal{X} . The diameter of such a set is denoted $\text{diam}_{\|\cdot\|_{\mathcal{X}}}(\mathcal{X}) := \max_{x,y \in \mathcal{X}} \|x - y\|_{\mathcal{X}}$. We let $\mathcal{N}(\mu, \Sigma)$ be the Gaussian density of specified mean and covariance. We denote the convex hull of a set S (when well-defined) by $\text{conv}(S)$. When $a, b \in \mathbb{R}^d$, we abuse notation and let $[a, b]$ be the line segment between a and b .

Norms. For $p \geq 1$, we let $\|\cdot\|_p$ applied to a vector-valued variable be the ℓ_p norm, namely $\|v\|_p = (\sum_{i \in [d]} |v_i|^p)^{1/p}$ for $v \in \mathbb{R}^d$; the ℓ_∞ norm is the maximum absolute value. We will use the well-known inequality

$$\|v\|_p \leq \|v\|_q \leq d^{\frac{1}{q} - \frac{1}{p}} \|v\|_p, \text{ for } v \in \mathbb{R}^d, q \leq p. \quad (2)$$

Matrices will be denoted in boldface throughout, and $\|\cdot\|_p$ applied to a matrix-valued variable \mathbf{M} is the Schatten- p norm, i.e. the ℓ_p norm of the singular values of \mathbf{M} .

Optimization. In the following discussion, fix some $f : \mathcal{X} \rightarrow \mathbb{R}$. We say f is G -Lipschitz in $\|\cdot\|_{\mathcal{X}}$ if for all $x, x' \in \mathcal{X}$, $|f(x) - f(x')| \leq G \|x - x'\|_{\mathcal{X}}$. We say f is μ -strongly convex in $\|\cdot\|_{\mathcal{X}}$ if for all $x, x' \in \mathcal{X}$ and $t \in [0, 1]$,

$$f(tx + (1 - t)y) \leq tf(x) + (1 - t)f(y) - \frac{\mu t(1 - t)}{2} \|x - y\|_{\mathcal{X}}^2.$$

Probability. For two densities π, π' , we define their total variation distance by $\|\pi - \pi'\|_{\text{TV}} := \frac{1}{2} \int |\pi(x) - \pi'(x)| dx$ and (when the Radon-Nikodym derivative exists) their KL divergence by $D_{\text{KL}}(\pi \parallel \pi') := \int \pi(x) \log \frac{\pi(x)}{\pi'(x)} dx$. We define the 2-Wasserstein distance by

$$W_2(\pi, \pi') = \inf_{\Gamma \in \Gamma(\pi, \pi')} \sqrt{\mathbb{E}_{(x, x') \sim \Gamma} \|x - x'\|_2^2},$$

where $\Gamma(\pi, \pi')$ is the set of couplings of π and π' . We note W_2 satisfies the following inequality.

Fact 1. Let $\text{Lip}_2(f)$ be the Lipschitz constant in the ℓ_2 norm of a function f . Then, for densities π, π' supported on \mathcal{X} ,

$$W_2(\pi, \pi') \geq \sup_{\text{Lip}_2(f) \leq 1} \int_{\mathcal{X}} f(x)(\pi(x) - \pi'(x))dx.$$

Proof. This follows from the dual characterization of the 1-Wasserstein distance (which shows $\sup_{\text{Lip}(f) \leq 1} \int_{\mathcal{X}} f(x)(\pi(x) - \pi'(x))dx = \inf_{\Gamma \in \Gamma(\pi, \pi')} \mathbb{E}_{(x, x') \sim \Gamma} \|x - x'\|_2$, and convexity of the square. \square

We use \propto to indicate proportionality, e.g. if π is a density and we write $\pi \propto \exp(-f)$, we mean $\pi(x) = \frac{\exp(-f)}{Z}$ where $Z := \int \exp(-f(x))dx$ and the integration is over the support of π .

We say that a measure π on \mathbb{R}^d is logconcave if for any $\lambda \in (0, 1)$ and compact $A, B \subset \mathbb{R}^d$,

$$\pi(\lambda A + (1 - \lambda)B) \geq \pi(A)^\lambda \pi(B)^{1-\lambda}.$$

We have the following equivalent characterization of logconcave measures.

Proposition 1 ([Bor75b]). *Let π be a measure on \mathbb{R}^d . Let E be the least affine subspace containing the support of π , and let m_E be the Lebesgue measure in E . Then π is logconcave if and only if $d\pi = f dm_E$, f is nonnegative and locally integrable, and $-\log f : E \rightarrow \mathbb{R} \cup \{+\infty\}$ is convex.*

In particular, Proposition 1 shows that the measure of any subspace of E according to π is zero. If in the characterization of [Bor75b] the function $-\log f$ is affine, we say π is logaffine. Following [Bor75b], we analogously define strong logconcavity with respect to a norm.

Definition 1 (strong logconcavity). *Let π be a measure on \mathbb{R}^d . Let E be the least affine subspace containing the support of π , and let m_E be the Lebesgue measure in E . We say π is μ -strongly logconcave with respect to $\|\cdot\|_{\mathcal{X}}$ if $d\pi = f dm_E$, f is nonnegative and locally integrable, and $-\log f : E \rightarrow \mathbb{R} \cup \{+\infty\}$ is μ -strongly convex in $\|\cdot\|_{\mathcal{X}}$.*

Privacy. Throughout, \mathcal{M} denotes a mechanism, and \mathcal{D} denotes a dataset. We say \mathcal{D} and \mathcal{D}' are neighboring if they differ in one entry. We say a mechanism \mathcal{M} satisfies (ϵ, δ) -differential privacy if it has output space Ω and for any neighboring $\mathcal{D}, \mathcal{D}'$,

$$\sup_{S \subseteq \Omega} \Pr[\mathcal{M}(\mathcal{D}) \in S] - \exp(\epsilon) \Pr[\mathcal{M}(\mathcal{D}') \in S] \leq \delta.$$

We define the privacy curve of two random variables X, Y supported on Ω by

$$\delta(X \| Y)(\epsilon) := \sup_{S \subseteq \Omega} \Pr[Y \in S] - \exp(\epsilon) \Pr[X \in S].$$

We say \mathcal{M} has a privacy curve $\delta : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ if for all neighboring $\mathcal{D}, \mathcal{D}'$, $\delta(\mathcal{M}(\mathcal{D}) \| \mathcal{M}(\mathcal{D}')) \leq \delta(\epsilon)$. For any $\epsilon \in \mathbb{R}_{\geq 0}$, it is clear that such a \mathcal{M} is $(\epsilon, \delta(\epsilon))$ -differentially private. We will frequently compare to the privacy curve of a Gaussian, so we recall the following bound from prior work.

Fact 2 (Gaussian privacy curve, Lemma 6.3, [GLL22]). *Let $\delta \in (0, \frac{1}{2})$ and $\epsilon > 0$. For any $|t| \leq \sqrt{2 \log \frac{1}{2\delta} + 2\epsilon} - \sqrt{2 \log \frac{1}{2\delta}} \leq \frac{\epsilon}{\sqrt{2 \log \frac{1}{2\delta}}}$, $\delta(\mathcal{N}(0, 1) \| \mathcal{N}(t, 1))(\epsilon) \leq \delta$.*

We will use Fact 2 after deriving our Gaussian differential privacy guarantees [DRS21] for strongly logconcave densities in Theorem 2.

3 Gaussian differential privacy in general norms

In this section, we give a generalization of Theorem 4.1 of [GLL22], which demonstrates that a regularized exponential mechanism for (Euclidean norm) Lipschitz losses achieves privacy guarantees comparable to an analogous instance of the Gaussian mechanism. The proof from [GLL22] was specialized to the Euclidean setup; to show our more general result, we draw upon the localization technique from convex geometry [LS93, KLS95]. We provide the relevant localization tools we will use in Section 3.1, and prove our Gaussian differential privacy result in Section 3.2.

3.1 Localization

We recall the localization lemma from [FG04]. We remark that the statement in [FG04] is more refined than our statement (in that [FG04] gives a complete characterization of extreme points, whereas we give a superset), but the following form of the [FG04] result suffices for our purposes.

Proposition 2 (Theorem 1, [FG04]). *Let $\mathcal{X} \subset \mathbb{R}^d$ be compact and convex, and let $f : \mathcal{X} \rightarrow \mathbb{R}$ be upper semi-continuous. Let $\mathcal{S}(f)$ be the set of logconcave densities ν supported in \mathcal{X} satisfying $\int_{\mathcal{X}} f d\nu \geq 0$. The set of extreme points of $\text{conv}(\mathcal{S}(f))$ satisfies one of the following.*

- ν is a Dirac measure at $x \in \mathcal{X}$ such that $f(x) \geq 0$.
- ν is logaffine and supported on $[a, b] \subset \mathcal{X}$ such that $\int_{[a, b]} f d\nu = 0$.

We next derive several extensions of Proposition 2.

Lemma 1 (Strongly logconcave localization). *Let $\mathcal{X} \subset \mathbb{R}^d$ be compact and convex, let $\beta : \mathcal{X} \rightarrow \mathbb{R}_{>0}$ be continuous, and let $f : \mathcal{X} \rightarrow \mathbb{R}$ be upper semi-continuous. Let $\mathcal{S}_{\mu, \beta}(f)$ be the set of probability densities π such that π is μ -strongly logconcave with respect to $\|\cdot\|_{\mathcal{X}}$ and supported in \mathcal{X} , such that $\pi' \propto \beta \pi$ is also μ -strongly logconcave, and $\int_{\mathcal{X}} f d\pi \geq 0$. The set of extreme points of $\text{conv}(\mathcal{S}_{\mu, \beta}(f))$ satisfy one of the following.*

- π is a Dirac measure at $x \in \mathcal{X}$ such that $f(x) \geq 0$.
- π is supported on $[a, b] \subset \mathcal{X}$.

Proof. Clearly, Dirac measures at x with $f(x) \geq 0$ are extreme points, so it suffices to consider other extreme points. Given any extreme point π which is not a Dirac measure, we prove the least affine subspace containing the support of π has dimension one, i.e. denoting the least affine subspace containing the support of π by S , we prove $\dim S = 1$.

