

Universality of Langevin Diffusion for Private Optimization, with Applications to Sampling from Rashomon Sets

Arun Ganesh
Google Research

ARUNGANESH@GOOGLE.COM

Abhradeep Thakurta
Google DeepMind

ATHAKURTA@GOOGLE.COM

Jalaj Upadhyay
Rutgers University

JALAJ.UPADHYAY@RUTGERS.EDU

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

In this paper we provide an algorithmic framework based on Langevin diffusion (LD) and its corresponding discretizations that allow us to simultaneously obtain: i) An algorithm for sampling from the exponential mechanism (McSherry and Talwar, 2007), whose privacy analysis does not depend on convexity and which can be stopped at anytime without compromising privacy, and ii) tight uniform stability guarantees for the exponential mechanism. As a direct consequence, we obtain optimal excess empirical and population risk guarantees for (strongly) convex losses under both pure and approximate differential privacy (DP). The framework allows us to design a DP uniform sampler from the Rashomon set. Rashomon sets are widely used in interpretable and robust machine learning, understanding variable importance, and characterizing fairness.

Keywords: Differential privacy, optimization, Langevin diffusion and dynamics.

1. Introduction

Differentially private empirical risk minimization (DP-ERM) (Chaudhuri et al., 2011; Chourasia et al., 2021; Iyengar et al., 2019; Kifer et al., 2012; Bassily et al., 2014; Smith et al., 2017; Song et al., 2013, 2020; Wu et al., 2017) and *differentially private stochastic optimization* (DP-SCO) (Asi et al., 2021; Bassily et al., 2019, 2020; Feldman et al., 2020; Kulkarni et al., 2021; Gopi et al., 2022) are two of the most widely studied problems in the differential privacy (DP) literature. The optimal algorithms for either of these settings are either based on *differentially private stochastic gradient descent* (DP-SGD) (Abadi et al., 2016; Bassily et al., 2014; Song et al., 2013), or sampling from an appropriate Gibbs distribution (a.k.a. the exponential mechanism (McSherry and Talwar, 2007)). In this paper we revisit the sampling perspective of DP optimization and study its implications.

At a high level, the Gibbs distribution sampling problem is to generate a sample θ from the distribution with density proportional to $\exp(-\beta\mathcal{L}(\theta; D))$. Here β is known as the *inverse temperature*, and $\mathcal{L}(\theta; D) = \frac{1}{n} \sum_{i=1}^n \ell(\theta; d_i)$ (with $\ell : \mathbb{R}^p \times \tau \rightarrow \mathbb{R}$) is the *empirical loss function*. Our main contribution is an algorithmic framework based on Langevin Diffusion (LD) (see Figure 1) to privately (approximately) sample from a Gibbs distribution, with the following implications:

1. Our framework recovers all existing (and tight) bounds for DP-ERM (Bassily et al., 2014), and in some cases improves on the best known bounds. The framework provides tight uniform stability guarantees for the Gibbs distribution of $O(1/n)$ on strongly convex losses and $O(1/\sqrt{n})$ on mildly regularized convex losses. This is much tighter than the generic $O(\varepsilon)$ uniform stability provided by the DP guarantee (Bassily et al., 2014)¹. This allows us to obtain optimal DP-SCO bounds under both ε -DP and (ε, δ) -DP², and improves on Asi et al. (2021) for the ε -DP case. In this sense, the LD framework is universal for DP optimization.
2. The privacy guarantee of our Langevin Diffusion (LD) based algorithm does not depend on convexity of the loss function, $\mathcal{L}(\theta; D)$. Therefore, it can be used in non-convex settings without compromising privacy.
3. In the (ε, δ) -DP setting, we can release the complete trajectory of the LD until it reaches the stationary Gibbs distribution. As a direct consequence, our algorithm is “anytime” DP, meaning the privacy is not compromised even if we stop the chain early. This is a very useful property in practice; in fact, subsequent works (Shejwalkar et al., 2022; Rabanser et al., 2023) have shown that the path of LD can be used to quantify predictive uncertainty.

Sampling from Rashomon sets: Our framework also allows us to go beyond what can be achieved by known algorithms for private learning. In particular, it allows us to uniformly and privately sample from the *Rashomon set* (Breiman, 2001) (Definition 1.4), which has been extensively studied in interpretable and robust machine learning (Rudin et al., 2022), understanding spectrum of variable importance (Fisher et al., 2019), decision making (Tulabandhula and Rudin, 2013, 2014a,b), measuring underspecification (Madras et al., 2019), and characterizing fairness (Coston et al., 2021).

There is a contemporary and most closely related work of Gopi et al. (2022) to ours. We defer the comparison to Section 2.5³, and provide detailed related work in Appendix B.

1.1. Our Contributions

Our main contribution is to design a Langevin diffusion (LD) based DP sampler for the following Gibbs distribution: $\exp(-\beta \max\{\mathcal{L}(\theta; D), \psi + \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D)\})$. As we will see in Section 1.2, by setting $\psi = 0$ we obtain optimal DP-ERM/DP-SCO algorithms, and for $\psi > 0$ we obtain a DP Rashomon set sampler. In this section, we first state the main result followed by the uniform stability result for LD. We end with a discussion of discretization of our LD, that outputs a sample within δ total variation distance of the stationary distribution of LD at the same privacy/utility trade-off.

We start with the LD algorithm, described in Figure 1, which forms the building block for all the algorithms considered in this paper. Intuitively, one should think of (1) as the limit of noisy gradient descent and (2) as the limit of projected noisy gradient descent, both as step size $\eta \rightarrow 0$.

Informal Theorem 1.1 (Corresponds to Theorems 2.1 and 2.3) *Assume that the loss functions are 1-Lipschitz, and the constraint set \mathcal{C} has diameter at most one. Then there exists a LD process $\{\theta_t\}_{t \geq 0}$ with stationary distribution Θ_∞ proportional to $\exp(-\beta \max\{\mathcal{L}(\theta; D), \psi + \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D)\})$,*

-
1. The uniform stability arguments in Bassily et al. (2014) gives the $O(\varepsilon + \frac{p}{\varepsilon n})$ SCO bound for convex losses under pure-DP. This is at best $O(\sqrt{p/n})$. In contrast, we obtain the optimal $O(1/\sqrt{n} + p/\varepsilon n)$.
 2. Through out the paper we will assume $\varepsilon = O(\log(1/\delta))$.
 3. The claim of contemporarity is also supported by the authors of Gopi et al. (2022).

Langevin diffusion (LD). Let W_t be a p -dimensional Brownian motion and $\beta > 0$ be the inverse temperature. Then LD is the following stochastic differential equation:

$$d\theta_t = -\beta \nabla \mathcal{L}(\theta_t; D) \cdot dt + \sqrt{2} \cdot dW_t. \quad (1)$$

“Projected” Langevin diffusion. Sometimes, we only have the Lipschitz guarantee within a constrained set. We can also consider the following “projected” version of LD:

$$d\theta_t = -\beta \nabla \mathcal{L}(\theta_t; D) \cdot dt + \sqrt{2} \cdot dW_t - \nu_t \mu(dt), \forall t \geq 0 : \theta_t \in \mathcal{C}. \quad (2)$$

where μ is a measure supported on $\{t : \theta_t \in \partial \mathcal{C}\}$ and ν_t is an outer unit normal vector at θ_t for all such θ_t . See (Bubeck et al., 2018, Section 2.1, 3.1) for a discussion of (2).

Figure 1: (Projected) Langevin diffusion

and (i) Θ_∞ satisfies ε -DP if $\beta = O(\varepsilon n)$. (ii) If the loss function is m -strongly convex and M -smooth, then for $t = \tilde{O}_\delta \left(\frac{1}{\beta m} \right)^4$ and $\beta = \tilde{O}_\delta (m \varepsilon^2 \min\{n^2, 1/M\psi\})$, releasing $\{\theta_{t'}\}_{0 \leq t' \leq t}$ is (ε, δ) -DP and θ_T is within total variation distance (TVD) δ of Θ_∞ . Furthermore, the privacy guarantee only requires Lipschitzness and smoothness.

A key takeaway from the privacy guarantee of Theorem 1.1 is that, as ψ becomes smaller, one can run the LD at a higher β and thus provide stronger risk guarantees (In particular, in Theorem 2.4 we show an excess empirical risk bound of p/β). In fact, if $\psi \leq 1/n^2$, then the choice of β is a factor $\approx n$ higher than that in the ε -DP case. We believe the relation $\beta \propto \min\{1/\psi, n^2\}$ is necessary. (See Section 2.1 for a formal reasoning.)

For (i) in Theorem 1.1, the privacy follows from the analysis of the exponential mechanism. For (ii), we use a continuous analog of the composition theorem for Rényi-DP (see Lemma 2.2 and more discussion on our continuous time composition theorem below). To show the sampling guarantee, we show that the stationary distribution satisfies a *log Sobolev inequality* (a measure of concentration; see Definition C.19) using standard techniques known as the *Bakry-Emery criterion* and *Holley-Stroock perturbation principle* (see e.g., Appendix A of Schlichting (2019)). The convergence guarantee then follows using the results in Vempala and Wibisono (2019).

Note that part (ii) of Theorem 1.1 also shows that Θ_∞ is private via analyzing privacy of the chain rather than the stationary distribution. This, in particular, shows that the entire trajectory of the LD is private (not just the *final iterate*). This is beneficial as works such as (Shejwalkar et al., 2022; Rabanser et al., 2023) have shown improved performance and uncertainty estimation from using intermediate values. There is another advantage of analyzing the privacy of the entire chain. Unlike the sampling/utility guarantee, the privacy in (ii) does not rely on convexity, i.e. we can use it for non-convex loss functions. Furthermore, by taking $\psi = 0$ and comparing to the Gaussian mechanism (Dwork et al., 2006a), we can see our privacy guarantee is tight up to log factors.

Uniform stability of Langevin diffusion (Section 2.3): While empirical guarantees are useful in their own regards, it is often desirable to get population risk guarantee. We derive this by showing the uniform stability property of the LD on thresholded losses. This implies that any empirical accuracy guarantee for the Gibbs sampler in Theorem 1.1 also extends to population level guarantees:

4. Here and throughout this paper, $O_\delta(\cdot)$ and $\tilde{O}_\delta(\cdot)$ hides polylog factors in $1/\delta$.

Informal Theorem 1.2 (Corresponds to Theorem 2.5) *Under the same assumptions and choice of β as in Theorem 1.1, the LD (at any time, including at its stationary distribution) satisfies $O(L\sqrt{\psi/m} + L^2/(mn))$ -uniform stability.*

The proof uses the fact that the time-independent uniform stability for finite-time LD implies the same uniform stability for its Gibbs distribution, which could be of independent interest. This in particular gives optimal SCO rates under ε -DP guarantee that matches non-private bound of $O(1/\sqrt{n})$ as $\varepsilon \rightarrow \infty$, thereby, improving on the state-of-the-art results (Section 1.2).

Continuous time composition for LD (Section 2.1): We cannot use standard composition theorems of DP (Dwork and Roth, 2014) because the underlying algorithm is a continuous time process. We quantify the Rényi divergence between two LD processes when run on neighboring data sets. A similar result was also provided in (Chourasia et al., 2021, Theorem 1) only for the last iterate, θ_t . In contrast, we prove a divergence bound between the entire histories $\{\theta_{t'}\}_{0 \leq t' \leq t}$, which enables us to output weighted averages of $\theta_{t'}$'s privately. Furthermore, it is proven using only tools from the differential privacy literature and *Fatou's lemma*, providing an arguably much simpler proof.

Discretizing our LD (Section 3): In general, sampling from a continuous-time object such as LD is intractable in practice. A common technique for approximately sampling from the distribution induced by an LD is the Stochastic Gradient Langevin Dynamics (SGLD), which has been extensively studied in the literature (e.g. Dalalyan (2017); Raginsky et al. (2017); Cheng and Bartlett (2018); Cheng et al. (2018); Vempala and Wibisono (2019); Ganesh and Talwar (2020); Erdogdu et al. (2021); Chewi et al. (2021); Welling and Teh (2011); Ryffel et al. (2022); Chourasia et al. (2021)). SGLD uses T steps of noisy gradient descent with step size η , which approximates running eq. (1) for time $t = T\eta$. Using results in Chewi et al. (2021), we show that SGLD provides a private approximation (w.r.t. TVD) of our Gibbs sampler in a polynomial number of gradient oracle calls. One disadvantage of this result is that the oracle complexity has a worse dependence on the problem parameters than DP-SGD with standard hyperparameters. For example, DP-SGD's iteration complexity in Bassily et al. (2019) is constant w.r.t. dimensions as it goes to infinity, whereas our SGLD iteration complexity has a linear dependence on dimensions.