Assume for the sake of contradiction that $\dim S \geq 2$. There exists x_0 in the relative interior of the support of π and a two-dimensional subspace $E \subset \mathbb{R}^d$ such that $x_0 + E \subseteq S$. Let $S_1(E)$ be the unit circle in E , and for any $u \in S_1(E)$ denote $H_u = \{x \in S : \langle x - x_0, u \rangle = 0\}$, $H_u^+ = \{x \in S : \langle x - x_0, u \rangle \geq 0\}$ and $H_u^- = \{x \in S : \langle x - x_0, u \rangle \leq 0\}$. Finally, define $\phi : S_1(E) \rightarrow \mathbb{R}$ by $\phi(u) := \int_{H_u^+} f d\pi - \frac{1}{2} (\int f d\pi)$, such that $\phi(u) = 0 \implies \int_{H_u^+} f d\pi = \frac{1}{2} \int f d\pi \geq 0$.

By Proposition 1, we know $\pi(H_u) = 0$. Moreover, ϕ is continuous since every hyperplane H_u has $\pi(H_u) = 0$. Since $\phi(u) = -\phi(-u)$, by the intermediate value theorem there exists $u_0 \in S_1(E)$ such that $\phi(u_0) = 0$. We can hence rewrite π as a convex combination of its restrictions to $H_{u_0}^+$ and $H_{u_0}^-$, both of which are μ -strongly logconcave, and whose (renormalized) multiplications by β are also μ -strongly logconcave. Since $\phi(u_0) = 0$ both of these restrictions belong to $\mathcal{S}_{\mu, \beta}(f)$, contradicting extremality of π . \square

We briefly remark that the proof technique used in Lemma 1 is quite general, and the only property we used about $\mathcal{S}_{\mu, \beta}$ is that it is a subset of logconcave densities, and it is closed under restrictions to convex subsets. Similar arguments hold for other density families with these properties.

Further, we note that restrictions to compact sets are upper semi-continuous; it is straightforward to verify our applications of Lemma 1 satisfy the upper semi-continuity assumption.

We prove the following two technical lemmas using Lemma 1.

Lemma 2. *Following notation of Lemma 1, fix a continuous function $\alpha : \mathcal{X} \rightarrow \mathbb{R}$ and a subset $S \subset \mathcal{X}$. For any probability density π on \mathcal{X} , define the renormalized density $\tilde{\pi} \propto e^{-\alpha}\pi$. Finally, let*

$$g(\pi) := \Pr_{x \sim \tilde{\pi}}[x \in S] - e^\epsilon \Pr_{x \sim \pi}[x \in S].$$

Then $\max_{\pi \in \mathcal{S}_{\mu,\beta}} g(\pi) = \max_{\pi \in \mathcal{S}_{\mu,\beta}^} g(\pi)$ where $\mathcal{S}_{\mu,\beta}^*$ is the subset of densities $\pi \in \mathcal{S}_{\mu,\beta}$ satisfying one of the following.*

- π is a Dirac measure at $x \in \mathcal{X}$.
- π is supported on $[a, b] \subset \mathcal{X}$.

Proof. Let $\mathcal{S}_{\mu,\beta}(f) \subseteq \mathcal{S}_{\mu,\beta}$ be the set of $\pi \in \mathcal{S}_{\mu,\beta}$ such that $\int f d\pi \geq 0$. We have

$$\begin{aligned} \max_{\pi \in \mathcal{S}_{\mu,\beta}} g(\pi) &= \max_{\pi \in \mathcal{S}_{\mu,\beta}} \int_{x \in S} d\tilde{\pi}(x) - e^\epsilon \int_{x \in S} d\pi(x) \\ &= \max_{\pi \in \mathcal{S}_{\mu,\beta}} \frac{\int_{x \in S} e^{-\alpha(x)} d\pi(x)}{\int_{x \in \mathcal{X}} e^{-\alpha(x)} d\pi(x)} - e^\epsilon \int_{x \in S} d\pi(x) \\ &= \max_{\pi \in \mathcal{S}_{\mu,\beta}} \max_{C \geq \int_{x \in \mathcal{X}} e^{-\alpha(x)} d\pi(x)} \frac{\int_{x \in S} e^{-\alpha(x)} d\pi(x)}{C} - e^\epsilon \int_{x \in S} d\pi(x) \\ &= \max_C \max_{\pi \in \mathcal{S}_{\mu,\beta}(C - e^{-\alpha})} \int_{x \in \mathcal{X}} \left(\frac{e^{-\alpha(x)}}{C} - e^\epsilon \right) \mathbf{1}_S(x) d\pi(x) \\ &= \max_C \max_{\pi \in \mathcal{S}_{\mu,\beta}(C - e^{-\alpha})^*} \int_{x \in \mathcal{X}} \left(\frac{e^{-\alpha(x)}}{C} - e^\epsilon \right) \mathbf{1}_S(x) d\pi(x), \end{aligned}$$

where $\mathcal{S}_{\mu,\beta}(C - e^{-\alpha})^*$ is the (super)set of extreme points of $\mathcal{S}_{\mu,\beta}(C - e^{-\alpha})$ given by the strongly logconcave localization lemma (Lemma 1). These candidate extreme points are Dirac measures at x such that $C \geq e^{-\alpha(x)}$, or are supported in $[a, b] \subset \mathcal{X}$. Hence, $\mathcal{S}_{\mu,\beta}(C - e^{-\alpha})^* \subseteq \mathcal{S}_{\mu,\beta}^*$, and we conclude

$$\max_{\pi \in \mathcal{S}_{\mu,\beta}} g(\pi) = \max_C \max_{\pi \in \mathcal{S}_{\mu,\beta}(C - e^{-\alpha})^*} \int_{x \in \mathcal{X}} \left(\frac{e^{-\alpha(x)}}{C} - e^\epsilon \right) \mathbf{1}_S(x) d\pi(x) \quad (3)$$

$$\leq \max_C \max_{\pi \in \mathcal{S}_{\mu,\beta}(C - e^{-\alpha})^*} \int_{x \in \mathcal{X}} \left(\frac{e^{-\alpha(x)}}{\int_{x \in \mathcal{X}} e^{-\alpha(x)} d\pi(x)} - e^\epsilon \right) \mathbf{1}_S(x) d\pi(x) \quad (4)$$

$$\leq \max_{\pi \in \mathcal{S}_{\mu,\beta}^*} \int_{x \in \mathcal{X}} \left(\frac{e^{-\alpha(x)}}{\int_{x \in \mathcal{X}} e^{-\alpha(x)} d\pi(x)} - e^\epsilon \right) \mathbf{1}_S(x) d\pi(x) = \max_{\pi \in \mathcal{S}_{\mu,\beta}^*} g(\pi). \quad (5)$$

The first inequality used that $C \geq \int_{x \in \mathcal{X}} e^{-\alpha(x)} d\pi(x)$ for $\pi \in \mathcal{S}_{\mu,\beta}(C - e^{-\alpha})^*$, and the second used that $\mathcal{S}_{\mu,\beta}(C - e^{-\alpha})^* \subseteq \mathcal{S}_{\mu,\beta}^*$ for any C . Since $\mathcal{S}_{\mu,\beta}^* \subseteq \mathcal{S}_{\mu,\beta}$, we have the claim. \square

Lemma 3. *Following notation of Lemma 1, fix continuous function $\alpha : \mathcal{X} \rightarrow \mathbb{R}$ and upper semi-continuous function $f : \mathcal{X} \rightarrow \mathbb{R}$. For any probability density π on \mathcal{X} , define $\tilde{\pi} \propto e^{-\alpha}\pi$ to be a renormalized density on \mathcal{X} . Finally, let*

$$g(\pi) := \int_{x \in \mathcal{X}} f(x) d(\pi - \tilde{\pi})(x).$$

Then $\max_{\pi \in \mathcal{S}_{\mu, \beta}} g(\pi) = \max_{\pi \in \mathcal{S}_{\mu, \beta}^*} g(\pi)$ where $\mathcal{S}_{\mu, \beta}^*$ is the subset of densities $\pi \in \mathcal{S}_{\mu, \beta}$ satisfying one of the following.

- π is a Dirac measure at $x \in \mathcal{X}$.
- π is supported on $[a, b] \subset \mathcal{X}$.

Proof. We follow the notation from Lemma 2. If π is a Dirac measure, $g(\pi) = 0$, so we only need to consider the case when $\max_{\pi \in \mathcal{S}_{\mu, \beta}} g(\pi) > 0$. We have

$$\begin{aligned} \max_{\pi \in \mathcal{S}_{\mu, \beta}} g(\pi) &= \max_{\pi \in \mathcal{S}_{\mu, \beta}} \int_{x \in \mathcal{X}} f(x) \left(1 - \frac{e^{-\alpha(x)}}{\int_{x \in \mathcal{X}} e^{-\alpha(x)} d\pi(x)} \right) d\pi(x) \\ &= \max_{\pi \in \mathcal{S}_{\mu, \beta}} \max_{C \leq \int_{x \in \mathcal{X}} e^{-\alpha(x)} d\pi(x)} \int_{x \in \mathcal{X}} \left(f(x) - \frac{e^{-\alpha(x)} f(x)}{C} \right) d\pi(x) \\ &= \max_{C} \max_{\pi \in \mathcal{S}_{\mu, \beta}(e^{-\alpha} - C)} \int_{x \in \mathcal{X}} \left(f(x) - \frac{e^{-\alpha(x)} f(x)}{C} \right) d\pi(x) \\ &= \max_{C} \max_{\pi \in \mathcal{S}_{\mu, \beta}(e^{-\alpha} - C)^*} \int_{x \in \mathcal{X}} \left(f(x) - \frac{e^{-\alpha(x)} f(x)}{C} \right) d\pi(x). \end{aligned}$$

The remainder of the proof is analogous to Lemma 2. \square

3.2 Gaussian differential privacy

Using an instantiation of the localization lemma, we prove Gaussian differential privacy in general norms by first reducing to one dimension and then using the result of [GLL22] to handle the one-dimensional case. Gaussian differential privacy was introduced by [DRS21] and is a useful tool to compare privacy curves. We first recall the (ℓ_2) Gaussian differential privacy result of [GLL22].

Proposition 3 (Theorem 4.1, [GLL22]). *Let $\mathcal{X} \subset \mathbb{R}^d$ be compact and convex, let $F, \tilde{F} : \mathcal{X} \rightarrow \mathbb{R}$ be μ -strongly convex in $\|\cdot\|_2$, and let $P \propto \exp(-F)$ and $Q \propto \exp(-\tilde{F})$. Suppose $\tilde{F} - F$ is G -Lipschitz in $\|\cdot\|_2$. For all $\epsilon \in \mathbb{R}_{\geq 0}$,*

$$\delta(P \parallel Q)(\epsilon) \leq \delta\left(\mathcal{N}(0, 1) \parallel \mathcal{N}\left(\frac{G}{\sqrt{\mu}}, 1\right)\right)(\epsilon).$$

We next give a simple comparison result between norms.

Lemma 4. *For $f : \mathcal{X} \rightarrow \mathbb{R}$, fix $a, b \in \mathcal{X}$, and let $\tilde{f} : [a, b] \rightarrow \mathbb{R}$ be the restriction of f to $[a, b]$.*

- (1) *If f is G -Lipschitz in $\|\cdot\|_{\mathcal{X}}$, \tilde{f} is $G \cdot \frac{\|b-a\|_{\mathcal{X}}}{\|b-a\|_2}$ -Lipschitz in $\|\cdot\|_2$.*
- (2) *If f is μ -strongly convex in $\|\cdot\|_{\mathcal{X}}$, \tilde{f} is $\mu \cdot \frac{\|b-a\|_{\mathcal{X}}^2}{\|b-a\|_2^2}$ -strongly convex in $\|\cdot\|_2$.*

Proof. To see the first claim, let $c = a + r(b - a)$ and $d = a + s(b - a)$ for $s, r \in [0, 1]$. We have by Lipschitzness of f that

$$|\tilde{f}(d) - \tilde{f}(c)| \leq G |s - r| \|b - a\|_{\mathcal{X}} = \left(G \cdot \frac{\|b - a\|_{\mathcal{X}}}{\|b - a\|_2} \right) \cdot \|d - c\|_2.$$

Similarly, to see the second claim, by strong convexity of f ,

$$\begin{aligned}
\tilde{f}(tc + (1-t)d) &\leq t\tilde{f}(c) + (1-t)\tilde{f}(d) - \frac{\mu t(1-t)}{2} \|c - d\|_{\mathcal{X}}^2 \\
&= t\tilde{f}(c) + (1-t)\tilde{f}(d) - \frac{\mu t(1-t)(r-s)^2}{2} \|a - b\|_{\mathcal{X}}^2 \\
&= t\tilde{f}(c) + (1-t)\tilde{f}(d) - \left(\mu \cdot \frac{\|b-a\|_{\mathcal{X}}^2}{\|b-a\|_2^2} \right) \left(\frac{t(1-t)}{2} \|d-c\|_2^2 \right).
\end{aligned}$$

□

We now present our main result on Gaussian differential privacy with respect to arbitrary norms.

Theorem 2. *Let $\mathcal{X} \subset \mathbb{R}^d$ be compact and convex, let $F, \tilde{F} : \mathcal{X} \rightarrow \mathbb{R}$ be μ -strongly convex in $\|\cdot\|_{\mathcal{X}}$, and let $P \propto \exp(-F)$ and $Q \propto \exp(-\tilde{F})$. Suppose $\tilde{F} - F$ is G -Lipschitz in $\|\cdot\|_{\mathcal{X}}$. For all $\epsilon \in \mathbb{R}_{\geq 0}$,*

$$\delta(P \parallel Q)(\epsilon) \leq \delta\left(\mathcal{N}(0, 1) \parallel \mathcal{N}\left(\frac{G}{\sqrt{\mu}}, 1\right)\right)(\epsilon).$$

Proof. Throughout this proof, fix some α which is G -Lipschitz in $\|\cdot\|_{\mathcal{X}}$ by assumption. We first claim that amongst all μ -strongly convex (in $\|\cdot\|_{\mathcal{X}}$) functions $F : \mathcal{X} \rightarrow \mathbb{R}$ such that $F + \alpha$ is also μ -strongly convex, defining $P \propto \exp(-F)$ and $Q \propto \exp(-(F + \alpha))$, some F maximizing $\delta(P \parallel Q)(\epsilon)$ is either a Dirac measure or supported on $[a, b] \subset \mathcal{X}$. We will prove this by contradiction.

Suppose otherwise, and let F be a μ -strongly convex function that maximizes $\delta(P \parallel Q)(\epsilon)$ defined above. Define $P \propto \exp(-F)$ and $Q \propto \exp(-(F + \alpha))$. Let $S^* \subseteq \mathcal{X}$ be the set achieving

$$\delta(P \parallel Q)(\epsilon) = \Pr_{X \sim P}[X \in S^*] - \exp(\epsilon) \Pr_{X \sim Q}[X \in S^*].$$

By Lemma 2, there is another μ -strongly logconcave π where the renormalized density $\propto \pi \exp(-\alpha)$ is also μ -strongly logconcave, such that (following notation of Lemma 2) $g(\pi) \geq g(P)$, where π is either a Dirac or supported on $[a, b]$. We conclude that $\delta(P \parallel Q)(\epsilon) \leq \delta(\pi \parallel \pi \exp(-\alpha))(\epsilon)$ (since the maximizing set for π is at least as good as S^*), a contradiction.

It hence suffices to prove the theorem statement for F, \tilde{F} , which are supported on some $[a, b] \in \mathcal{X}$. By Lemma 4, we have that $\tilde{F} - F$ is $G \cdot \frac{\|b-a\|_{\mathcal{X}}}{\|b-a\|_2}$ -Lipschitz in $\|\cdot\|_2$ and F, \tilde{F} are $\mu \cdot \frac{\|b-a\|_{\mathcal{X}}^2}{\|b-a\|_2^2}$ -strongly convex in $\|\cdot\|_2$. We conclude by Proposition 3 which shows

$$\begin{aligned}
\delta(P \parallel Q)(\epsilon) &\leq \delta\left(\mathcal{N}(0, 1) \parallel \mathcal{N}\left(\frac{G}{\sqrt{\mu}} \cdot \frac{\|b-a\|_{\mathcal{X}}}{\|b-a\|_2} \cdot \frac{\|b-a\|_2}{\|b-a\|_{\mathcal{X}}}, 1\right)\right)(\epsilon) \\
&= \delta\left(\mathcal{N}(0, 1) \parallel \mathcal{N}\left(\frac{G}{\sqrt{\mu}}, 1\right)\right)(\epsilon).
\end{aligned}$$

□

Our proof strategy is a reduction to an application of Proposition 3 in one dimension. It is an interesting open question to obtain a simpler direct proof of Proposition 3 in the one-dimensional setting (without using the machinery of [GLL22]), which is tight up to constant factors.

4 Private ERM and SCO in general norms

In this section, we derive our results for private ERM and SCO in general norms. We will state our results for private ERM (Section 4.1) and SCO (Section 4.2) with respect to an arbitrary compact convex subset \mathcal{X} of a d -dimensional normed space, satisfying Assumption 1. We then use this to derive guarantees for a variety of settings of import in Section 4.3.

4.1 Private ERM under Assumption 1

To develop our private ERM algorithms, we recall the following risk guarantee from [dKL18] of sampling from Gibbs distributions (improving upon [KV06, BST14]).

Proposition 4 ([dKL18], Corollary 1). *Let $\mathcal{X} \subset \mathbb{R}^d$ be compact and convex, let $F : \mathcal{X} \rightarrow \mathbb{R}$ be convex, and let $k > 0$. If $\nu \propto \exp(-kF)$,*

$$\mathbb{E}_{x \sim \nu}[F(x)] \leq \min_{x \in \mathcal{X}} F(x) + \frac{d}{k}.$$

We conclude by a simple combination of Proposition 4 (providing a risk guarantee) and Theorem 2 (providing a privacy guarantee), which yields our main result on private ERM.

Theorem 3 (Private ERM). *Under Assumption 1 and following notation (1), drawing a sample x from the density $\nu \propto \exp(-k(F_{\mathcal{D}} + \mu r))$ for*

$$k = \frac{\sqrt{dn\epsilon}}{G\sqrt{2\Theta \log \frac{1}{2\delta}}}, \quad \mu = \frac{G\sqrt{2d \log \frac{1}{2\delta}}}{\sqrt{\Theta n\epsilon}},$$

is (ϵ, δ) -differentially private, and produces x such that

$$\mathbb{E}_{x \sim \nu}[F_{\mathcal{D}}(x)] - \min_{x \in \mathcal{X}} F_{\mathcal{D}}(x) \leq G\sqrt{\Theta} \cdot \frac{\sqrt{8d \log \frac{1}{2\delta}}}{n\epsilon}.$$

Proof. Let $F_{\mathcal{D}'}$ be the realization of (1) when \mathcal{D} is replaced with a neighboring dataset \mathcal{D}' which agrees in all entries except some sample $s'_i \neq s_i$. By Assumption 1, we have $k(F_{\mathcal{D}} - F_{\mathcal{D}'})$ is $\frac{kG}{n}$ -Lipschitz, and both $k(F_{\mathcal{D}} + \mu r)$ and $k(F_{\mathcal{D}'} + \mu r)$ are $k\mu$ -strongly convex (all with respect to $\|\cdot\|_{\mathcal{X}}$). Hence, combining Theorem 2 and Fact 2 shows the mechanism is (ϵ, δ) -differentially private, since

$$\mu = \frac{2G^2 k \log \frac{1}{2\delta}}{n^2 \epsilon^2} \implies \frac{G\sqrt{k}}{n\sqrt{\mu}} \leq \frac{\epsilon}{\sqrt{2 \log \frac{1}{2\delta}}}. \quad (6)$$

Let $x_{\mathcal{D}}^* := \operatorname{argmin}_{x \in \mathcal{X}} F_{\mathcal{D}}(x)$. We obtain the risk guarantee by the calculation (see Proposition 4)

$$\begin{aligned} \mathbb{E}_{x \sim \nu}[F_{\mathcal{D}}(x)] &\leq F_{\mathcal{D}}(x_{\mathcal{D}}^*) + (\mu r(x_{\mathcal{D}}^*) - \mathbb{E}_{x \sim \nu} \mu r(x)) + \frac{d}{k} \\ &\leq F_{\mathcal{D}}(x_{\mathcal{D}}^*) + \mu\Theta + \frac{d}{k} \end{aligned}$$

and plugging in our choices of μ and k . \square

4.2 Private SCO under Assumption 1

We first give a generic comparison result between population risk and empirical risk under Assumption 1. To do so, we use two helper results from prior work. The first was derived in [GLL22] by combining a transportation inequality and a log-Sobolev inequality (see e.g. [OV00]).