1.2. Applications of Our Algorithmic Framework

Recovering DP-ERM/DP-SCO bounds (Section 4): We show that setting $\psi = 0$ for the Gibbs sampler in Theorem 1.1 retrieves the optimal DP-ERM/DP-SCO bounds, i.e., using only the LD sampler as a primitive, one can achieve all the existing bounds for DP-ERM/DP-SCO and improve some prior results (see Table 1) as corollaries. Since most of these results are known in the literature, in the next theorem, we only present the improvement exhibited in this paper.

Informal Theorem 1.3 (Corresponds to Theorems E.1 and E.2) *For L -Lipschitz convex losses over \mathcal{C} , there exists an ε -DP algorithm with $O\left(\frac{L\|\mathcal{C}\|_2 p}{\varepsilon n} + \frac{L\|\mathcal{C}\|_2}{\sqrt{n}}\right)$ excess population risk. Further, if the loss function is also m -strongly convex, then there is an ε -DP algorithm that has excess population risk of $O\left(\frac{L^2 p^2 \log n}{m \varepsilon^2 n^2} + \frac{L^2}{mn}\right)$.*

The best prior bounds were $O\left(\frac{p \log n}{\varepsilon n} + \frac{\log^{3/2} n}{\sqrt{n}}\right)$ for convex losses, and $O\left(\frac{p^2 \log^2 n}{m \varepsilon^2 n^2} + \frac{\log^3 n}{mn}\right)$ for m -strongly convex losses, both due to Asi et al. (2021), which lacked the ability to match the optimal non-private bounds of $O(1/\sqrt{n})$ (and $O(1/mn)$ respectively) as $\varepsilon \rightarrow \infty$. (To the best of

our knowledge, this gap is inherent in the technique of [Asi et al. \(2021\)](#).) In most cases, the best DP-ERM/SCO bounds can be achieved by setting $\psi = 0$ in Theorems 1.1 and 1.2 (in the convex case, after adding a quadratic regularizer to enforce strong convexity). The second result in Theorem 1.3 is the only one which does not directly apply the Rashomon sampler’s risk bounds to the (regularized) loss. Instead, we use an *iterated exponential mechanism*, which samples from a sequence of Gibbs distributions defined over an adaptively chosen sequence of sets $\mathcal{C}_k \subset \mathcal{C}_{k-1} \subseteq \dots \subseteq \mathcal{C}_0 = \mathcal{C}$. To analyze it, we compose the privacy/utility analysis obtained independently for the Gibbs distribution over each of these sets. Our analysis simplifies the analysis in [Bassily et al. \(2014\)](#) and does not require running two different algorithms (i.e., output perturbation and exponential mechanism) to obtain the optimal trade-off. We give a summary of the iterated exponential mechanism and its analysis in Section 4 and defer a full discussion on retrieving DP-ERM/SCO bounds to Appendix E.

Additionally, for DP-ERM, in Appendix F we provide a lower bound for non-convex losses which demonstrate that, unlike for convex losses, it is not possible to achieve better utility (up to factors in $\log(1/\delta)$) than $\tilde{O}(p/\varepsilon n)$ even if we relax privacy guarantee to (ε, δ) -DP. We prove this result by appealing to the lower bound in [Steinke and Ullman \(2017\)](#).

Sampling from Rashomon Sets (Sections 2 and Appendix A): Our motivation to design uniform sampler for the Rashomon set stems from quantifying predictive multiplicity for Rashomon sets (described later) ([Hsu and Calmon, 2022](#))⁵ and that it is a strict generalization of the problems of DP-ERM and DP-SCO. Given the wide applicability of Rashomon set mentioned earlier (also see Appendix B for more discussion), we believe studying this problem will have more implications. We start with the formal definition of uniform sampling from Rashomon set.

Definition 1.4 ((λ, ψ, γ)-Rashomon sampler) Given a loss function $\mathcal{L}(\theta; D) = \frac{1}{n} \sum_{i=1}^n \ell(\theta; d_i)$ defined over a data set D , a constraint set \mathcal{C} , and a threshold ψ , an algorithm \mathcal{A} is (λ, ψ, γ) -sampler for the Rashomon set $\mathcal{G} = \left\{ \theta \in \mathcal{C} : \mathcal{L}(\theta; D) \leq \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi \right\}$ if distribution of $\theta^{\text{priv}} \leftarrow \mathcal{A}$ is λ -far in total variation distance (TVD) from a distribution π such that

1. **Uniform sampling:** The marginal distribution of π over \mathcal{G} is uniform.
2. **Maximality condition:** For any $\theta \in \mathcal{G}$ and $\theta' \notin \mathcal{G}$, the distribution π assigns probability density to θ which is greater than or equal to that of θ' .
3. **Excess empirical risk guarantee:** $\mathbb{E}_{\mathcal{A}} [\mathcal{L}(\theta^{\text{priv}}; D)] \leq \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi + \gamma$.

The *maximality* condition rules out trivial and uninteresting solutions. It ensures that (ignoring privacy constraints), it is possible to combine the Rashomon sampler with rejection sampling to efficiently get an uniform sample from \mathcal{G} . In particular, we show in Theorem D.6 that, if $\psi = \omega(p\gamma)$, then the Gibbs sampler has $1 - o(1)$ probability of hitting the Rashomon set for the values of β stated in Theorem 1.1. We presented Definition 1.4 with respect to empirical accuracy guarantee for the ease of presentation. One can also define *excess population risk guarantee*: $\mathbb{E}_{\mathcal{A}, d \sim \mathcal{D}} [\ell(\theta^{\text{priv}}; d)] \leq \psi + \min_{\theta \in \mathcal{C}} \mathbb{E}_{d \sim \mathcal{D}} [\ell(\theta; d)] + \gamma'$, where γ' is the *population level slack* and \mathcal{D} is the distribution from which the data samples in the data set D are drawn i.i.d.

5. Predictive multiplicity refers to models that achieve statistically-indistinguishable performance on a test set assign wildly different predictions to an input sample. Therefore, if we naively pick a model from a Rashomon set, it can have highly disparate impact on the predictions on an individual test sample resulting in unfair and potentially individual level harm ([Smith, 2020](#); [Creel and Hellman, 2022](#)).

In Theorems 2.1 and 2.3, we show that our LD based Gibbs sampler from Theorem 1.1 is a $(0, \psi, p/\beta)$ -Rashomon sampler. Furthermore, in Theorem 2.6, we show that the excess population risk (i.e., $\psi + \gamma'$) for the Rashomon set sampler, is bounded by $\frac{p}{\beta} + O\left(\psi + L\sqrt{\psi/m} + L^2/mn\right)$. We give a summary of the bounds in Table 2.

We next discuss the use case of predictive multiplicity eluded above. Consider a classification problem with c classes and Δ_c be the probability simplex. Let (y, \mathbf{x}) be a test sample and let $f(\mathbf{x}; \theta) \in \Delta_c$ be the prediction function that provides a probability distribution across the c classes. The objective is to estimate the variance, $\text{Var}(f(\mathbf{x}, \theta))$, where $\theta \sim_{\text{unif}} \mathcal{G}$. One can get its approximate estimate by sampling $\{\theta_1, \dots, \theta_k\} \sim_{\text{iid}} \mathcal{G}$, and estimating the standard deviation of $\{f(\mathbf{x}, \theta_1), \dots, f(\mathbf{x}, \theta_k)\}$. While there is a privacy cost in sampling k models, (Brawner and Honaker, 2018, Section 6.4) shows that the variance estimation comes at *no additional privacy cost* when compared to that of outputting a single model. Given the standard deviation, one can then decide on the class to predict for \mathbf{x} , either via further randomization, or other strategy, including more sophisticated measures like *Rashomon capacity* (Hsu and Calmon, 2022). We leave exploring their DP variants for future research.

Organization: For the ease of presentation and owing to the generality of the Rashomon set sampling problem, we state all our main results in that context in Section 2. In Section 3 we provide the details for the discretization of the LD algorithm. In Section 4 we provide a summary of the DP-ERM/DP-SCO results obtained by setting $\psi = 0$ for the Rashomon set sampler. Finally, in Section 5 we end with discussions and future directions. We enumerate our notations in Table 3.

2. Rashomon Set Sampling

In this section, we provide the privacy analysis (Sections 2.1 and 2.4)⁶, and the utility analysis (Sections 2.2 and 2.3) for the Rashomon set sampler based on the Gibbs distribution proportional to $\exp(-\beta\tilde{\mathcal{L}}(\theta; D))$, where $\tilde{\mathcal{L}}(\theta; D) := \max\{\mathcal{L}(\theta; D), \psi + \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D)\}$. Both the privacy and the utility guarantee primarily depend on the choice of the inverse temperature β . Under (ε, δ) -DP, we can operate with higher values of β than in the case of ε -DP (see Theorems 2.1 and 2.3): this translates to a better utility for certain choices of ψ .

In our algorithmic version of the Gibbs distribution, we instantiate it with the LD described in Figure 1. The LD that is used to instantiate the Gibbs distribution in the ε -DP case unfortunately requires running the algorithm for $t \rightarrow \infty$ to reach within ε of the stationary distribution in terms of ℓ_∞ -distance⁷. However, if we are willing to tolerate (ε, δ) -DP guarantee, then the LD can be run in time $\approx \frac{\log(1/\delta)}{\beta m}$, where m is the strong convexity parameter. A standard discretization of the LD we use in this section (via SGLD), that makes the algorithm run on a finite precision machine, can be found in Section 3. All the missing proofs of this section appear in Appendix D.

2.1. Privacy Guarantees and Running Time

The privacy in ε -DP case follows from the fact that $\max\{\mathcal{L}(\theta; D), \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi\}$ cannot change by more than $L \|\mathcal{C}\|_2 / n$ when we change one data point because each $\ell \in [0, L \|\mathcal{C}\|_2]$.

6. The analysis in Section 2.1 is via analyzing the privacy of the path of the LD that converges to the Gibbs distribution, and an alternate one in Section 2.4 via directly analyzing the privacy of the Gibbs distribution.

7. Bassily et al. (2014); Mangoubi and Vishnoi (2021) provides rejection sampling based polytime algorithms, but lack the generalization properties of LD.