Proposition 5 ([GLL22], Theorem 6.7, Lemma 6.8). *Let $\mathcal{X} \subseteq \mathbb{R}^d$ be compact and convex, let $F, \tilde{F} : \mathcal{X} \rightarrow \mathbb{R}$ be μ -strongly convex in $\|\cdot\|_2$, and let $P \propto \exp(-F)$ and $Q \propto \exp(-\tilde{F})$. Suppose $\tilde{F} - F$ is H -Lipschitz in $\|\cdot\|_2$. Then, $W_2(P, Q) \leq \frac{H}{\mu}$.*

Corollary 1. Let $\mathcal{X} \subset \mathbb{R}^d$ be compact and convex, and let $\alpha, f : \mathcal{X} \rightarrow \mathbb{R}$ be H -Lipschitz and G -Lipschitz respectively in $\|\cdot\|_{\mathcal{X}}$. Let $\mathcal{S}_{\mu, \exp(-\alpha)}$ be the set of densities π over \mathcal{X} such that π is μ -strongly logconcave with respect to $\|\cdot\|_{\mathcal{X}}$, and $\tilde{\pi} \propto \pi \exp(-\alpha)$ is also μ -strongly logconcave. For any $\pi \in \mathcal{S}_{\mu, \exp(-\alpha)}$ define $g(\pi) := \int_{\mathcal{X}} f(x) d(\pi - \tilde{\pi})(x)$ where $\tilde{\pi} \propto \pi \exp(-\alpha)$. Then, $g(\pi) \leq \frac{GH}{\mu}$.

Proof. By Lemma 3 (and following its notation), it suffices to show $g(\pi) \leq \frac{GH}{\mu}$ for all $\pi \in \mathcal{S}_{\mu, \exp(-\alpha)}^*$. Clearly this is true for a Dirac measure π as then $g(\pi) = 0$, so consider the other case where π is supported on $[a, b]$, such that $\pi \propto \exp(-F)$ and F is μ -strongly convex in $\|\cdot\|_{\mathcal{X}}$. Further, define $\tilde{F} = F + \alpha$, so that \tilde{F} is also μ -strongly convex and supported on $[a, b]$.

By Lemma 4, restricting to $[a, b]$, F and \tilde{F} are $\mu \cdot \frac{\|b-a\|_{\mathcal{X}}^2}{\|b-a\|_2^2}$ -strongly convex in $\|\cdot\|_2$, $F - \tilde{F}$ is $H \cdot \frac{\|b-a\|_{\mathcal{X}}}{\|b-a\|_2}$ -Lipschitz in ℓ_2 and f is $\frac{\|b-a\|_{\mathcal{X}}}{\|b-a\|_2}$ -Lipschitz in $\|\cdot\|_2$. Hence, where the inequalities are by Fact 1 and Proposition 5 respectively,

$$g(\pi) = \int_{\mathcal{X}} f(x) d(\pi - \tilde{\pi})(x) \leq GW_2(\pi, \tilde{\pi}) \leq \frac{GH}{\mu}.$$

□

The second relates the population risk to the empirical risk on an independent sample.

Proposition 6 (Lemma 7, [BE02]). Suppose $\mathcal{D} = \{s_i\}_{i \in [n]}$ is drawn independently from \mathcal{P} , let $s \sim \mathcal{P}$ be drawn independently from \mathcal{D} , and let $\mathcal{D}' := \{s\} \cup \{s_i\}_{i \in [n] \setminus \{1\}}$ be \mathcal{D} where s_1 is swapped with s . Then, for any symmetric³ mechanism $\mathcal{M} : \text{supp}(\mathcal{P})^n \rightarrow \mathbb{R}^d$,

$$\mathbb{E}[F_{\text{pop}}(\mathcal{M}(\mathcal{D})) - F_{\mathcal{D}}(\mathcal{M}(\mathcal{D}))] = \mathbb{E}[f(\mathcal{M}(\mathcal{D}); s) - f(\mathcal{M}(\mathcal{D}'); s)],$$

where expectations are over \mathcal{M} and the randomness used in producing \mathcal{D} and s .

By applying Corollary 1 and Proposition 6 (which bound the generalization error of our mechanism), we provide the following extension of Theorem 3, our main result on private SCO.

Theorem 4 (Private SCO). Under Assumption 1 and following notation (1), drawing a sample x from the density $\nu \propto \exp(-k(F_{\mathcal{D}} + \mu r))$ for

$$k = \sqrt{\frac{d + C_2}{C_1}}, \quad \mu = \frac{2G^2 k \log \frac{1}{2\delta}}{n^2 \epsilon^2}, \quad C_1 := \frac{2G^2 \Theta \log \frac{1}{2\delta}}{n^2 \epsilon^2}, \quad C_2 := \frac{n \epsilon^2}{2 \log \frac{1}{2\delta}},$$

is (ϵ, δ) -differentially private, and produces x such that

$$\mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n, x \sim \nu}[F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} F_{\text{pop}}(x) \leq G\sqrt{\Theta} \cdot \left(\frac{\sqrt{8d \log \frac{1}{2\delta}}}{n \epsilon} + \sqrt{\frac{8}{n}} \right).$$

Proof. For the given choice of k, μ , the privacy proof follows identically to Theorem 3, so we focus on the risk proof. We follow the notation of Proposition 6 and let $s \sim \mathcal{P}$ independently from \mathcal{D} .

³Here, a symmetric mechanism is one which only depends on the set of inputs rather than their order.

By exchanging the expectation and minimum and using that $\mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n} F_{\mathcal{D}} = F_{\text{pop}}$,

$$\begin{aligned}
\mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n, x \sim \nu} [F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} F_{\text{pop}}(x) &\leq \mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n} \left[\mathbb{E}_{x \sim \nu} [F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} [F_{\mathcal{D}}(x)] \right] \\
&\leq \mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n} [\mathbb{E}_{x \sim \nu} [F_{\text{pop}}(x) - F_{\mathcal{D}}(x)]] \\
&\quad + \mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n} \left[\mathbb{E}_{x \sim \nu} [F_{\mathcal{D}}(x)] - \min_{x \in \mathcal{X}} [F_{\mathcal{D}}(x)] \right] \\
&\leq \mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n} [\mathbb{E}_{x \sim \nu} [F_{\text{pop}}(x) - F_{\mathcal{D}}(x)]] + \mu\Theta + \frac{d}{k},
\end{aligned}$$

where we bounded the empirical risk in the proof of Theorem 3. Next, let ν' be the density $\propto \exp(-k(F_{\mathcal{D}'} + \mu r))$. Our mechanism is symmetric, and hence by Proposition 6,

$$\mathbb{E} [F_{\text{pop}}(x) - F_{\mathcal{D}}(x)] = \mathbb{E} [\mathbb{E}_{x \sim \nu} [f(x; s)] - \mathbb{E}_{x \sim \nu'} [f(x; s)]]$$

where the outer expectations are over the randomness of drawing \mathcal{D}, s . Finally, for any fixed realization of \mathcal{D}, s , the densities ν, ν' satisfy the assumption of Corollary 1 with $H = \frac{G}{n}$, and $f(\cdot; s)$ is G -Lipschitz, so Corollary 1 shows that

$$\mathbb{E}_{x \sim \nu} [f(x; s)] - \mathbb{E}_{x \sim \nu'} [f(x; s)] \leq \frac{G^2}{n\mu}.$$

Combining the above three displays bounds the population risk by

$$\begin{aligned}
\mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n, x \sim \nu} [F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} F_{\text{pop}}(x) &\leq \frac{G^2}{n\mu} + \mu\Theta + \frac{d}{k} \\
&= C_1 k + \frac{C_2 + d}{k},
\end{aligned}$$

for our given value of μ . The conclusion follows by optimizing over k yielding a risk of $2\sqrt{C_1(C_2 + d)}$, and using the scalar inequality $\sqrt{a + b} \leq \sqrt{a} + \sqrt{b}$ for nonnegative a, b . \square

4.3 Applications

To derive our private optimization algorithms for ℓ_p -norm and Schatten- p norm geometries, we recall the following results on the existence of bounded strongly convex regularizers.

Proposition 7 ([BCL94]). *For $1 < p \leq 2$, letting $\|\cdot\|_p$ be the ℓ_p norm of a vector, $r(v) := \frac{1}{2(p-1)}\|v\|_p^2$ is 1-strongly convex in $\|\cdot\|_p$. Similarly, for $1 < p \leq 2$, letting $\|\cdot\|_p$ be the Schatten- p norm of a matrix, $r(\mathbf{M}) := \frac{1}{2(p-1)}\|\mathbf{M}\|_p^2$ is 1-strongly convex in $\|\cdot\|_p$.*

We state a useful result on efficiently sampling from Lipschitz, strongly logconcave densities under value oracle access given by [GLL22] (building upon the framework of [LST21]). We slightly specialize the result of [GLL22] by giving a rephrasing sufficient for our purposes.

Proposition 8 ([GLL22], Theorem 2.3). *Let $\mathcal{X} \subset \mathbb{R}^d$ be compact and convex with $\text{diam}_{\|\cdot\|_2}(\mathcal{X}) \leq D$. Let $\mathcal{D} = \{s_i\}_{i \in [n]}$ and let $\tilde{F}_{\mathcal{D}}(x) = \frac{1}{n} \sum_{i \in [n]} f(x; s_i) + \psi(x)$ such that all $f(\cdot; s_i) : \mathcal{X} \rightarrow \mathbb{R}$ are G -Lipschitz in $\|\cdot\|_2$ and convex, and $\psi(x) : \mathcal{X} \rightarrow \mathbb{R}$ is μ -strongly convex in $\|\cdot\|_2$. For $\delta \in (0, \frac{1}{2})$, we can generate a sample within total variation δ of the density $\propto \exp(-\tilde{F}_{\mathcal{D}})$ in N value queries to some $f(\cdot; s_i)$ and samples from densities $\propto \exp\left(-\psi - \frac{1}{2\eta}\|\cdot - v\|_2^2\right)$ for some $\eta > 0$, $v \in \mathbb{R}^d$, where*

$$N = O\left(\frac{G^2}{\mu} \log^2\left(\frac{G^2(D^2 + \mu^{-1})d}{\delta}\right)\right).$$

ℓ_p norms. We state our results on private convex optimization under ℓ_p geometry. As a preliminary, we combine norm equivalence bounds (2) and Proposition 8 to give the following algorithmic result on sampling from a logconcave distribution under value oracle access under ℓ_p geometry.

Proposition 9. *Let $p \geq 1$ and let $\mathcal{X} \subset \mathbb{R}^d$ be compact and convex with $\text{diam}_{\|\cdot\|_p}(\mathcal{X}) \leq D$. Let $\mathcal{D} = \{s_i\}_{i \in [n]}$ and let $\tilde{F}_{\mathcal{D}}(x) = \frac{1}{n} \sum_{i \in [n]} f(x; s_i) + \psi(x)$ such that all $f(\cdot; s_i) : \mathcal{X} \rightarrow \mathbb{R}$ are G -Lipschitz in $\|\cdot\|_p$ and convex, and $\psi(x) : \mathcal{X} \rightarrow \mathbb{R}$ is μ -strongly convex in $\|\cdot\|_p$. For $\delta \in (0, \frac{1}{2})$, we can generate a sample within total variation δ of the density $\propto \exp(-\tilde{F}_{\mathcal{D}})$ in N value queries to some $f(\cdot; s_i)$ and samples from densities $\propto \exp\left(-\psi - \frac{1}{2\eta}\|\cdot - v\|_2^2\right)$ for some $\eta > 0$, $v \in \mathbb{R}^d$, where*

$$N = O\left(\frac{G^2 d^{\frac{2}{p}-1}}{\mu} \log^2\left(\frac{G^2(D^2 + \mu^{-1})d}{\delta}\right)\right) \text{ if } p \in [1, 2],$$

$$N = O\left(\frac{G^2 d^{1-\frac{2}{p}}}{\mu} \log^2\left(\frac{G^2(D^2 + \mu^{-1})d}{\delta}\right)\right) \text{ if } p \in [2, \infty).$$

Proof. For $p \in [1, 2]$, note that each $f(\cdot; s_i)$ is $d^{\frac{1}{p}-\frac{1}{2}}G$ -Lipschitz in the ℓ_2 norm by combining (2) and the definition of Lipschitzness. Moreover, because the ℓ_p norm is larger than the ℓ_2 norm, ψ remains μ -strongly convex in the ℓ_2 norm. The diameter D is only affected by $\text{poly}(d)$ factors when converting norms, which is accounted for by the logarithmic term. Hence, the complexity bound follows by applying Proposition 8 under this change of parameters. For the other case of $p \in [2, \infty)$, the Lipschitz bound is G , and the strong convexity bound is $d^{\frac{2}{p}-1}\mu$ by a similar argument. \square

In the following discussion, we primarily focus on the value oracle query complexity of our samplers. Generic results on logconcave sampling (see e.g. [LV07], or more recent developments by [JLLV21, Che21, KL22]) imply the samples from the densities $\propto \exp(-\psi - \frac{1}{2\eta}\|\cdot - v\|_2^2)$ can be performed in polynomial time, for all the ψ that are relevant in our applications (which are all squared ℓ_p distances). We expect samplers which run in nearly-linear time (in d) may be designed for applications where \mathcal{X} is structured, such as an ℓ_p ball, but for brevity we omit this discussion.

Corollary 2. *Let $1 < p \leq 2$ be a constant, and let $\epsilon > 0$, $\delta \in (0, 1)$. Let $\mathcal{X} \subset \mathbb{R}^d$ have $\text{diam}_{\|\cdot\|_p}(\mathcal{X}) \leq D$, and let $F_{\text{pop}} = \mathbb{E}_{s \sim \mathcal{P}}[f(\cdot; s)]$ where all $f(\cdot; s) : \mathbb{R}^d \rightarrow \mathbb{R}$ are convex and G -Lipschitz in $\|\cdot\|_p$. Finally, let $\mathcal{D} = \{s_i\}_{i \in [n]} \sim \mathcal{P}^n$ independently and $F_{\mathcal{D}} := \frac{1}{n} \sum_{i \in [n]} f(\cdot; s_i)$.*

(1) *There is an (ϵ, δ) -differentially private algorithm \mathcal{M} which produces x such that*

$$\mathbb{E}_{\mathcal{M}}[F_{\mathcal{D}}(x)] - \min_{x \in \mathcal{X}} F_{\mathcal{D}}(x) \leq 2GD \cdot \frac{\sqrt{d \log \frac{1}{2\delta}}}{n\epsilon\sqrt{p-1}} \quad (7)$$

using

$$O\left(\frac{n^2\epsilon^2 d^{\frac{2}{p}-1}}{\log \frac{1}{\delta}} \log^2\left(\frac{GDdn\epsilon}{\delta}\right)\right) \text{ value queries to some } f(\cdot; s_i). \quad (8)$$

(2) *There is an (ϵ, δ) -differentially private algorithm \mathcal{M} which produces x such that*

$$\mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n, \mathcal{M}}[F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} F_{\text{pop}}(x) \leq 2GD \cdot \left(\sqrt{\frac{1}{n(p-1)}} + \frac{\sqrt{d \log \frac{1}{2\delta}}}{n\epsilon\sqrt{p-1}}\right). \quad (9)$$

using

$$O\left(\frac{n^2\epsilon^2d^{\frac{2}{p}-1}}{\log\frac{1}{\delta}}\log^2\left(\frac{GDdne}{\delta}\right)\right) \text{ value queries to some } f(\cdot; s_i).$$

Proof. We will parameterize Assumption 1 with the function $r(x) := \frac{1}{2(p-1)}\|x - x_0\|_p^2$, where $x_0 \in \mathcal{X}$ is an arbitrary point, and strong convexity follows from Proposition 7. By assumption, we may set $\Theta = \frac{1}{2(p-1)}D^2$. The conclusions follow by combining Theorem 3, Theorem 4. To obtain (ϵ, δ) -differential privacy, it suffices to run the mechanism with privacy level $\delta \leftarrow \frac{\delta}{2}$, run to total variation $\frac{\delta}{2}$ using Proposition 9, and take a union bound. For both ERM and SCO, note that our choices of k and μ satisfy the relation (6), namely $\frac{kG^2}{\mu} = O(n^2\epsilon^2/\log\frac{1}{\delta})$. Since both the Lipschitz and strong convexity parameters are scaled up by k in our application of Proposition 9, we have the leading-order term is $\frac{kG^2}{\mu}$ which yields the conclusion. \square

For any p such that $p-1$ is bounded away from 0, Corollary 2 matches the information-theoretic lower bound of [BGN21] (and its subsequent sharpening by [LL22]). When this is not the case, we use norm equivalence (2) to obtain a weaker bound.

Corollary 3. *Let $\epsilon > 0$, $\delta \in (0, 1)$. Let $\mathcal{X} \subset \mathbb{R}^d$ have $\text{diam}_{\|\cdot\|_1}(\mathcal{X}) \leq D$, and let $F_{\text{pop}} = \mathbb{E}_{s \sim \mathcal{P}}[f(\cdot; s)]$ where all $f(\cdot; s) : \mathbb{R}^d \rightarrow \mathbb{R}$ are convex and G -Lipschitz in $\|\cdot\|_1$. Finally, let $\mathcal{D} = \{s_i\}_{i \in [n]} \sim \mathcal{P}^n$ independently and $F_{\mathcal{D}} := \frac{1}{n} \sum_{i \in [n]} f(\cdot; s_i)$.*

(1) *There is an (ϵ, δ) -differentially private algorithm \mathcal{M} which produces x such that*

$$\mathbb{E}_{\mathcal{M}}[F_{\mathcal{D}}(x)] - \min_{x \in \mathcal{X}} F_{\mathcal{D}}(x) \leq 6GD\sqrt{\log d} \cdot \frac{\sqrt{d \log \frac{1}{2\delta}}}{n\epsilon} \quad (10)$$

using

$$O\left(\frac{n^2\epsilon^2d}{\log\frac{1}{\delta}}\log^2\left(\frac{GDdne}{\delta}\right)\right) \text{ value queries to some } f(\cdot; s_i). \quad (11)$$

(2) *There is an (ϵ, δ) -differentially private algorithm \mathcal{M} which produces x such that*

$$\mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n, \mathcal{M}}[F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} F_{\text{pop}}(x) \leq 6GD\sqrt{\log d} \cdot \left(\sqrt{\frac{1}{n}} + \frac{\sqrt{d \log \frac{1}{2\delta}}}{n\epsilon}\right) \quad (12)$$

using

$$O\left(\frac{n^2\epsilon^2d}{\log\frac{1}{\delta}}\log^2\left(\frac{GDdne}{\delta}\right)\right) \text{ value queries to some } f(\cdot; s_i).$$

Proof. We will parameterize Assumption 1 with the function $r(x) := \frac{e^2}{2(q-1)}\|x - x_0\|_q^2$, where $q = 1 + \frac{1}{\log d}$. By combining Proposition 7 (which shows r is e^2 -strongly convex in ℓ_q) and (2), we have that r is 1-strongly convex in ℓ_1 . The remainder of the proof follows identically to Corollary 2. \square

The term scaling as $\sqrt{\log d/n}$ in (12), namely the non-private population risk, is known to be optimal from existing lower bounds on SCO [DJW14]. Up to a $\sqrt{\log d}$ factor, the non-private empirical risk is optimal with respect to current private optimization lower bounds [BGN21, LL22].

Corollary 4. Let $p \geq 2$, and let $\epsilon > 0$, $\delta \in (0, 1)$. Let $\mathcal{X} \subset \mathbb{R}^d$ have $\text{diam}_{\|\cdot\|_p}(\mathcal{X}) \leq D$, and let $F_{\text{pop}} = \mathbb{E}_{s \sim \mathcal{P}}[f(\cdot; s)]$ where all $f(\cdot; s) : \mathbb{R}^d \rightarrow \mathbb{R}$ are convex and G -Lipschitz in $\|\cdot\|_p$. Finally, let $\mathcal{D} = \{s_i\}_{i \in [n]} \sim \mathcal{P}^n$ independently and $F_{\mathcal{D}} := \frac{1}{n} \sum_{i \in [n]} f(\cdot; s_i)$.