Theorem 2.1 (ε -DP sampler; Theorem 3.1 of Bassily et al. (2014)) *Suppose we have a constraint set \mathcal{C} and a loss function $\ell(\theta; d)$ such that for all d , $\ell(\cdot; d) \in \mathbb{R}^+$, and is L -Lipschitz within \mathcal{C} . Then, sampling θ^{priv} from the Gibbs distribution $\exp(-\beta \max\{\mathcal{L}(\theta; D), \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi\})$ is ε -differentially private for $\beta = O(\frac{\varepsilon n}{L\|\mathcal{C}\|_2})$ and all ψ .*

(ε, δ) -DP Sampler: If $\ell(\theta; \cdot)$ is m -strongly convex and our goal is (ε, δ) -DP, we can use larger values of β . However, for the settings when $\psi > 0$, we would additionally require $\ell(\theta; \cdot)$ to be M -smooth. We first need a ‘‘continuous’’ composition theorem that bounds the Rényi divergence between two instances of LDs run on adjacent databases:

Lemma 2.2 *Let θ_0, θ'_0 have the same distribution Θ_0 , θ_t be the solution to (2) given θ_0 and data set D (and correspondingly θ'_0 and θ'_t for a data set D'). Let $\Theta_{[0,t]}$ ($\Theta'_{[0,t]}$) be the distribution of the trajectory of LD $\{\theta_{t'}\}_{t' \in [0,t]}$ ($\{\theta'_{t'}\}_{t' \in [0,t]}$, respectively). Suppose we have that $\|\nabla \mathcal{L}(\theta; D) - \nabla \mathcal{L}(\theta; D')\|_2 \leq \Delta$ for all θ . Then $\forall \alpha \geq 1$: $R_\alpha(\Theta_{[0,t]}, \Theta'_{[0,t]}) \leq \frac{\alpha \beta^2 \Delta^2 t}{4}$.*

The idea behind the proof is to use a bound on the divergence between Gaussians and RDP composition to provide a bound on the divergence between the projected noisy gradient descents on datasets D and D' . Then, taking the limit as the step size in gradient descent goes to 0 and applying Fatou’s lemma (Lemma C.14), we get the bound above. A full proof is deferred to Appendix D.1. Rényi divergence bounds imply (ε, δ) -DP privacy guarantees (Fact C.4), which we use to prove the following theorem in Appendix D.2:

Theorem 2.3 ((ε, δ) -DP sampler) *Suppose we have a constraint set \mathcal{C} and a loss function $\ell(\theta; d)$ such that for all d , $\ell(\cdot; d) \in \mathbb{R}^+$ is m -strongly convex, M -smooth, and is L -Lipschitz within \mathcal{C} . For $t = \tilde{O}_\delta\left(\frac{1}{\beta m}\right)$ and an appropriate choice of θ_0 , let Θ_t be the distribution over θ_t given by running (2) on $\max\{\mathcal{L}(\theta; D), \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi\}$. Then Θ_t is within total variation distance δ of the Gibbs distribution on $\max\{\mathcal{L}(\theta; D), \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi\}$, and is (ε, δ) -DP if any of the following holds: (i) $\psi \in \left(\frac{2L^2}{Mn^2}, \frac{L\|\mathcal{C}\|_2}{2}\right]$, and $\beta = \tilde{\Theta}\left(\frac{\varepsilon^2(m/M)}{\log(L\|\mathcal{C}\|_2/\delta^2) \log(1/\delta)} \cdot \frac{1}{\psi}\right)$, (ii) $\psi \leq \frac{2L^2}{Mn^2}$ and $\beta = \tilde{\Theta}\left(\frac{\varepsilon^2 n^2 m}{L^2 \log(L\|\mathcal{C}\|_2/\delta^2) \log(1/\delta)}\right)$. Furthermore, the privacy holds even if we release the entire trajectory $\{\Theta_{t'}\}_{t' \in [0,t]}$, and even without convexity.*

If ψ has a dependence $o(1/n)$ on n , this gives a better dependence of β on n than Theorem 2.1. For privacy, we first bound the sensitivity of the thresholded loss by $O(\max\{L/n, \sqrt{M\psi}\})$. We then appeal to Lemma 2.2 and the translation from Rényi divergence bounds to (ε, δ) -DP bounds (Fact C.4). Since Lemma 2.2 does not require convexity and allows for releasing the entire trajectory, the same is true for the privacy guarantee of Theorem 2.3. The sampling guarantee is given by showing that the Gibbs distribution satisfies log-Sobolev inequality (LSI; see Definition C.19), a measure of concentration. This implies convergence of LD to the Gibbs distribution by results in Vempala and Wibisono (2019). A few comments are in order about the theorem.

Direct analysis of Gibbs distribution: Note that by triangle inequality, Theorem 2.3 also implies that, under strong convexity, the Gibbs distribution is $(\varepsilon, 2\delta)$ -DP. We can also analyze the Gibbs distribution directly: Since the Gibbs distribution satisfies LSI, one can obtain an isoperimetric inequality for the probability measure via Ledoux (1999). Using this isoperimetric inequality and the coupling between Gibbs distributions that we later state in Theorem 2.5, one can obtain a bound

on β that improves Theorem 2.3 by log factors, giving tight bounds in the case $\psi = O\left(\frac{L^2}{Mn^2}\right)$. However, this privacy proof relies heavily on the convexity and does not give an “anytime” private sampler like our LD-based proof. For the sake of completeness, we provide a proof of this result in Section 2.4 (see Theorem 2.7 for a precise statement).

Dependence of β on ψ : We believe the relation $\beta \propto \min\{1/\psi, n^2\}$ in Theorem 2.3 is necessary. To see this, consider mean estimation on an all zero databases D and on neighboring D' that contain just a single 1. Fix $\varepsilon = 1$ for simplicity. The Gibbs samplers for these datasets have densities proportional to $\exp(-\beta \min\{\|x\|_2^2/2, \psi\})$ and $\exp(-\beta \min\{\|x - 1/n\|_2^2/2, \psi\})$, respectively. For the case $\psi = 0$, this is just a Gaussian mechanism with sensitivity $1/n$ and variance $1/\beta$, and one way to argue this mechanism is differentially private is to provide a tail bound on x , which gives a tail bound on the privacy loss. For $\psi > 0$, the unnormalized density decreases by a factor of $\exp(-\Omega(\beta\psi))$ in the range $[-\sqrt{2\psi}, \sqrt{2\psi}]$ when we apply thresholding to the loss function. In turn, if most of the probability mass is in this interval, a tail bound on events outside this interval can get worse by a factor of $\exp(-\Omega(\beta\psi))$. Suppose we use $\beta \approx n^2$, which gives $(1, \delta)$ -DP for $\psi = 0$. Then the interval $[-\sqrt{2\psi}, \sqrt{2\psi}]$ contains all points within $\sqrt{\psi/\beta} \approx n\sqrt{\psi}$ standard deviations of the mean. So this interval contains most of the probability mass of the mechanism when $\psi = \tilde{\Omega}(1/n^2)$. This roughly matches the “transition point” in Theorem 2.3. In order for a tail bound that holds w.p. $1 - \delta$ on the Gaussian mechanism to be non-vacuous after thresholding, we need $\beta\psi = O(\log(1/\delta))$. This roughly matches our choice of β after the “transition point.”

Running time: In the following discussion, we will assume $L = \|\mathcal{C}\|_2 = \Theta(1)$ for brevity. The LD for the ε -DP sampler (mentioned in Theorem 2.1) runs for time $t \rightarrow \infty$ to obtain the privacy/utility trade-off obtained by Theorems 2.1 and 2.4. However, if we are willing to tolerate (ε, δ) -DP, then assuming the loss function $\mathcal{L}(\theta; D)$ is m -strongly convex, one can obtain asymptotically the same privacy/utility trade-off, and run for time $t_\beta = \tilde{O}_\delta\left(\max\left\{\frac{1}{\beta m}, \frac{\psi}{m}\right\}\right)$. This follows from Lemma D.5. Setting $\beta = \varepsilon n$ from Theorem 2.1, we have $t_\beta = \tilde{O}_\delta\left(\max\left\{\frac{1}{m\varepsilon n}, \frac{\psi}{m}\right\}\right)$. The (ε, δ) -DP sampler from Theorem 2.3 also runs in time t_β , where β is chosen based on Theorem 2.3. Hence, to obtain the best (ε, δ) -DP Rashomon sampler, one needs $t_\beta = \tilde{O}_\delta\left(\max\left\{\frac{1}{m\varepsilon^2 n^2}, \frac{\varepsilon^2 M\psi}{m}\right\}\right)$. In Section 3, we discuss how one can obtain an approximate sampler using only discrete noisy gradient steps.

Probability of hitting the Rashomon Set: The Gibbs measure induced by the LD algorithm (i.e., $\exp(-\beta \max\{\mathcal{L}(\theta; D), \psi + \min_{\theta \in \mathcal{C}} \mathcal{L}(\theta; D)\})$) trivially satisfies the maximality condition for the Rashomon sampler. In Theorem D.6, we show that the ε -DP Rashomon set sampler hits the set \mathcal{G} with probability at least $1 - o(1)$, if $\psi = \tilde{\omega}(p^2/(\varepsilon n))$. (Here, we assumed all other parameters to be constant.) Although all our Rashomon set samplers are forced to provide non-zero probability mass on the Rashomon set due to the maximality condition, it is unclear how to obtain a similar guarantee as Theorem D.6 for our (ε, δ) -DP sampler. We leave it as an open question.

2.2. Excess Empirical Risk Guarantees

The excess empirical risk guarantee follows from Theorem 3.2 in Bassily et al. (2014):

Theorem 2.4 *Assume each loss function is convex. Then sampling θ^{priv} from the Gibbs distribution $\exp(-\beta \max\{\mathcal{L}(\theta; D), \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi\})$ is a $(0, \psi, p/\beta)$ -Rashomon sampler.*

Using Theorem 2.4, to get the best risk bound it suffices to find the largest value of β that preserves DP under given assumptions on the loss function. Theorem 2.3 suggest that when $\psi = 0$,

the setting of β that is required for (ε, δ) -DP is independent of the smoothness parameter M , and hence can also be applied to non-smooth functions. Combining Theorems 2.1, 2.3, and 2.4 we get the existence of the following three Rashomon samplers: i) $\left(0, \psi, \tilde{O}\left(p \cdot \frac{L\|C\|_2}{\varepsilon n}\right)\right)$ -sampler for all ψ and is ε -DP, ii) $\left(0, \psi, \tilde{O}\left(p\psi \cdot \frac{M \log(L\|C\|_2/\delta^2) \log(1/\delta)}{m\varepsilon^2}\right)\right)$ -sampler when $\psi \in \left(\frac{2L^2}{Mn^2}, \frac{L\|C\|_2}{2}\right]$ and is (ε, δ) -DP, and iii) $\left(0, \psi, \tilde{O}\left(p \cdot \frac{L^2 \log(L\|C\|_2/\delta^2) \log(1/\delta)}{m\varepsilon^2 n^2}\right)\right)$ -sampler when $\psi \leq \frac{2L^2}{Mn^2}$ and is (ε, δ) -DP. In the (ε, δ) -DP setting if instead of sampling from the Gibbs distribution, we operate with the LD, then in the Rashomon sampler we set $\lambda = \delta$ instead of $\lambda = 0$ as done until now. The loss guarantee worsens by at most $O(\delta L^2/m)$ if we use LD instead of the Gibbs sampler, so the earlier bounds remain unchanged for $\delta = O(1/\varepsilon^2 n^2)$.

Ignoring the polylogarithmic terms, if $\psi = \tilde{\Omega}\left(\frac{\varepsilon L\|C\|_2}{(M/m)} \cdot \frac{1}{n}\right)$, then our privacy analysis of the ε -DP Rashomon sampler gives a better bound on β than our analysis of the (ε, δ) -DP sampler, and vice versa when $\psi = \tilde{o}\left(\frac{\varepsilon L\|C\|_2}{(M/m)} \cdot \frac{1}{n}\right)$. As previously mentioned, we believe that the weaker bound on β for higher values of ψ in our analysis of the (ε, δ) -DP sampler is fundamental to the problem.

2.3. Excess Population Risk Guarantees

For the same Rashomon samplers, we can derive bounds on their population risks under strong convexity. We give a bound on uniform stability (see Definition C.15) of the Gibbs distribution:

Theorem 2.5 *Suppose we have \mathcal{C} and $\ell(\cdot; d) \in \mathbb{R}^+$ such that, for all d , $\ell(\cdot; d)$ is m -strongly convex, and is L -Lipschitz within \mathcal{C} . Then sampling θ^{priv} from the Gibbs distribution proportional to $\exp(-\beta \max\{\mathcal{L}(\theta; D), \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi\})$ satisfies $\left(4L\sqrt{\frac{2\psi}{m}} + \frac{2L^2}{mn}\right)$ -uniform stability.*

Proof Let θ_t, θ'_t be the solutions to running (2) from the same initialization on $\mathcal{L}(\theta; D)$ and $\mathcal{L}(\theta; D')$, respectively. Similarly, let $\tilde{\theta}_t$ be the solutions to running eq. (2) on $\max\left\{\psi + \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D), \mathcal{L}(\theta; D)\right\}$ and $\tilde{\theta}'_t$ be the solutions to running eq. (2) on $\max\left\{\psi + \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D'), \mathcal{L}(\theta; D')\right\}$. Let W_∞ denote the the ∞ -Wasserstein distance. Then we will show that for all t , $W_\infty(\tilde{\theta}_t, \tilde{\theta}'_t) \leq 4\sqrt{\frac{2\psi}{m}} + \frac{2L}{mn}$. Taking the limit as t goes to infinity, we get a ∞ -Wasserstein distance bound between the Gibbs samplers. The theorem now follows by using L -Lipschitzness.