(1) There is an (ϵ, δ) -differentially private algorithm \mathcal{M} which produces x such that

$$\mathbb{E}_{\mathcal{M}}[F_{\mathcal{D}}(x)] - \min_{x \in \mathcal{X}} F_{\mathcal{D}}(x) \leq 2GD \cdot \frac{d^{1-\frac{1}{p}} \sqrt{\log \frac{1}{2\delta}}}{n\epsilon} \quad (13)$$

using

$$O\left(\frac{n^2 \epsilon^2 d^{1-\frac{2}{p}}}{\log \frac{1}{\delta}} \log^2\left(\frac{GDdne}{\delta}\right)\right) \text{ value queries to some } f(\cdot; s_i). \quad (14)$$

(2) There is an (ϵ, δ) -differentially private algorithm \mathcal{M} which produces x such that

$$\mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n, \mathcal{M}}[F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} F_{\text{pop}}(x) \leq 2GD \cdot \left(\frac{d^{\frac{1}{2}-\frac{1}{p}}}{\sqrt{n}} + \frac{d^{1-\frac{1}{p}} \sqrt{\log \frac{1}{2\delta}}}{n\epsilon}\right). \quad (15)$$

using

$$O\left(\frac{n^2 \epsilon^2 d^{1-\frac{2}{p}}}{\log \frac{1}{\delta}} \log^2\left(\frac{GDdne}{\delta}\right)\right) \text{ value queries to some } f(\cdot; s_i).$$

Proof. We will parameterize Assumption 1 with the function $r(x) := \frac{1}{2} \|x - x_0\|_2^2$. By combining Proposition 7 (which shows r is 1-strongly convex in ℓ_2 , and hence also ℓ_p) and (2), we may set $\Theta = \frac{1}{2} d^{1-2/p} D^2$. The remainder of the proof follows identically to Corollary 2. \square

Schatten- p norms. Our results extend immediately to matrix spaces equipped with Schatten- p norm geometries. We record our relevant results in the following.

Corollary 5. Let $p \in [1, \infty)$, $\epsilon > 0$, $\delta \in (0, 1)$, and let $d_1, d_2 \in \mathbb{N}$ have $d_1 > d_2$. Let $\mathcal{X} \subset \mathbb{R}^{d_1 \times d_2}$ have $\text{diam}_{\|\cdot\|_p}(\mathcal{X}) \leq D$, and let $F_{\text{pop}} = \mathbb{E}_{s \sim \mathcal{P}}[f(\cdot; s)]$ where all $f(\cdot; s) : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}$ are convex and G -Lipschitz in $\|\cdot\|_p$. Finally, let $\mathcal{D} = \{s_i\}_{i \in [n]} \sim \mathcal{P}^n$ independently and $F_{\mathcal{D}} := \frac{1}{n} \sum_{i \in [n]} f(\cdot; s_i)$.

(1) For constant $1 < p \leq 2$, there is an (ϵ, δ) -differentially private algorithm \mathcal{M} which produces \mathbf{M} such that

$$\begin{aligned} \mathbb{E}_{\mathcal{M}}[F_{\mathcal{D}}(\mathbf{M})] - \min_{\mathbf{M} \in \mathcal{X}} F_{\mathcal{D}}(\mathbf{M}) &\leq 2GD \cdot \frac{\sqrt{d_1 d_2 \log \frac{1}{2\delta}}}{n\epsilon\sqrt{p-1}}, \\ \mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n, \mathcal{M}}[F_{\text{pop}}(\mathbf{M})] - \min_{\mathbf{M} \in \mathcal{X}} F_{\text{pop}}(\mathbf{M}) &\leq 2GD \cdot \left(\sqrt{\frac{1}{n(p-1)}} + \frac{\sqrt{d_1 d_2 \log \frac{1}{2\delta}}}{n\epsilon\sqrt{p-1}}\right). \end{aligned}$$

The value oracle complexity of the algorithm is bounded as in (8) for $d \leftarrow d_2$ in the non-logarithmic term, and $d \leftarrow d_1$ in the logarithmic term.

(2) For $p = 1$, there is an (ϵ, δ) -differentially private algorithm \mathcal{M} which produces \mathbf{M} such that

$$\begin{aligned} \mathbb{E}_{\mathcal{M}}[F_{\mathcal{D}}(\mathbf{M})] - \min_{\mathbf{M} \in \mathcal{X}} F_{\mathcal{D}}(\mathbf{M}) &\leq 6GD\sqrt{\log d_2} \cdot \frac{\sqrt{d_1 d_2 \log \frac{1}{2\delta}}}{n\epsilon\sqrt{p-1}}, \\ \mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n, \mathcal{M}}[F_{\text{pop}}(\mathbf{M})] - \min_{\mathbf{M} \in \mathcal{X}} F_{\text{pop}}(\mathbf{M}) &\leq 6GD\sqrt{\log d_2} \cdot \left(\sqrt{\frac{1}{n}} + \frac{\sqrt{d_1 d_2 \log \frac{1}{2\delta}}}{n\epsilon\sqrt{p-1}}\right). \end{aligned}$$

The value oracle complexity of the algorithm is bounded as in (11) for $d \leftarrow d_2$ in the non-logarithmic term, and $d \leftarrow d_1$ in the logarithmic term.

(3) For $p \geq 2$, there is an (ϵ, δ) -differentially private algorithm \mathcal{M} which produces \mathbf{M} such that

$$\begin{aligned} \mathbb{E}_{\mathcal{M}}[F_{\mathcal{D}}(\mathbf{M})] - \min_{\mathbf{M} \in \mathcal{X}} F_{\mathcal{D}}(\mathbf{M}) &\leq 2GD \cdot \frac{d_2^{\frac{1}{2} - \frac{1}{p}} \sqrt{d_1 d_2 \log \frac{1}{2\delta}}}{n\epsilon}, \\ \mathbb{E}_{\mathcal{D} \sim \mathcal{P}^n, \mathcal{M}}[F_{\text{pop}}(\mathbf{M})] - \min_{\mathbf{M} \in \mathcal{X}} F_{\text{pop}}(\mathbf{M}) &\leq 2GD \cdot \left(\frac{d_2^{\frac{1}{2} - \frac{1}{p}}}{\sqrt{n}} + \frac{d_2^{\frac{1}{2} - \frac{1}{p}} \sqrt{d_1 d_2 \log \frac{1}{2\delta}}}{n\epsilon} \right). \end{aligned}$$

The value oracle complexity of the algorithm is bounded as in (14) for $d \leftarrow d_2$ in the non-logarithmic term, and $d \leftarrow d_1$ in the logarithmic term.

Proof. The privacy and utility proofs follow identically to Corollaries 2, 3, and 4, where we use the second portion of Proposition 7 instead of the first. We note that the “dimension-dependent” term in the risk inherited from Proposition 4 scales as $d_1 d_2$ (the dimensionality of the matrix space). However, the terms in the risk due to the size of regularizers (inherited from the tradeoffs in (2), for $p = 1$ and $p > 2$) scales as a power of d_2 , the maximum dimension of singular values. To obtain the value oracle complexity, we note that by definition of the Schatten norm, it satisfies the relationship (2) as well. Moreover, the Schatten-2 norm agrees with the vector ℓ_2 norm (when the matrix is flattened into a vector), since they are both the Frobenius norm. Hence, we may directly apply Proposition 8 after paying a norm conversion, in the same way as was done in Proposition 9. \square

Remark on high-probability bounds. One advantage of using a sampling-based algorithm is an immediate high-probability bound which follows due to the good concentration of Lipschitz functions over samples from strongly logconcave distributions, stated below.

Lemma 5 (Concentration of Lipschitz functions, [Led99], Section 2.3 and [BL00], Proposition 3.1). *Let ℓ be a G -Lipschitz function and $X \sim \exp(-F)$ for a μ -strongly convex function F , all with respect to the same norm $\|\cdot\|_{\mathcal{X}}$. For all $t \geq 0$,*

$$\Pr[\ell(X) - \mathbb{E}[\ell(X)] \geq t] \leq \exp\left(-\frac{t^2\mu}{2G^2}\right).$$

In particular, we have demonstrated that the population and empirical risks (which are Lipschitz) have good expectations. Naively combining Lemma 5 and our main results on the expectation utility bound then yields tight concentration around the mean in some parameter regimes, but we suspect the resulting bound is loose in general. We leave it as an interesting open problem to obtain tight high-probability bounds in all parameter regimes.

Acknowledgments

We would like to thank the anonymous SODA reviewers for their helpful suggestions. We would also like to thank a reviewer and Cristóbal Guzmán for pointing out a mistake in the utility guarantees in Corollary 5 in the original version, which is corrected in this version.

References

- [Abo16] John M. Abowd. The challenge of scientific reproducibility and privacy protection for statistical agencies. Technical report, Census Scientific Advisory Committee, 2016. [1](#)
- [ABRW12] Alekh Agarwal, Peter L. Bartlett, Pradeep Ravikumar, and Martin J. Wainwright. Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization. *IEEE Trans. Inf. Theory*, 58(5):3235–3249, 2012. [1.1](#), [1](#), [1.1](#)
- [AC21] Kwangjun Ahn and Sinho Chewi. Efficient constrained sampling via the mirror-langevin algorithm. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 28405–28418, 2021. [1.3](#)
- [ACG⁺16] Martín Abadi, Andy Chu, Ian J. Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 308–318. ACM, 2016. [1](#)
- [AFKT21] Hilal Asi, Vitaly Feldman, Tomer Koren, and Kunal Talwar. Private stochastic convex optimization: Optimal rates in L1 geometry. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 393–403. PMLR, 2021. [\(document\)](#), [1](#), [1.1](#), [1.3](#), [1.3](#)
- [AGL⁺18] Zeyuan Allen-Zhu, Ankit Garg, Yuanzhi Li, Rafael Mendes de Oliveira, and Avi Wigderson. Operator scaling via geodesically convex optimization, invariant theory and polynomial identity testing. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 172–181. ACM, 2018. [1](#)
- [AHK12] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory Comput.*, 8(1):121–164, 2012. [1](#), [1.1](#)
- [AKPS19] Deeksha Adil, Rasmus Kyng, Richard Peng, and Sushant Sachdeva. Iterative refinement for p-norm regression. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1405–1424. SIAM, 2019. [1](#)
- [ANW10] Alekh Agarwal, Sahand N. Negahban, and Martin J. Wainwright. Fast global convergence rates of gradient methods for high-dimensional statistical recovery. In John D. Lafferty, Christopher K. I. Williams, John Shawe-Taylor, Richard S. Zemel, and Aron Culotta, editors, *Advances in Neural Information Processing Systems 23: 24th Annual Conference on Neural Information Processing Systems 2010. Proceedings of a meeting held 6-9 December 2010, Vancouver, British Columbia, Canada*, pages 37–45. Curran Associates, Inc., 2010. [1](#)
- [BC12] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and non-stochastic multi-armed bandit problems. *Found. Trends Mach. Learn.*, 5(1):1–122, 2012. [1](#)