We now show the desired bound on $W_\infty(\tilde{\theta}_t, \tilde{\theta}'_t)$. By the triangle inequality, we have $W_\infty(\tilde{\theta}_t, \tilde{\theta}'_t) \leq W_\infty(\tilde{\theta}_t, \theta_t) + W_\infty(\tilde{\theta}'_t, \theta'_t) + W_\infty(\theta_t, \theta'_t)$. So, it suffices to prove that for all t , $W_\infty(\tilde{\theta}_t, \theta_t) \leq 2\sqrt{\frac{2\psi}{m}}$, $W_\infty(\tilde{\theta}'_t, \theta'_t) \leq 2\sqrt{\frac{2\psi}{m}}$, and $W_\infty(\theta_t, \theta'_t) \leq \frac{2L}{mn}$.

We first prove the desired bound on $W_\infty(\tilde{\theta}_t, \theta_t)$, the bound on $W_\infty(\tilde{\theta}'_t, \theta'_t)$ follows identically.

Bounding $W_\infty(\tilde{\theta}_t, \theta_t)$: We show that conditioned on the value of a *shared* Brownian motion W_t , $\|\theta_t - \tilde{\theta}_t\|_2 \leq 2\sqrt{\frac{2\psi}{m}}$ holds deterministically, which implies the desired Wasserstein distance bound. We split $[0, \infty)$ into intervals of maximal length for which one of the following three holds throughout each interval: (i) $\tilde{\theta}_t \notin \mathcal{G}$, [(ii) $\tilde{\theta}_t \in \mathcal{G}, \theta_t \in \mathcal{G}$, and (iii) $\tilde{\theta}_t \in \mathcal{G}, \theta_t \notin \mathcal{G}$.

In case (ii), by strong convexity throughout the interval we have $\|\theta_t - \tilde{\theta}_t\|_2 \leq \|\mathcal{G}\|_2 \leq 2\sqrt{\frac{2\psi}{m}}$. So it suffices to show that in cases (i) and (iii), $\|\theta_t - \tilde{\theta}_t\|_2$ is non-increasing throughout the interval. Then the desired Wasserstein distance bound holds by induction, since initially $\theta_0 = \tilde{\theta}_0$.

In case (i), since \mathcal{L} is convex, projection is contractive (which implies $\frac{d\|\theta_t - \tilde{\theta}_t\|_2^2}{dt}$ can only increase if we use (1) instead of (2)), and that we are using a shared Brownian motion W_t , we have

$$\frac{1}{2} \cdot \frac{d}{dt} \|\theta_t - \tilde{\theta}_t\|_2^2 \leq \left\langle \frac{d\theta_t}{dt} - \frac{d\tilde{\theta}_t}{dt}, \theta_t - \tilde{\theta}_t \right\rangle = \beta \langle -(\nabla \mathcal{L}(\theta_t; D) - \nabla \mathcal{L}(\tilde{\theta}_t; D)), \theta_t - \tilde{\theta}_t \rangle \leq 0.$$

Similarly, in case (iii), by convexity and since $\tilde{\theta}_t$ is in the Rashomon set and θ_t is not (or $\mathcal{L}(\tilde{\theta}_t; D) \leq \mathcal{L}(\theta_t; D)$), we have

$$\frac{1}{2} \cdot \frac{d\|\theta_t - \tilde{\theta}_t\|_2^2}{dt} \leq \left\langle \frac{d\theta_t}{dt} - \frac{d\tilde{\theta}_t}{dt}, \theta_t - \tilde{\theta}_t \right\rangle = \beta \langle \nabla \mathcal{L}(\theta_t; D), \tilde{\theta}_t - \theta_t \rangle \leq \beta(\mathcal{L}(\tilde{\theta}_t; D) - \mathcal{L}(\theta_t; D)) \leq 0.$$

Bounding $W_\infty(\theta_t, \theta'_t)$: We again show that conditioned on the value of a *shared* Brownian motion W_t , $\|\theta_t - \theta'_t\|_2 \leq \frac{2L}{mn}$ holds deterministically. We again use the fact that projection is contractive, so we can consider using (1) instead of (2). Then by m -strong convexity and L -Lipschitzness:

$$\begin{aligned} \frac{1}{2} \cdot \frac{d\|\theta_t - \theta'_t\|_2^2}{dt} &\leq \left\langle \frac{d\theta_t}{dt} - \frac{d\theta'_t}{dt}, \theta_t - \theta'_t \right\rangle = -\beta \langle \nabla \mathcal{L}(\theta_t; D) - \nabla(\mathcal{L}(\theta'_t, D')), \theta_t - \theta'_t \rangle \\ &= -\beta (\langle \nabla \mathcal{L}(\theta_t; D) - \nabla(\mathcal{L}(\theta'_t, D)), \theta_t - \theta'_t \rangle + \langle \nabla \mathcal{L}(\theta'_t; D) - \nabla(\mathcal{L}(\theta'_t, D')), \theta_t - \theta'_t \rangle) \\ &\leq \beta \left(-m \|\theta_t - \theta'_t\|_2^2 + \frac{2L}{n} \|\theta_t - \theta'_t\|_2 \right). \end{aligned} \quad (3)$$

Then, if $\|\theta_t - \theta'_t\|_2 > \frac{2L}{mn}$, then we have (3) < 0 . That is, we have $\frac{d\|\theta_t - \theta'_t\|_2^2}{dt} < 0$. This implies the desired Wasserstein distance bound completing the proof of Theorem 2.5. \blacksquare

Remark 1 We note that in the second part of the proof, one can instead take the L^2/mn -uniform stability of (noisy) gradient descent on strongly convex losses proved in [Hardt et al. \(2016\)](#), and then take the limit as $\eta \rightarrow 0, T \rightarrow \infty$ to conclude the same uniform stability bound holds for the Gibbs distribution, rather than appeal to the derivative of the distance between θ_t and θ'_t .

It is straightforward to see that a $(0, \psi, \gamma)$ -Rashomon sampler has expected excess empirical risk at most $\psi + \gamma$. Then, an algorithm's excess population risk is at most its excess empirical risk plus its uniform stability (see Lemma C.16), giving us the following result:

Theorem 2.6 Assume that each of the individual loss function $\ell(\theta; \cdot) \in \mathbb{R}^+$ is m -strongly convex and L -Lipschitz within the constraint set \mathcal{C} . Then

- There exists an ε -DP rashomon sampler with threshold ψ and excess population risk of $O\left(\frac{L\|\mathcal{C}\|_{2^p}}{\varepsilon n} + \psi + L\sqrt{\frac{\psi}{m} + \frac{L^2}{mn}}\right)$.
- There exists an (ε, δ) -DP rashomon sampler with threshold ψ and excess population risk (assuming M -smoothness)

$$\tilde{O}\left(p \max\left\{\psi, \frac{L^2}{Mn^2}\right\} \cdot \frac{M \log(L\|\mathcal{C}\|_2/\delta^2) \log(1/\delta)}{m\varepsilon^2} + \psi + L\sqrt{\frac{\psi}{m} + \frac{L^2}{mn}}\right).$$

2.4. Interlude: Proof of Privacy for the Gibbs Sampler via Isoperimetry

Theorem 2.7 Suppose ℓ is m -strongly convex, M smooth, and L -Lipschitz within \mathcal{C} . Let $\mathcal{L}'(\theta; D) := \max\{\mathcal{L}(\theta; D), \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi\}$. Then for $\beta = O\left(\frac{\varepsilon m}{\max\{L^2/n^2, M\psi\} \log(1/\delta)}\right)$, sampling from the distribution with density proportional to $\exp(-\beta \mathcal{L}'(\theta; D)) \cdot \mathbf{1}(\theta \in \mathcal{C})$ satisfies (ε, δ) -DP.

Proof Let P be the probability measure for the distribution induced by D and Q be the same for D' . Let $g(\theta) = \log(P(\theta)/Q(\theta))$. Then g is a $6\beta \max\{L/n, \sqrt{M\psi}\}$ -Lipschitz function. It suffices to show that $\Pr_{\theta \sim P}[g(\theta) > \varepsilon] \leq \delta$. We first bound $E_{\theta \sim P}[g(\theta)]$ as a function of β . This is simply the KL divergence between P, Q , which must be non-negative. By symmetry, $E_{\theta \sim Q}[g(\theta)]$ is non-positive. In addition, we can bound the difference between these two expressions: in the proof of Theorem 2.5, we showed that the ∞ -Wasserstein distance between P and Q is at most $4\sqrt{\frac{2\psi}{m} + \frac{2L}{mn}}$. Then, by Lipschitzness of g ,

$$\begin{aligned} E_{\theta \sim P}[g(\theta)] &\leq E_{\theta \sim Q}[g(\theta)] + 6\beta \max\{L/n, \sqrt{M\psi}\} \cdot \left(4\sqrt{\frac{2\psi}{m} + \frac{2L}{mn}}\right) \\ &\leq 6\beta \max\{L/n, \sqrt{M\psi}\} \cdot \left(4\sqrt{\frac{2\psi}{m} + \frac{2L}{mn}}\right) \leq 36\sqrt{2}\beta \max\left\{\frac{L^2}{mn^2}, \sqrt{\frac{M}{m}}\psi\right\}. \end{aligned} \quad (4)$$

We next give a high probability bound on g as a function of β . P satisfies LSI with constant $\beta m \exp(-\beta\psi)$ by Lemma D.4. Then by Proposition 2.3 in Ledoux (1999), plugging in our LSI constant and Lipschitzness bound for g we have:

$$\Pr_{\theta \sim P}[g(\theta) > \mathbb{E}_{\theta \sim P}[g(\theta)] + \beta r] \leq \exp\left(-\frac{\beta m \exp(-\beta\psi) r^2}{18 \max\{L^2/n^2, M\psi/2\}}\right) \quad (5)$$

Setting $r = \frac{3\sqrt{2} \max\{L/n, \sqrt{M\psi/2}\}}{\sqrt{\beta m \exp(-\beta\psi/2)}} \sqrt{\log(1/\delta)}$, the right hand side becomes δ . Setting $\beta \leq \frac{\varepsilon m}{c \max\{L^2/n^2, M\psi\} \log(1/\delta)}$, where c is a sufficiently large constant, gives that the upper bound in (4) is at most $\varepsilon/2$ and $\beta r \leq \varepsilon/2$. Here, we use that for this choice of β , $\beta\psi \leq 1$. Plugging in to (5) we get $\Pr_{\theta \sim P}[g(\theta) > \varepsilon] \leq \delta$ exactly as desired. This completes the proof of Theorem 2.7. \blacksquare

2.5. Comparison with Contemporary Work of Gopi et al. (2022)

All our results for DP-ERM and DP-SCO, and their relation with Langevin diffusion is contemporary to Gopi et al. (2022) (which is also formally acknowledged by Gopi et al. (2022)). Our results specific to Rashomon sets (i.e., with setting $\psi > 0$) and SGLD are subsequent to their work. In

the following, we discuss the technical differences between two works, and highlight the settings in which one might be better over the other.

On the privacy aspect: Since [Gopi et al. \(2022\)](#) gives Gaussian DP guarantees, which do not translate to ε -DP guarantees, we will restrict the discussion specific to (ε, δ) -DP. Unlike [Gopi et al. \(2022\)](#), our privacy guarantee (for LD as well as SGLD) is independent of convexity. This is highly desirable for broader applicability in settings where the loss function may not be globally convex, but has local convexity properties (e.g., losses emerging from deep learning settings ([Izmailov et al., 2018](#); [Keskar et al., 2016](#))). Our privacy holds for the entire path of the optimization. In contrast, [Gopi et al. \(2022\)](#) only guarantee privacy for the final model release. In particular, this implies that we can stop and output the model at any point of the optimization trajectory. While this might not yield optimal model, the privacy is never compromised. The ability to release the path has been used in subsequent works ([Shejwalkar et al., 2022](#); [Rabanser et al., 2023](#)) for uncertainty quantification.