[BCL94] Keith Ball, Eric A. Carlen, and Elliott H. Lieb. Sharp uniform convexity and smoothness estimates for trace norms. *Inventiones mathematicae*, 115(1):463–482, 1994. [1.1](#), [7](#)

[BDKT12] Aditya Bhaskara, Daniel Dadush, Ravishankar Krishnaswamy, and Kunal Talwar. Unconditional differentially private mechanisms for linear queries. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 1269–1284, 2012. [1.3](#)

[BE02] Olivier Bousquet and André Elisseeff. Stability and generalization. *J. Mach. Learn. Res.*, 2:499–526, 2002. [6](#)

[BEM⁺17] Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnés, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th Symposium on Operating Systems Principles, Shanghai, China, October 28-31, 2017*, pages 441–459. ACM, 2017. [1](#)

[BFGT20] Raef Bassily, Vitaly Feldman, Cristóbal Guzmán, and Kunal Talwar. Stability of stochastic gradient descent on nonsmooth convex losses. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. [1.3](#)

[BFTT19] Raef Bassily, Vitaly Feldman, Kunal Talwar, and Abhradeep Guha Thakurta. Private stochastic convex optimization with optimal rates. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 11279–11288, 2019. [1.3](#)

[BGM21] Raef Bassily, Cristóbal Guzmán, and Michael Menart. Differentially private stochastic optimization: New results in convex and non-convex settings. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 9317–9329, 2021. [1](#), [1.3](#)

[BGN21] Raef Bassily, Cristóbal Guzmán, and Anupama Nandi. Non-euclidean differentially private stochastic convex optimization. In Mikhail Belkin and Samory Kpotufe, editors, *Conference on Learning Theory, COLT 2021, 15-19 August 2021, Boulder, Colorado, USA*, volume 134 of *Proceedings of Machine Learning Research*, pages 474–499. PMLR, 2021. ([document](#)), [1](#), [1](#), [1.1](#), [1.3](#), [4.3](#), [4.3](#)

[BL00] Sergey G Bobkov and Michel Ledoux. From brunn-minkowski to brascamp-lieb and to logarithmic sobolev inequalities. *GAFA, Geometric and Functional Analysis*, 10:1028–1052, 2000. [1.2](#), [5](#)

[Bor75a] Christer Borell. The brunn-minkowski in gauss space. *Inventiones mathematicae*, 30:207–216, 1975. [1.2](#)

[Bor75b] Christer Borell. Convex set functions in space. *Periodica Mathematica Hungarica*, 6(2):111–136, 1975. [1](#), [2](#)

[BST14] Raef Bassily, Adam D. Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 464–473. IEEE Computer Society, 2014. [1](#), [1.1](#), [1.3](#), [4.1](#)

[Bub15] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Found. Trends Mach. Learn.*, 8(3-4):231–357, 2015. [1](#)

[CDWY20] Yuansi Chen, Raaz Dwivedi, Martin J. Wainwright, and Bin Yu. Fast mixing of metropolized hamiltonian monte carlo: Benefits of multi-step gradients. *J. Mach. Learn. Res.*, 21:92:1–92:72, 2020. [1](#)

[CE17] Dario Cordero-Erausquin. Transport inequalities for log-concave measures, quantitative forms, and applications. *Canada J. Math.*, 69(3):481–501, 2017. [1.2](#)

[Che21] Yuansi Chen. An almost constant lower bound of the isoperimetric coefficient in the kls conjecture. *GAFA, Geometric and Functional Analysis*, 31:34–61, 2021. [4.3](#)

[CM08] Kamalika Chaudhuri and Claire Monteleoni. Privacy-preserving logistic regression. In Daphne Koller, Dale Schuurmans, Yoshua Bengio, and Léon Bottou, editors, *Advances in Neural Information Processing Systems 21, Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 8-11, 2008*, pages 289–296. Curran Associates, Inc., 2008. [1](#)

[CMS11] Kamalika Chaudhuri, Claire Monteleoni, and Anand D. Sarwate. Differentially private empirical risk minimization. *J. Mach. Learn. Res.*, 12:1069–1109, 2011. [1](#), [1.3](#)

[CRT06] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory*, 52(2):489–509, 2006. [1](#), [1.1](#)

[DFO20] Jelena Diakonikolas, Maryam Fazel, and Lorenzo Orecchia. Fair packing and covering on a relative scale. *SIAM J. Optim.*, 30(4):3284–3314, 2020. [1](#)

[DG21] Jelena Diakonikolas and Cristóbal Guzmán. Complementary composite minimization, small gradients in general norms, and applications to regression problems. *CoRR*, abs/2101.11041, 2021. [1](#)

[DJW14] John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Privacy aware learning. *J. ACM*, 61(6):38:1–38:57, 2014. [1.1](#), [1](#), [1.1](#), [4.3](#)

[dKL18] Etienne de Klerk and Monique Laurent. Comparison of lasserre’s measure-based bounds for polynomial optimization to bounds obtained by simulated annealing. *Math. Oper. Res.*, 43(4):1317–1325, 2018. [4.1](#), [4](#)

[DKM⁺06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In Serge Vaudeyay, editor, *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, volume 4004 of *Lecture Notes in Computer Science*, pages 486–503. Springer, 2006. [1](#)

[DKY17] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 3571–3580, 2017. [1](#)

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006. [1](#)

[DRS21] Jinshuo Dong, Aaron Roth, and Weijie J. Su. Gaussian differential privacy. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 84(1):3–37, 2021. [1.2](#), [2](#), [3.2](#)

[EPK14] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: randomized aggregatable privacy-preserving ordinal response. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 1054–1067. ACM, 2014. [1](#)

[FG04] Matthieu Fradelizi and Olivier Guédon. The extreme points of subsets of s-concave probabilities and a geometric localization theorem. *Discrete & Computational Geometry*, 31(2):327–335, 2004. [1.2](#), [3.1](#), [2](#)

[FKT20] Vitaly Feldman, Tomer Koren, and Kunal Talwar. Private stochastic convex optimization: optimal rates in linear time. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 439–449. ACM, 2020. [1.3](#)

[GLL22] Sivakanth Gopi, Yin Tat Lee, and Daogao Liu. Private convex optimization via exponential mechanism. In Po-Ling Loh and Maxim Raginsky, editors, *Conference on Learning Theory, 2-5 July 2022, London, UK*, volume 178 of *Proceedings of Machine Learning Research*, pages 1948–1989. PMLR, 2022. ([document](#)), [1.1](#), [1.2](#), [1.2](#), [1.3](#), [1.3](#), [2](#), [3](#), [3.2](#), [3](#), [3.2](#), [4.2](#), [5](#), [4.3](#), [8](#)

[GTU22] Arun Ganesh, Abhradeep Thakurta, and Jalaj Upadhyay. Langevin diffusion: An almost universal algorithm for private euclidean (convex) optimization. *CoRR*, abs/2204.01585, 2022. [1.2](#), [1.3](#)

[HKRC18] Ya-Ping Hsieh, Ali Kavis, Paul Rolland, and Volkan Cevher. Mirrored langevin dynamics. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 2883–2892, 2018. [1.3](#)

[HLL⁺22] Yuxuan Han, Zhicong Liang, Zhipeng Liang, Yang Wang, Yuan Yao, and Jiheng Zhang. Private streaming sco in ℓ_p geometry with applications in high dimensional online decision making. In *International Conference on Machine Learning*, pages 8249–8279. PMLR, 2022. [1.3](#)

[HT10] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 705–714, 2010. [1.3](#)

[Jia21] Qijia Jiang. Mirror langevin monte carlo: the case under isoperimetry. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 715–725, 2021. [1.3](#)

[JLLV21] He Jia, Aditi Laddha, Yin Tat Lee, and Santosh S. Vempala. Reducing isotropy and volume to KLS: an $o^*(n^3\psi^2)$ volume algorithm. In Samir Khuller and Virginia Vasilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 961–974. ACM, 2021. [4.3](#)

[JLT20] Arun Jambulapati, Jerry Li, and Kevin Tian. Robust sub-gaussian principal component analysis and width-independent schatten packing. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. [1](#)

[JT14] Prateek Jain and Abhradeep Guha Thakurta. (near) dimension independent risk bounds for differentially private learning. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, volume 32 of *JMLR Workshop and Conference Proceedings*, pages 476–484. JMLR.org, 2014. [1](#)

[KJ16] Shiva Prasad Kasiviswanathan and Hongxia Jin. Efficient private empirical risk minimization for high-dimensional learning. In Maria-Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 488–497. JMLR.org, 2016. [1](#)

[KL22] Bo'az Klartag and Joseph Lehec. Bourgain's slicing problem and kls isoperimetry up to polylog. *arXiv preprint arXiv:2203.15551*, 2022. [4.3](#)

[KLL21] Janardhan Kulkarni, Yin Tat Lee, and Daogao Liu. Private non-smooth ERM and SCO in subquadratic steps. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 4053–4064, 2021. [1.3](#)

[KLOS14] Jonathan A. Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 217–226. SIAM, 2014. [1](#)

[KLS95] Ravi Kannan, László Lovász, and Miklós Simonovits. Isoperimetric problems for convex bodies and a localization lemma. *Discrete & Computational Geometry*, 13(3):541–559, 1995. [1.2](#), [3](#)

[Kol11] Alexander V. Kolesnikov. Mass transportation and contractions. *arXiv preprint arXiv:1103.1479*, 2011. [1.2](#)

[KPSW19] Rasmus Kyng, Richard Peng, Sushant Sachdeva, and Di Wang. Flows in almost linear time via adaptive preconditioning. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 902–913. ACM, 2019. [1](#)