On proof technique: In [Theorem 2.7](#), we demonstrated that the proof technique of [Gopi et al. \(2022\)](#) can be extended to obtain privacy for the Gibbs distribution used in Rashomon set sampling problem (i.e., when $\psi > 0$). However, unlike [Lemma 2.2](#), the proof of [Theorem 2.7](#) requires an isoperimetric inequality based on LSI ([Ledoux, 1999](#)), and bounding the LSI constant for the Rashomon set sampler. In addition, while both our work and ([Gopi et al., 2022](#)) shows uniform stability property for the Gibbs distribution, our proof is arguably simpler: [Theorem 2.5](#) uses only well-known properties of gradient descent and avoids tools such as the Talagrand transportation inequality. Since the proof techniques of [Gopi et al. \(2022\)](#) and ours are completely independent, we believe that both the techniques will find adoption in subsequent works on DP optimization.

On run time: In the DP-SCO setting, the gradient oracle complexity of [Gopi et al. \(2022\)](#) is better than our SGLD discretization of the Langevin diffusion. In particular, they achieve oracle complexity with logarithmic dependence on p and the total variation distance error δ , whereas we have a dependence p/δ^2 on these two parameters. Furthermore, we require exact gradient oracles, whereas they only require an unbiased function oracle.

3. Discrete Approximation of the Langevin Diffusion and SGLD

While exactly sampling from the Gibbs distribution may be computationally intractable, under some additional assumptions, a number of papers study polynomial-time algorithms for approximate sampling from the Gibbs distribution with various metric of approximation, such as Renyi divergence ([Vempala and Wibisono, 2019](#); [Ganesh and Talwar, 2020](#); [Erdogdu et al., 2021](#); [Chewi et al., 2021](#)) and ∞ -divergence ([Bassily et al., 2014](#); [Mangoubi and Vishnoi, 2021](#)). For sampling from the distribution $\exp(-\beta\mathcal{L})$, a popular approximate sampler is the Stochastic Gradient Langevin Dynamics (SGLD). The SGLD approximates a finite-time solution to [\(1\)](#) via the discrete updates: $\theta_{t+1} = \theta_t - \eta\beta\nabla\mathcal{L}(\theta_t; D) + \xi_t$, where $\xi_t \sim \mathcal{N}(0, 2\eta I_p)$. This update can be seen as equivalent to [eq. \(1\)](#) if we use $\nabla\mathcal{L}(\theta_{\eta\lfloor t/\eta\rfloor}; D)$ instead of $\nabla\mathcal{L}(\theta_t; D)$, i.e., instead of continuously, update the gradient drift term in [eq. \(1\)](#) every η time.

Many works study SGLD as an approximate sampler, and show that, for sufficiently small η and large T , θ_T is an approximate sample from the stationary distribution $\exp(-\beta\mathcal{L})$ of [\(1\)](#). For an appropriate definition of approximation, such as total variation distance, privacy of the stationary distribution then implies privacy of SGLD. SGLD also converges in polynomial time for stronger notions of convergence such as the Rényi divergence, but since divergences are “one-directional”

bounds, whereas privacy requires “bi-directional” bounds, these results combined with the privacy of the stationary distribution do not necessarily ensure privacy of SGLD. One could also use the result of [Altschuler and Talwar \(2022c\)](#), which shows SGLD is an approximate sampler for the *discrete chain’s* stationary distribution. However, one would then need to show a bound on the bias due to the discretization that preserves privacy and utility. Instead, we appeal to the privacy of the noisy gradient steps taken since SGLD is just a reparameterization of DP-SGD. Using the composition theorem for Rényi divergences (Fact C.8) and translation from Rényi divergence bounds to (ε, δ) -DP (Fact C.4), we get the following.

Lemma 3.1 *If $\|\nabla\mathcal{L}(\theta; D) - \nabla\mathcal{L}(\theta; D')\|_2 \leq \Delta$ and $\beta \leq \frac{2\varepsilon}{\Delta\sqrt{T\eta\log(1/\delta)}}$, then outputting θ_T sampled from SGLD is (ε, δ) -DP.*

This statement matches Theorem 2.2 as $\eta \rightarrow 0$. For the utility guarantee of SGLD as an approximate sampler, we use results from [Chewi et al. \(2021\)](#), though these require smoothness. The thresholded loss does not satisfy smoothness for $\psi > 0$. In Appendix D.4, we show that using standard smoothing techniques, one can still get a discrete approximate sampler for the Gibbs distribution of a smoothed version of the thresholded loss, which implies the following sampler:

Theorem 3.2 *Suppose $\|\nabla\mathcal{L}(\theta; D) - \nabla\mathcal{L}(\theta; D')\|_2 \leq \Delta$, and that each individual loss function l is m -strongly convex and M -smooth. Let $0 \leq \lambda \leq \sqrt{\frac{\psi}{Mp}}$ and Q be the (unconstrained) Gibbs distribution of $\beta\tilde{\mathcal{L}}'$, where $\tilde{\mathcal{L}}' := \mathbb{E}_{\xi \sim N(0, \lambda^2 \mathbb{I}_p)} \left[\min\{\mathcal{L}(\theta + \xi; D), \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi\} \right]$. Then for $\beta = \tilde{O}\left(\frac{\varepsilon^2 m}{\max\{\Delta^2, M\psi\} \log^2(1/\delta)}\right)$, $T = \tilde{\Omega}\left(\frac{p(M^2 + \frac{M\psi}{\lambda^2}) \max\{\Delta^4, M^2\psi^2\}}{\varepsilon^4 m^4 \delta^2}\right)$ there exists an algorithm using T iterations that satisfies (ε, δ) -DP and returns a sample from a distribution whose TVD from Q is δ . Furthermore, the privacy guarantee holds even without convexity. In particular, if $\psi = 0$, for $\lambda = 0$ the bounds are $\beta = \tilde{O}\left(\frac{\varepsilon^2 m}{\Delta^2 \log^2(1/\delta)}\right)$, $T = \tilde{\Omega}\left(\frac{pM^2\Delta^4}{\varepsilon^4 m^4 \delta^2}\right)$.*

In Theorem 3.2, we do not know if one can exactly compute the values of $\nabla\tilde{\mathcal{L}}'(\theta)$, but one can make approximate oracle calls to the gradients of $\tilde{\mathcal{L}}'$ via Monte Carlo sampling. That is, one samples ξ_1, \dots, ξ_k , and then uses $\frac{1}{k} \sum_{i \in [k]} \nabla(\min\{\mathcal{L}(\theta + \xi; D), \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi\})$ as an estimate of $\nabla\tilde{\mathcal{L}}'(\theta)$. It is easy to check that using Monte Carlo sampling instead of exact gradients does not affect our (worst-case) privacy analysis. Of course, in practice, the assumptions in Theorem 3.2 may not hold anyway, and using the convolved loss function and using the unperturbed loss function will lead to similar outcomes for small λ , and also have the same privacy guarantees.

4. New Results for DP-SCO

For the purposes of brevity, in this section we only summarize our new results for ε -DP-ERM/SCO under strong convexity. We defer the other results for DP-ERM/DP-SCO to Appendix E, with a summary in Table 1⁸. Our algorithm is an *iterated exponential mechanism*. Letting $\mathcal{C}_0 = \mathcal{C}$, for $i = 1$ to k , the iterated exponential mechanism samples θ_i from an ε_i -DP exponential mechanism on \mathcal{C}_{i-1} , then sets $\mathcal{C}_i := \{\theta \in \mathcal{C}_{i-1} : \|\theta - \theta_i\|_2 \leq \sqrt{cL(p + 3 \log n) \|\mathcal{C}_{i-1}\|_2 / (m\varepsilon_i n)}\}$. The radii of these sets, as well as our choice $\varepsilon_i = \varepsilon \cdot 2^{i-k-1}$, is chosen such that with high probability, all

8. The bounds are derived via LD, and for (ε, δ) -DP can be improved to bounds tight within constants via direct analysis of privacy via Theorem 2.7. However, then the analysis heavily relies on convexity for privacy.

	Assumption	ε -DP	(ε, δ) -DP
DP-ERM	Convex	$\frac{Lp}{\varepsilon n}$	$\tilde{O}_\delta\left(\frac{L\sqrt{p}}{\varepsilon n}\right)$
	SC	$\frac{L^2(p^2+p\log n)}{m\varepsilon^2n^2}$	$\tilde{O}_\delta\left(\frac{L^2p}{m\varepsilon^2n^2}\right)$
DP-SCO	Convex	$\frac{L}{\sqrt{n}} + \frac{Lp}{\varepsilon n}$	$\frac{L}{\sqrt{n}} + \tilde{O}_\delta\left(\frac{L\sqrt{p}}{\varepsilon n}\right)$
	SC	$\frac{L^2}{mn} + \frac{L^2p^2\log n}{m\varepsilon^2n^2}$	$\frac{L^2}{mn} + \tilde{O}_\delta\left(\frac{L^2p}{m\varepsilon^2n^2}\right)$

Table 1: Summary of results that can be (re-)derived using the Rashomon sampler. The bounds marked in blue were not known even via different algorithms, and all other bounds are tight. Convex: Convex bounded Lipschitz losses, SC: Convex with $\nabla^2\ell(\theta; \cdot) \succcurlyeq m\mathbb{I}$.

\mathcal{C}_i contain the empirical minimizer of \mathcal{L} over \mathcal{C} (see Lemma E.3), but the radii of the sequence are rapidly decreasing. In turn, we can show after $k \approx \log \log(\varepsilon n)$ iterated exponential mechanisms, $\|\mathcal{C}_{k-1}\|_2 = O\left(\frac{L(p+\log n)}{m\varepsilon n}\right)$, and θ_k satisfies our desired DP-ERM bound by Theorem 2.4. Furthermore, we satisfy ε -DP by composition. For DP-SCO we use the fact that for strongly convex losses the empirical and population minimizers are within distance $O\left(\frac{L\log n}{m\sqrt{n}}\right)$ with high probability by concentration of sums of random vectors. In turn, by increasing the radius of each \mathcal{C}_i by $O\left(\frac{L\log n}{m\sqrt{n}}\right)$, the population minimizer remains within the sequence. Then, we can apply Theorem 2.5 and Theorem 2.4 to get the desired DP-SCO bound for the final exponential mechanism.

5. Discussion and Future Directions

In this work we demonstrated the power of Langevin diffusion (LD) by simultaneously obtaining tight guarantees for DP-ERM, DP-SCO, and obtaining the the first private uniform sampling algorithms from Rashomon sets. Furthermore, we demonstrated that, via SGLD, it is possible to maintain the same privacy/utility trade-offs whiling allowing the algorithm to implemented on a finite precision machine. We believe the idea of using a LD to analyze the privacy of the mechanism that samples from the Gibbs distribution, and using a LD to analyze the uniform stability of this mechanism has wider applicability in the DP literature. We leave it for future exploration. Furthermore, the gradient complexity of our SGLD algorithm is inferior to that of Gopi et al. (2022). It is an important open question if it is at all possible to close this gap while not relying on the convexity of the loss function as in our case. This would have a significant real world implications where we aim to ensure privacy and also train learning models that are inherently non-convex. Finally, we believe that exploring other applications of Rashomon sets would would facilitate wider adoption of machine learning models that are differentially private.

Acknowledgments

We would like to thank Walid Krichene, Dvijotham Krishnamurthy, Ryan McKenna, Sewoong Oh, Adam Smith, and Thomas Steinke for their helpful comments. AG was supported in part by NSF CCF-1816861. JU was supported by the Decanal Research (GN: 302918) from Rutgers University.