[KST12] Daniel Kifer, Adam D. Smith, and Abhradeep Thakurta. Private convex optimization for empirical risk minimization with applications to high-dimensional regression. In Shie Mannor, Nathan Srebro, and Robert C. Williamson, editors, *COLT 2012 - The*

25th Annual Conference on Learning Theory, June 25-27, 2012, Edinburgh, Scotland, volume 23 of *JMLR Proceedings*, pages 25.1–25.40. JMLR.org, 2012. [1](#), [1.3](#)

[KV06] Adam Tauman Kalai and Santosh S. Vempala. Simulated annealing for convex optimization. *Math. Oper. Res.*, 31(2):253–266, 2006. [4.1](#)

[Led99] Michel Ledoux. *Concentration of measure and logarithmic Sobolev inequalities*. Séminaire de probabilités XXXIII, 1999. [1.2](#), [5](#)

[LL22] Daogao Liu and Zhou Lu. Lower bounds for differentially private erm: Unconstrained and non-euclidean. *arXiv preprint arXiv:2105.13637*, 2022. [1](#), [1.1](#), [4.3](#), [4.3](#)

[LS93] László Lovász and Miklós Simonovits. Random walks in a convex body and an improved volume algorithm. *Random structures & algorithms*, 4(4):359–412, 1993. [1](#), [3](#)

[LST21] Yin Tat Lee, Ruoqi Shen, and Kevin Tian. Structured logconcave sampling with a restricted gaussian oracle. In *Conference on Learning Theory*, pages 2993–3050. PMLR, 2021. [1.1](#), [1.3](#), [4.3](#)

[LTVW22] Ruilin Li, Molei Tao, Santosh S. Vempala, and Andre Wibisono. The mirror langevin algorithm converges with vanishing bias. In Sanjoy Dasgupta and Nika Haghtalab, editors, *International Conference on Algorithmic Learning Theory, 29-1 April 2022, Paris, France*, volume 167 of *Proceedings of Machine Learning Research*, pages 718–742. PMLR, 2022. [1.3](#)

[LV07] László Lovász and Santosh S. Vempala. The geometry of logconcave functions and sampling algorithms. *Random Struct. Algorithms*, 30(3):307–358, 2007. [4.3](#)

[MS08] Emanuel Milman and Sasha Sodin. An isoperimetric inequality for uniformly log-concave measures and uniformly convex bodies. *Journal of Functional Analysis*, 254(5):1235–1268, 2008. [1.2](#)

[MT07] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI, USA, Proceedings*, pages 94–103. IEEE Computer Society, 2007. [1.3](#)

[Nem04] Arkadi Nemirovski. Interior point polynomial time methods in convex programming. *Lecture notes*, 42(16):3215–3224, 2004. [1](#)

[NY83] A. Nemirovski and D. B. Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley, 1983. [1](#), [1.1](#)

[OV00] Felix Otto and Cédric Villani. Generalization of an inequality by talagrand and links with the logarithmic sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, 2000. [4.2](#)

[Sha12] Shai Shalev-Shwartz. Online learning and online convex optimization. *Found. Trends Mach. Learn.*, 4(2):107–194, 2012. [1](#), [1.1](#)

[ST74] Vladimir Sudakov and Boris Tsirelson. Extremal properties of half-spaces for spherically invariant measures. *J. Soviet Math*, 9:9–18, 1974. [1.2](#)

[Tea17] Apple Differential Privacy Team. Learning with privacy at scale. Technical report, Apple, 2017. [1](#)

[ZPFP20] Kelvin Shuangjian Zhang, Gabriel Peyré, Jalal Fadili, and Marcelo Pereyra. Wasserstein control of mirror langevin monte carlo. In Jacob D. Abernethy and Shivani Agarwal, editors, *Conference on Learning Theory, COLT 2020, 9-12 July 2020, Virtual Event [Graz, Austria]*, volume 125 of *Proceedings of Machine Learning Research*, pages 3814–3841. PMLR, 2020. [1.3](#)

A Private ERM and SCO under strong convexity

In this section, we derive our results for private ERM and SCO in general norms under the assumption that the sample losses are strongly convex. We will state our results for private ERM (Theorem 5) and SCO (Theorem 6) with respect to an arbitrary compact convex subset \mathcal{X} of a d -dimensional normed space, satisfying the following Assumption 2.

Assumption 2. *We make the following assumptions.*

- (1) *There is a compact, convex subspace $\mathcal{X} \subset \mathbb{R}^d$ equipped with a norm $\|\cdot\|_{\mathcal{X}}$.*
- (2) *There is a set Ω such that for any $s \in \Omega$, there is a function $f(\cdot; s) : \mathcal{X} \rightarrow \mathbb{R}$ which is G -Lipschitz and μ -strongly convex in $\|\cdot\|_{\mathcal{X}}$.*

Theorem 5 (Private ERM). *Under Assumption 2 and following notation (1), drawing a sample x from the density $\nu \propto \exp(-kF_{\mathcal{D}})$ for*

$$k = \frac{n^2\epsilon^2\mu}{2G^2 \log \frac{1}{2\delta}},$$

is (ϵ, δ) -differentially private, and produces x such that

$$\mathbb{E}_{x \sim \nu}[F_{\mathcal{D}}(x)] - \min_{x \in \mathcal{X}} F_{\mathcal{D}}(x) \leq \frac{2dG^2 \log \frac{1}{2\delta}}{n^2\epsilon^2\mu}.$$

Proof. Let $F_{\mathcal{D}'}$ be the realization of (1) when \mathcal{D} is replaced with a neighboring dataset \mathcal{D}' which agrees in all entries except some sample $s'_i \neq s_i$. By Assumption 2, we have $k(F_{\mathcal{D}} - F_{\mathcal{D}'})$ is $\frac{kG}{n}$ -Lipschitz, and both $kF_{\mathcal{D}}$ and $kF_{\mathcal{D}'}$ are $k\mu$ -strongly convex (all with respect to $\|\cdot\|_{\mathcal{X}}$). Hence, combining Theorem 2 and Fact 2 shows the mechanism is (ϵ, δ) -differentially private, since

$$k = \frac{n^2\epsilon^2\mu}{2G^2 \log \frac{1}{2\delta}} \implies \frac{G\sqrt{k}}{n\sqrt{\mu}} \leq \frac{\epsilon}{\sqrt{2 \log \frac{1}{2\delta}}}.$$

Let $x_{\mathcal{D}}^* := \operatorname{argmin}_{x \in \mathcal{X}} F_{\mathcal{D}}(x)$. We obtain the risk guarantee by the calculation (see Proposition 4)

$$\mathbb{E}_{x \sim \nu}[F_{\mathcal{D}}(x)] \leq F_{\mathcal{D}}(x_{\mathcal{D}}^*) + \frac{d}{k} \leq F_{\mathcal{D}}(x_{\mathcal{D}}^*) + \frac{2dG^2 \log \frac{1}{2\delta}}{n^2\epsilon^2\mu}.$$

□

Theorem 6 (Private SCO). *Under Assumption 2 and following notation (1), drawing a sample x from the density $\nu \propto \exp(-kF_{\mathcal{D}})$ for*

$$k = \frac{n^2\epsilon^2\mu}{2G^2 \log \frac{1}{2\delta}}$$

is (ϵ, δ) -differentially private, and produces x such that

$$\mathbb{E}_{\mathcal{D} \sim \pi^n, x \sim \nu}[F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} F_{\text{pop}}(x) \leq \frac{G^2}{n\mu} \left(1 + \frac{2d \log \frac{1}{2\delta}}{n\epsilon^2} \right).$$

Proof. For the given choice k, μ , the privacy proof follows identically to Theorem 3, so we focus on the risk proof. We follow the notation of Proposition 6 and let $s \sim \pi$ independently from π . By exchanging the expectation and minimum and using that $\mathbb{E}_{\mathcal{D} \sim \pi^n} F_{\mathcal{D}} = F_{\text{pop}}$,

$$\begin{aligned} \mathbb{E}_{\mathcal{D} \sim \pi^n, x \sim \nu} [F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} F_{\text{pop}}(x) &\leq \mathbb{E}_{\mathcal{D} \sim \pi^n} \left[\mathbb{E}_{x \sim \nu} [F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} [F_{\mathcal{D}}(x)] \right] \\ &\leq \mathbb{E}_{\mathcal{D} \sim \pi^n} [\mathbb{E}_{x \sim \nu} [F_{\text{pop}}(x) - F_{\mathcal{D}}(x)]] \\ &\quad + \mathbb{E}_{\mathcal{D} \sim \pi^n} \left[\mathbb{E}_{x \sim \nu} [F_{\mathcal{D}}(x)] - \min_{x \in \mathcal{X}} [F_{\mathcal{D}}(x)] \right] \\ &\leq \mathbb{E}_{\mathcal{D} \sim \pi^n} [\mathbb{E}_{x \sim \nu} [F_{\text{pop}}(x) - F_{\mathcal{D}}(x)]] + \frac{d}{k}, \end{aligned}$$

where we bounded the empirical risk in the proof of Theorem 5. Next, let ν' be the density $\propto \exp(-kF_{\mathcal{D}'})$. Our mechanism is symmetric, and hence by Proposition 6,

$$\mathbb{E} [F_{\text{pop}}(x) - F_{\mathcal{D}}(x)] = \mathbb{E} [\mathbb{E}_{x \sim \nu} [f(x; s)] - \mathbb{E}_{x \sim \nu'} [f(x; s)]]$$

where the outer expectations are over the randomness of drawing \mathcal{D}, s . Finally, for any fixed realization of \mathcal{D}, s , the densities ν, ν' satisfy the assumption of Corollary 1 with $H = \frac{G}{n}$, and $f(\cdot; s)$ is G -Lipschitz, so Corollary 1 shows that

$$\mathbb{E}_{x \sim \nu} [f(x; s)] - \mathbb{E}_{x \sim \nu'} [f(x; s)] \leq \frac{G^2}{n\mu}.$$

Combining the above three displays bounds the population risk by

$$\mathbb{E}_{\mathcal{D} \sim \pi^n, x \sim \nu} [F_{\text{pop}}(x)] - \min_{x \in \mathcal{X}} F_{\text{pop}}(x) \leq \frac{G^2}{n\mu} + \frac{d}{k} = \frac{G^2}{n\mu} \left(1 + \frac{2d \log \frac{1}{2\delta}}{n\epsilon^2} \right).$$

for our given value of k . □