References

- Martín Abadi, Andy Chu, Ian J. Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proc. of the 2016 ACM SIGSAC Conf. on Computer and Communications Security (CCS'16)*, pages 308–318, 2016.
- Jason M Altschuler and Kunal Talwar. Concentration of the langevin algorithm’s stationary distribution. *arXiv preprint arXiv:2212.12629*, 2022a.
- Jason M Altschuler and Kunal Talwar. Privacy of noisy stochastic gradient descent: More iterations without more privacy loss. *arXiv preprint arXiv:2205.13710*, 2022b.
- Jason M Altschuler and Kunal Talwar. Resolving the mixing time of the langevin algorithm to its stationary distribution for log-concave sampling. *arXiv preprint arXiv:2210.08448*, 2022c.
- Hilal Asi and John C Duchi. Instance-optimality in differential privacy via approximate inverse sensitivity mechanisms. *Advances in neural information processing systems*, 33:14106–14117, 2020a.
- Hilal Asi and John C Duchi. Near instance-optimality in differential privacy. *arXiv preprint arXiv:2005.10630*, 2020b.
- Hilal Asi, Daniel Asher Nathan Levy, and John Duchi. Adapting to function difficulty and growth conditions in private optimization. In *Advances in Neural Information Processing Systems*, 2021.
- Dominique Bakry and Michel Émery. Diffusions hypercontractives. *Séminaire de probabilités de Strasbourg*, 19:177–206, 1985. URL <http://eudml.org/doc/113511>.
- Peter L Bartlett, Olivier Bousquet, and Shahar Mendelson. Local rademacher complexities. 2005.
- Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Proc. of the 2014 IEEE 55th Annual Symp. on Foundations of Computer Science (FOCS)*, pages 464–473, 2014.
- Raef Bassily, Vitaly Feldman, Kunal Talwar, and Abhradeep Thakurta. Private stochastic convex optimization with optimal rates. In *Advances in Neural Information Processing Systems*, pages 11279–11288, 2019.
- Raef Bassily, Vitaly Feldman, Cristóbal Guzmán, and Kunal Talwar. Stability of stochastic gradient descent on nonsmooth convex losses. *arXiv preprint arXiv:2006.06914*, 2020.
- Olivier Bousquet and André Elisseeff. Stability and generalization. *The Journal of Machine Learning Research*, 2:499–526, 2002.
- Thomas Brawner and James Honaker. Bootstrap inference and differential privacy: Standard errors for free. 2018.
- Leo Breiman. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science*, 16(3):199–231, 2001.

- Sébastien Bubeck, Ronen Eldan, and Joseph Lehec. Sampling from a log-concave distribution with projected langevin monte carlo. *Discrete and Computational Geometry*, 59, 06 2018. doi: 10.1007/s00454-018-9992-1.
- Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. *SIAM Journal on Computing*, 47(5):1888–1938, 2018.
- Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(Mar):1069–1109, 2011.
- Xiang Cheng and Peter Bartlett. Convergence of langevin mcmc in kl-divergence. In Firdaus Janoos, Mehryar Mohri, and Karthik Sridharan, editors, *Proceedings of Algorithmic Learning Theory*, volume 83 of *Proceedings of Machine Learning Research*, pages 186–211. PMLR, 07–09 Apr 2018. URL <https://proceedings.mlr.press/v83/cheng18a.html>.
- Xiang Cheng, Niladri S Chatterji, Yasin Abbasi-Yadkori, Peter L Bartlett, and Michael I Jordan. Sharp convergence rates for langevin dynamics in the nonconvex setting. *arXiv preprint arXiv:1805.01648*, 2018.
- Sinho Chewi, Murat A Erdogdu, Mufan Bill Li, Ruoqi Shen, and Matthew Zhang. Analysis of langevin monte carlo from poincaré to log-sobolev. *arXiv preprint arXiv:2112.12662*, 2021.
- Rishav Chourasia, Jiayuan Ye, and Reza Shokri. Differential privacy dynamics of langevin diffusion and noisy gradient descent. In *Advances in Neural Information Processing Systems*, 2021.
- Beau Coker, Cynthia Rudin, and Gary King. A theory of statistical inference for ensuring the robustness of scientific results. *Management Science*, 67(10):6174–6197, 2021.
- Amanda Coston, Ashesh Rambachan, and Alexandra Chouldechova. Characterizing fairness over the set of good models under selective labels. In *International Conference on Machine Learning*, pages 2144–2155. PMLR, 2021.
- Kathleen Creel and Deborah Hellman. The algorithmic leviathan: Arbitrariness, fairness, and opportunity in algorithmic decision-making systems. *Canadian Journal of Philosophy*, pages 1–18, 2022.
- Arnak S. Dalalyan. Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 79(3):651–676, 2017. ISSN 13697412, 14679868. URL <http://www.jstor.org/stable/44681805>.
- Arnak S Dalalyan and Lionel Riou-Durand. On sampling from a log-concave density using kinetic langevin diffusions. *Bernoulli*, 26(3):1956–1988, 2020.
- Alain Durmus and Eric Moulines. Nonasymptotic convergence analysis for the unadjusted langevin algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017.
- Alain Durmus, Szymon Majewski, and Błażej Miasojedow. Analysis of langevin monte carlo via convex optimization. *The Journal of Machine Learning Research*, 20(1):2666–2711, 2019.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014.

Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology–EUROCRYPT*, pages 486–503, 2006a.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proc. of the Third Conf. on Theory of Cryptography (TCC)*, pages 265–284, 2006b. URL http://dx.doi.org/10.1007/11681878_14.

Murat A. Erdogdu, Rasa Hosseinzadeh, and Matthew S. Zhang. Convergence of langevin monte carlo in chi-squared and renyi divergence, 2021.

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

Vitaly Feldman, Tomer Koren, and Kunal Talwar. Private stochastic convex optimization: Optimal rates in linear time. In *Proc. of the Fifty-Second ACM Symp. on Theory of Computing (STOC’20)*, 2020.

Aaron Fisher, Cynthia Rudin, and Francesca Dominici. All models are wrong, but many are useful: Learning a variable’s importance by studying an entire class of prediction models simultaneously. *J. Mach. Learn. Res.*, 20(177):1–81, 2019.

Arun Ganesh and Kunal Talwar. Faster differentially private samplers via rényi divergence analysis of discretized langevin mcmc. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 7222–7233. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/50cf0fe63e0ff857e1c9d01d827267c>

Sivakanth Gopi, Yin Tat Lee, and Daogao Liu. Private convex optimization via exponential mechanism. *arXiv preprint arXiv:2203.00263*, 2022.

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.

Moritz Hardt, Benjamin Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML’16*, page 1225–1234. JMLR.org, 2016.

Thomas P. Hayes. A large-deviation inequality for vector-valued martingales. 2003.

Richard Holley and Daniel W. Stroock. Logarithmic sobolev inequalities and stochastic ising models. *Journal of Statistical Physics*, 46:1159–1194, 1987.

Samuel B Hopkins, Gautam Kamath, Mahbod Majid, and Shyam Narayanan. Robustness implies privacy in statistical estimation. *STOC*, 2023.

- Hsiang Hsu and Flavio P Calmon. Rashomon capacity: Measuring predictive multiplicity in probabilistic classification. In *NeurIPS*, 2022.
- Roger Iyengar, Joseph P Near, Dawn Song, Om Thakkar, Abhradeep Thakurta, and Lun Wang. Towards practical differentially private convex optimization. In *2019 IEEE Symposium on Security and Privacy (SP)*, 2019.
- Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*, 2018.
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.
- Daniel Kifer, Adam Smith, and Abhradeep Thakurta. Private convex empirical risk minimization and high-dimensional regression. In *Conference on Learning Theory*, pages 25–1, 2012.
- Alexander Kolesnikov and Emanuel Milman. Riemannian metrics on convex sets with applications to poincaré and log-sobolev inequalities. *Calculus of Variations and Partial Differential Equations*, 55, 06 2016. doi: 10.1007/s00526-016-1018-3.
- Janardhan Kulkarni, Yin Tat Lee, and Daogao Liu. Private non-smooth erm and sco in subquadratic steps. *Advances in Neural Information Processing Systems*, 34, 2021.
- Michel Ledoux. Concentration of measure and logarithmic sobolev inequalities. In *Seminaire de probabilites XXXIII*, pages 120–216. Springer, 1999.
- Michel Ledoux. Logarithmic sobolev inequalities for unbounded spin systems revisited. *Séminaire de Probabilités XXXV*, pages 167–194, 2001.
- Benjamin Letham, Portia A Letham, Cynthia Rudin, and Edward P Browne. Prediction uncertainty and optimal experimental design for learning dynamical systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 26(6):063110, 2016.
- David Madras, James Atwood, and Alex D’Amour. Detecting underspecification with local ensembles. *arXiv e-prints*, pages arXiv–1910, 2019.
- Oren Mangoubi and Nisheeth K Vishnoi. Sampling from log-concave distributions with infinity-distance guarantees. In *Advances in Neural Information Processing Systems*, 2021.
- Oren Mangoubi and Nisheeth K Vishnoi. Re-analyze gauss: Bounds for private matrix approximation via dyson brownian motion. *arXiv preprint arXiv:2211.06418*, 2022.
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS’07)*, pages 94–103. IEEE, 2007.
- Nicolai Meinshausen and Peter Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473, 2010.

- Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.
- Daniel Nevo and Ya’acov Ritov. Identifying a minimal class of models for high-dimensional data. *The Journal of Machine Learning Research*, 18(1):797–825, 2017.
- Stephan Rabanser, Anvith Thudi, Abhradeep Thakurta, Krishnamurthy Dvijotham, and Nicolas Papernot. Training private models that know what they don’t know. *arXiv preprint arXiv:2305.18393*, 2023.
- Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky. Non-convex learning via stochastic gradient langevin dynamics: a nonasymptotic analysis. In *Proceedings of the 2017 Conference on Learning Theory*, pages 1674–1703, 2017.
- Cynthia Rudin, Chaofan Chen, Zhi Chen, Haiyang Huang, Lesia Semenova, and Chudi Zhong. Interpretable machine learning: Fundamental principles and 10 grand challenges. *Statistics Surveys*, 16:1–85, 2022.
- Théo Ryffel, Francis Bach, and David Pointcheval. Differential privacy guarantees for stochastic gradient langevin dynamics. *arXiv preprint arXiv:2201.11980*, 2022.
- André Schlichting. Poincaré and log–sobolev inequalities for mixtures. *Entropy*, 21(1):89, jan 2019. doi: 10.3390/e21010089. URL <https://doi.org/10.3390/e21010089>.
- Lesia Semenova, Cynthia Rudin, and Ronald Parr. On the existence of simpler machine learning models. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1827–1858, 2022.
- Virat Shejwalkar, Arun Ganesh, Rajiv Mathews, Om Thakkar, and Abhradeep Thakurta. Recycling scraps: Improving private learning by leveraging intermediate checkpoints. *arXiv preprint arXiv:2210.01864*, 2022.
- Adam Smith, Abhradeep Thakurta, and Jalaj Upadhyay. Is interaction necessary for distributed private learning? In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 58–77. IEEE, 2017.
- Helen Smith. Algorithmic bias: should students pay the price? *AI & society*, 35(4):1077–1078, 2020.
- Shuang Song, Kamalika Chaudhuri, and Anand D Sarwate. Stochastic gradient descent with differentially private updates. In *2013 IEEE Global Conference on Signal and Information Processing*, pages 245–248. IEEE, 2013.
- Shuang Song, Om Thakkar, and Abhradeep Thakurta. Characterizing private clipped gradient descent on convex generalized linear problems. *arXiv preprint arXiv:2006.06783*, 2020.
- Nathan Srebro, Karthik Sridharan, and Ambuj Tewari. Smoothness, low noise and fast rates. *Advances in neural information processing systems*, 23, 2010.
- Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *arXiv preprint arXiv:1501.06095*, 2015.

Thomas Steinke and Jonathan Ullman. Tight lower bounds for differentially private selection. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 552–563. IEEE, 2017.

Hiroshi Tanaka. Stochastic differential equations with reflecting boundary condition in convex regions. *Hiroshima Mathematical Journal*, 9(1):163 – 177, 1979. doi: 10.32917/hmj/1206135203. URL <https://doi.org/10.32917/hmj/1206135203>.

Theja Tulabandhula and Cynthia Rudin. Machine learning with operational costs. 2013.

Theja Tulabandhula and Cynthia Rudin. On combining machine learning with decision making. *Machine learning*, 97(1):33–64, 2014a.

Theja Tulabandhula and Cynthia Rudin. Robust optimization using machine learning for uncertainty sets. *arXiv preprint arXiv:1407.1097*, 2014b.

T. van Erven and P. Harremoës. Rényi divergence and Kullback-Leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, July 2014. ISSN 1557-9654. doi: 10.1109/TIT.2014.2320500.

Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted langevin algorithm: Isoperimetry suffices. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/65a99bb7a3115fdede20da98b08a370>

Di Wang, Changyou Chen, and Jinhui Xu. Differentially private empirical risk minimization with non-convex loss functions. In *International Conference on Machine Learning*, pages 6526–6535. PMLR, 2019.

Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688, 2011.

Xi Wu, Fengang Li, Arun Kumar, Kamalika Chaudhuri, Somesh Jha, and Jeffrey F. Naughton. Bolt-on differential privacy for scalable stochastic gradient descent-based analytics. In Semih Salihoglu, Wenchao Zhou, Rada Chirkova, Jun Yang, and Dan Suciu, editors, *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD*, 2017.

Appendix A. Summary of Risk Bounds for Rashomon Set Sampling

In Table 2 we provide a summary of the empirical/population risk guarantees for our Rashomon samplers.

Assumption	ε -DP	(ε, δ) -DP
Excess empirical risk (γ)	$\frac{Lp}{\varepsilon n}$	$\frac{L^2 p \log(1/\delta)}{m \varepsilon^2 n^2}$
Excess population risk	$\frac{Lp}{\varepsilon n} + \psi + L\sqrt{\frac{\psi}{m}} + \frac{L^2}{mn}$	$\frac{L^2 p \log(1/\delta)}{m \varepsilon^2 n^2} + \psi + L\sqrt{\frac{\psi}{m}} + \frac{L^2}{mn}$

Table 2: Summary of our Rashomon sampler guarantees. In all results, λ (the sampling error) is 0. In bounds where the parameters appear, we assume L -Lipschitzness, m -strong convexity, and M -smoothness within \mathcal{C} .

Appendix B. Related Work

Rashomon effect, Rashomon set, and predictive multiplicity: Rashomon set has been extensively studied in applied machine learning since its conception (Coker et al., 2021; Fisher et al., 2019; Meinshausen and Bühlmann, 2010; Letham et al., 2016; Nevo and Ritov, 2017; Semenova et al., 2022; Srebro et al., 2010; Tulabandhula and Rudin, 2014b) (see the survey (Rudin et al., 2022) for more references), culminating in a recent work of Semenova et al. (2022). For example, Fisher et al. (2019) leverages the Rashomon set in order to understand the spectrum of variable importance and other statistics across the set of good models, and Tulabandhula and Rudin (2013, 2014a,b) uses the Rashomon set to assist with decision making. However, it is computationally inefficient to find the simplest model in the Rashomon set, so a natural question is when should we even search for a simpler model. In a recent work, Semenova et al. (2022) showed that, if there is a large Rashomon set of almost-equally-accurate models, a simple model may also be contained in it and that model is guaranteed to generalize well. For this, they defined *Rashomon ratio*, which is the ratio of the volume of the set of accurate models to the volume of the hypothesis space. They then used the insights gained from Rashomon ratio to infer whether a simpler model exists or not.

Rashomon ratio are defined in terms of volumes of the hypothesis (or model parameter) space. In a recent work, Hsu and Calmon (2022) proposed a different metric known as *Rashomon capacity*. It aims to measure the multiplicity of classifier outputs for individual samples, i.e., it measures the spread of the scores with divergence measures for probability distributions (also known as *predictive multiplicity*). In particular, this helps to distinguish Rashomon sets where different predictions are a result of highly different predicted probabilities vs sets where the predictions are different but they come from similar soft outputs.

Some applications of Rashomon sets: Rashomon set has many applications, such as in interpretable and robust machine learning (Fisher et al., 2019; Rudin et al., 2022), understanding spectrum of variable importance (Fisher et al., 2019), decision making (Tulabandhula and Rudin, 2013, 2014a,b), measuring underspecification (Madras et al., 2019), and characterizing fairness (Coston et al., 2021) to name a few. In fact, if we assume that the loss function is smoothness of a loss function, then one can obtain a tighter excess risk bound through local Rademacher complexity (Bartlett et al., 2005). A line of work has also related Rashomon sets with p -hacking and robustness of estimation.

The central argument there is that the Rashomon set is a set on which one might conduct a sensitivity analysis for choices made by an analyst. We refer the interested readers to the survey by [Rudin et al. \(2022\)](#).

Sampling and differential privacy: Several lines of work have designed Markov chains that generate samples from distributions that are close to a given log-concave distribution. For example, there are works that give sampling algorithms with bounds on the distance to the target density in terms of Wasserstein distance ([Dalalyan and Riou-Durand, 2020](#); [Durmus and Moulines, 2017](#)), KL-divergence ([Durmus et al., 2019](#)), and Renyi divergence ([Vempala and Wibisono, 2019](#)). For privacy, we need a bound in terms of ℓ_∞ distance. For this, the first work that performs sampling with bounded ℓ_∞ distance is by [Hardt and Talwar \(2010\)](#). This was extended to bounded Lipschitz log-concave distribution by [Bassily et al. \(2014\)](#). Since then, several works have shown efficient algorithms for sampling from log-concave distribution ([Mangoubi and Vishnoi, 2021](#); [Ganesh and Talwar, 2020](#)). There has been some work that finds polynomial time sampling algorithms for special loss functions. For example, [Asi and Duchi \(2020a,b\)](#) showed efficient algorithms to sample from the exponential mechanism when the score function has a specific structure, which they call *path-length function*. The motivation of [Asi and Duchi \(2020a,b\)](#) was to study instance-optimality of certain wide class of statistical problems. A variant of this function (which is quasi-convex) has been recently used by [Hopkins et al. \(2023\)](#) for robust high-dimensional parameter estimation problems, including mean and covariance estimation, when the data is picked from multivariate Gaussian distribution. Using the insights from differential privacy, a recent line of work ([Altschuler and Talwar, 2022b,c](#)) have improved (and characterized) the mixing time of the discretization of Langevin diffusion.

There has been some recent work to study the asymptotic bias introduced by the discretization of Langevin diffusion. In a recent work, [Altschuler and Talwar \(2022a\)](#) showed that the stationary distribution of discretization of the Langevin diffusion is sub-Gaussian when the potential function is strongly convex, and is sub-exponential when the potential function is convex.

Differential privacy and dynamical systems: The connection between dynamical systems and differential privacy is also not new. [Chourasia et al. \(2021\)](#) and [Ryffel et al. \(2022\)](#) study discretization of the LD algorithm as DP-(Stochastic) Gradient Langevin Dynamics (DP-SGLD). They show that under smoothness and strong convexity on the loss function $\mathcal{L}(\theta; D)$, the privacy cost of DP-SGLD converges to a stationary finite value, even when the number of time steps goes to ∞ . [Wang et al. \(2019\)](#) used the result by [Raginsky et al. \(2017\)](#) to prove a sub-optimal excess empirical risk of $\tilde{O}\left(\frac{p \log(1/\delta)}{\varepsilon^2 \log(n)}\right)$ for non-convex loss functions. In a concurrent, and complementary work on convex losses, [Gopi et al. \(2022\)](#) study private optimization and show the universality of exponential mechanisms for both stochastic convex optimization and empirical risk minimization. Their analysis takes the sampling perspective when the diffusion process has completed.

It is probably important to mention that objective perturbation ([Chaudhuri et al., 2011](#); [Kifer et al., 2012](#)) can be potentially thought of as a (near) universal algorithm for the problem classes considered in this paper, albeit the following two caveats: i) The instantiation of the algorithm for ε -DP and (ε, δ) -DP require two different noise models to be drawn from, namely, Gamma distribution, and Normal distribution, and ii) It requires the loss functions $\ell(\theta; \cdot)$ to be twice-continuously differ-

entiable, and $\nabla_{\theta}^2 \ell(\theta; \cdot)$ to have a near constant rank. As mentioned in the remainder of our paper, Langevin diffusion does not require any such assumptions.⁹

Recently, [Mangoubi and Vishnoi \(2022\)](#) used continuous-time viewpoint to study the error incurred by adding a symmetric Gaussian matrix to input covariance matrix. In particular, they viewed the the perturbed matrix as a continuous-time symmetric matrix diffusion, where each entry of the perturbed matrix is the value reached by a Brownian motion after the time equals to the scaling of variance required for privacy. In particular, the corresponding Brownian motion is well studied in statistical quantum physics and is known as *Dyson brownian motion*.

Appendix C. Notation and Preliminaries

In this section, we give a brief exposition of the concepts and results used in the rest of the paper. In Table 3 we provide a summary of the notation used in the paper.

Rényi divergence and differential privacy. Rényi divergence is the generalization of KL divergence to higher order and satisfies many useful properties ([van Erven and Harremos, 2014](#)). More formally,

Definition C.1 (Rényi Divergence) For $0 < \alpha < \infty$, $\alpha \neq 1$ and distributions P, Q , such that $\text{supp}(P) = \text{supp}(Q)$ the α -Rényi divergence between P and Q is

$$R_{\alpha}(P, Q) = \frac{1}{\alpha - 1} \ln \int_{\text{supp}(Q)} \frac{P(x)^{\alpha}}{Q(x)^{\alpha-1}} dx = \frac{1}{\alpha - 1} \ln \mathbb{E}_{x \sim Q} \left[\frac{P(x)^{\alpha}}{Q(x)^{\alpha}} \right].$$

The α -Rényi divergence for $\alpha = 1$ (resp. ∞) is defined by taking the limit of $R_{\alpha}(P, Q)$ as α approaches 1 (resp. ∞) and equals the KL divergence (resp. max divergence).

We next define differential privacy, our choice of the notion of data privacy. Central to the notion of differential privacy is the definition of *adjacent* or *neighboring* datasets. Two datasets D and D' are called adjacent if they differ in exactly one data point.

Definition C.2 (Approximate Differential privacy ([Dwork et al., 2006b](#))) A randomized mechanism $\mathcal{M} : \mathcal{D}^n \rightarrow \mathcal{R}$ is said to have (ϵ, δ) -differential privacy, or (ϵ, δ) -DP for short, if for any adjacent $D, D' \in \mathcal{D}^n$ and measurable subset $S \subset \mathcal{R}$, it holds that

$$\Pr[\mathcal{M}(D) \in S] \leq e^{\epsilon} \Pr[\mathcal{M}(D') \in S] + \delta.$$

When $\delta = 0$, it is known as pure differential privacy, and we denote it by ϵ -DP.

Definition C.3 (Renyi Differential privacy ([Mironov, 2017](#))) A randomized mechanism $\mathcal{M} : \mathcal{D}^n \rightarrow \mathcal{R}$ is said to have (α, ϵ) -Rényi differential privacy, or (α, ϵ) -RDP for short, if for any adjacent $D, D' \in \mathcal{D}^n$ it holds that

$$R_{\alpha}(\mathcal{M}(D), \mathcal{M}(D')) \leq \epsilon.$$

9. In particular, we can always ensure twice differentiability by convolving the loss function with the bump kernel ([Kifer et al., 2012](#)), and then make the smoothness parameter finite but arbitrarily large which does not affect the Lipschitzness.

Notation	
$D = \{d_1, \dots, d_n\}$	data set
\mathcal{D}	data distribution
τ	domain set of data
$\mathcal{C} \subset \mathbb{R}^p$	convex set/parameter space
ℓ	loss function
\mathcal{L}	empirical loss function
Risk_{ERM}	excess empirical risk
Risk_{SCO}	excess population risk
m	strong convexity parameter
M	smoothness parameter
L	Lipschitz constant
θ^{priv}	private model output
θ^*	optimal model
θ^{emp}	model for ERM
β	inverse temperature
W_t	standard Brownian motion
$R_\alpha(\cdot, \cdot)$	Rényi divergence of order α
T	time
$A \succeq 0$	A is positive semidefinite
$A \succeq B$	$A - B$ is positive semidefinite
\mathbb{I}_p	$p \times p$ identity matrix
(ε, δ)	Privacy parameters
γ	Empirical loss slack
γ'	Population loss slack
ψ	Threshold of Rashomon set
λ	Sampling error wrt TVD

Table 3: Notation Table

It is easy to see that ε -DP is merely (∞, ε) -RDP. Similarly, the following fact relates (ε, δ) -DP to (α, ε) -RDP:

Fact C.4 (Proposition 3 in Mironov (2017)) *If \mathcal{M} satisfies (α, ε) -RDP, then \mathcal{M} is $(\varepsilon + \frac{\log 1/\delta}{\alpha-1}, \delta)$ -differentially private for any $0 < \delta < 1$.*

Rényi divergences satisfy a number of other useful properties, which we list here.

Fact C.5 (Monotonicity (van Erven and Harremoës, 2014, Theorem 3)) *For any distributions P, Q and $0 \leq \alpha_1 \leq \alpha_2$ we have $R_{\alpha_1}(P, Q) \leq R_{\alpha_2}(P, Q)$.*

Fact C.6 (Post-Processing (van Erven and Harremoës, 2014, Theorem 9)) *For any sample spaces \mathcal{X}, \mathcal{Y} , distributions P, Q over \mathcal{X} , and any function $f : \mathcal{X} \rightarrow \mathcal{Y}$ we have $R_\alpha(f(P), f(Q)) \leq R_\alpha(P, Q)$.*

Lemma C.7 (Gaussian dichotomy (van Erven and Harremos, 2014, Example 3)) Let $\mathcal{P} = \mathcal{P}_1 \times \mathcal{P}_2 \times \dots$ and $\mathcal{Q} = \mathcal{Q}_1 \times \mathcal{Q}_2 \times \dots$, where \mathcal{P}_i and \mathcal{Q}_i are unit variance Gaussian distributions with mean μ_i and ν_i , respectively. Then

$$R_\alpha(\mathcal{P}_i, \mathcal{Q}_i) = \frac{\alpha}{2}(\mu_i - \nu_i)^2,$$

and by additivity for $\alpha > 0$,

$$R_\alpha(\mathcal{P}, \mathcal{Q}) = \frac{\alpha}{2} \sum_{i=1}^{\infty} (\mu_i - \nu_i)^2.$$

As a corollary, we have:

$$R_\alpha(N(0, \sigma^2 \mathbb{I}_p), N(\mathbf{x}, \sigma^2 \mathbb{I}_p)) \leq \frac{\alpha \|\mathbf{x}\|_2^2}{2\sigma^2}.$$

Fact C.8 (Adaptive Composition Theorem (Mironov, 2017, Proposition 1)) Let $\mathcal{X}_0, \mathcal{X}_1, \dots, \mathcal{X}_k$ be arbitrary sample spaces. For each $i \in [k]$, let $f_i, f'_i : \Delta(\mathcal{X}_{i-1}) \rightarrow \Delta(\mathcal{X}_i)$ be maps from distributions over \mathcal{X}_{i-1} to distributions over \mathcal{X}_i such that for any distribution X_{i-1} over \mathcal{X}_{i-1} , $R_\alpha(f_i(X_{i-1}), f'_i(X_{i-1})) \leq \varepsilon_i$. Then, for $F, F' : \Delta(\mathcal{X}_0) \rightarrow \Delta(\mathcal{X}_k)$ defined as $F(\cdot) = f_k(f_{k-1}(\dots f_1(\cdot) \dots))$ and $F'(\cdot) = f'_k(f'_{k-1}(\dots f'_1(\cdot) \dots))$ we have $R_\alpha(F(X_0), F'(X_0)) \leq \sum_{i=1}^k \varepsilon_i$ for any $X_0 \in \Delta(\mathcal{X}_0)$.

Fact C.9 (Weak Triangle Inequality (Mironov, 2017, Proposition 11)) For any $\alpha > 1$, $q > 1$ and distributions $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3$ with the same support:

$$R_\alpha(\mathcal{P}_1, \mathcal{P}_3) \leq \frac{\alpha - 1/q}{\alpha - 1} R_{q\alpha}(\mathcal{P}_1, \mathcal{P}_2) + R_{\frac{q\alpha-1}{q-1}}(\mathcal{P}_2, \mathcal{P}_3).$$

We discuss two differentially private mechanisms for optimization in this paper. The first one is the *exponential mechanism*. (McSherry and Talwar, 2007). Given some arbitrary domain \mathfrak{D} and range \mathfrak{R} , the exponential mechanism is defined with respect to some loss function, $\ell : \mathfrak{D} \times \mathfrak{R} \rightarrow \mathbb{R}$.

Definition C.10 (Exponential mechanism (McSherry and Talwar, 2007)) Given a privacy parameter ε , the range \mathfrak{R} and a loss function $\ell : \mathfrak{D} \times \mathfrak{R} \rightarrow \mathbb{R}$, the exponential mechanism samples a single element from \mathfrak{R} based on the probability distribution

$$\pi_D(r) = \frac{e^{-\varepsilon \ell(D, r) / 2\Delta_\ell}}{\sum_{r \in \mathfrak{R}} e^{-\varepsilon \ell(D, r) / 2\Delta_\ell}}$$

where Δ_ℓ is the sensitivity of u , defined as $\Delta_\ell := \max_{D \sim D', r \in \mathfrak{R}} |u(D, r) - u(D', r)|$. If \mathfrak{R} is continuous, we instead sample from the distribution with pdf:

$$p_D(r) = \frac{e^{-\varepsilon \ell(D, r) / 2\Delta_\ell}}{\int_{r \in \mathfrak{R}} e^{-\varepsilon \ell(D, r) / 2\Delta_\ell} dr}.$$

Theorem C.11 Assume each of the individual loss function in $\mathcal{L}(\theta; D)$ is L -Lipschitz within the constraint set \mathcal{C} , individual loss function $\ell(\theta; \cdot)$ is convex, and the constraint set \mathcal{C} is convex. Then,

Algorithm 1 Exponential mechanism

Input: Loss function \mathcal{L} , constraint set \mathcal{C} , Lipschitz constant L , number of iterations k , privacy parameter ε , data set D of n -samples.

- 1: Sample and **output** a point θ^{priv} from the constraint set \mathcal{C} w.p. $\propto \exp\left(-\frac{\varepsilon n}{2L\|\mathcal{C}\|_2} \cdot \mathcal{L}(\theta; D)\right)$.

Algorithm 1 is ε -differentially private. Furthermore, for θ^{priv} as specified in Algorithm 1, over the randomness of the algorithm,

$$\mathbb{E}_{\theta^{\text{priv}}} [\text{Risk}_{\text{ERM}}(\theta^{\text{priv}})] = O\left(\frac{Lp \cdot \|\mathcal{C}\|_2}{\varepsilon n}\right).$$

Equivalence of Algorithm 1 and Langevin diffusion: The following lemma, which is implied by, e.g. (Tanaka, 1979, Theorem 4.1), shows that one can implement Algorithm 1 using only solutions to eq. (2); note that this does not necessarily mean solutions to eq. (2) are efficiently sampleable.

Lemma C.12 *Let \mathcal{L} be a M -smooth function for some finite M . Then if $\beta_t = \beta$ for all t , then the stationary distribution of (2) has pdf proportional to $\exp(-\beta\mathcal{L}(\theta; D)) \cdot \mathbb{1}(\theta \in \mathcal{C})$.*

We recall that one can ensure smoothness by convolving \mathcal{L} (appropriately extended to all of \mathbb{R}^p) with the Gaussian kernel of finite variance (Feldman et al., 2018, Appendix C). In particular, since we only need M to be finite, we can take the convolution with the Gaussian kernel $\mathcal{N}(\mathbf{0}, \lambda^2 \mathbb{I}_p)$ for arbitrarily small $\lambda > 0$, and in turn the result of the convolution is L/λ -smooth (which is perhaps arbitrarily large but still finite) and differs from \mathcal{L} by an arbitrarily small amount everywhere in \mathcal{C} .

We use the result by Steinke and Ullman (2017) for our lower bound proof. We use their equivalent result for empirical mean (see equation (2) in Steinke and Ullman (2017)) and for privacy parameters (ε, δ) using a standard reduction (Bun et al., 2018; Steinke and Ullman, 2015)¹⁰:

Theorem C.13 *Fix $n, s, k \in \mathbb{N}$. Set $\beta = 1 + \frac{1}{2} \log\left(\frac{s}{8 \max\{2k, 28\}}\right)$. Let $P^1, \dots, P^s \sim \text{Beta}(\beta, \beta)$ and let $X := \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ be such that $\mathbf{x}_i \in \{0, 1\}^s$ for all $i \in [n]$, $\mathbf{x}_{i,j}$ is independent (conditioned on P) and $\mathbb{E}[\mathbf{x}_{i,j}] = P^j$ for all $i \in [n]$ and $j \in [s]$. Let $\mathcal{M} : (\{0, 1\}^s)^n \rightarrow \{0, 1\}^d$ be $(1, \frac{1}{ns})$ -differentially private. Suppose $\|\mathcal{M}(x)\|_1 = \|\mathcal{M}(x)\|_2^2 = k$ for all X with probability 1 and*

$$\mathbb{E}_{\mathcal{M}} \left[\frac{1}{n} \sum_{i=1}^n \sum_{\substack{j \in [s] \\ \mathcal{M}(x)^j = 1}} \mathbf{x}_{i,j} \right] \geq \frac{1}{n} \max_{\substack{S \subseteq [d] \\ |S|=k}} \sum_{i=1}^n \sum_{u \in S} \mathbf{x}_{i,u} - \frac{k}{20}. \quad (6)$$

Then $n \in \Omega\left(\sqrt{k} \log\left(\frac{s}{k}\right)\right)$.

Results from statistics and machine learning: We will sometimes use Fatou's lemma in our proofs. The form we will use is stated here for convenience:

¹⁰. Steinke and Ullman (2017) present their result in the terms of population mean and privacy parameters $(1, \frac{1}{ns})$.

Lemma C.14 (Fatou’s Lemma) *Let $\{X_i\}$ be a sequence of random variables such that there is some constant c such that for all i , $\Pr[X_i \geq c] = 1$. Then:*

$$\mathbb{E} \left[\liminf_{i \rightarrow \infty} X_i \right] \leq \liminf_{i \rightarrow \infty} \mathbb{E} [X_i].$$

For our SCO bounds, we will use uniform stability. Uniform stability of a learning algorithm is a notion of algorithmic stability introduced to derive high-probability bounds on the generalization error. Formally, it is defined as follows:

Definition C.15 (Uniform stability (Bousquet and Elisseeff, 2002)) *A mechanism \mathcal{M} is $\mu(n)$ -uniformly stable with respect to ℓ if for any pair of databases D, D' of size n differing in at most one individual:*

$$\sup_{d \in \tau} [\mathbb{E}_{\mathcal{M}} [\ell(\mathcal{M}(D), d)] - \mathbb{E}_{\mathcal{M}} [\ell(\mathcal{M}(D'), d)]] \leq \mu(n).$$

In this paper, we will need the following result.

Lemma C.16 (Bousquet and Elisseeff (2002)) *Suppose \mathcal{M} is $\mu(n)$ -uniformly stable. Then:*

$$\mathbb{E}_{D \sim \mathcal{D}^n, \mathcal{M}} [\text{Risk}_{\text{SCO}}(\mathcal{M}(D))] \leq \mathbb{E}_{D \sim \mathcal{D}^n, \mathcal{M}} [\text{Risk}_{\text{ERM}}(\mathcal{M}(D))] + \mu(n).$$

We also use KL divergence and TVD, and the relation between the two.

Definition C.17 *The KL divergence (equivalent to the 1-Rényi divergence) between two distributions P, Q is given by $R_1(P, Q) := \mathbb{E}_{x \sim P} \log \left(\frac{P(x)}{Q(x)} \right)$.*

Lemma C.18 (Pinsker’s inequality) *Let $\text{TVD}(P, Q)$ be the total variation distance between P and Q . Then*

$$\text{TVD}(P, Q) \leq \sqrt{\frac{1}{2} R_1(P, Q)}.$$

Next, we define the LSI constant of a distribution.

Definition C.19 *A distribution P satisfies LSI with constant c if for all smooth functions $g : \mathbb{R}^p \rightarrow \mathbb{R}$ with $\mathbb{E}_{x \sim P}[g(x)^2] < \infty$:*

$$\mathbb{E}_{x \sim P}[g(x)^2 \log(g(x)^2)] - \mathbb{E}_{x \sim P}[g(x)^2] \cdot \mathbb{E}_{x \sim P}[\log(g(x)^2)] \leq \frac{2}{c} \mathbb{E}_{x \sim P}[\|\nabla g(x)\|_2^2].$$

Appendix D. Proofs from Section 2 and Section 3

D.1. Proof of Lemma 2.2

Proof For ease of presentation, we will show a divergence bound between Θ_t, Θ'_t which are the distributions of θ_t, θ'_t , and then describe how to modify the proof to show the same bound between $\Theta_{[0,t]}, \Theta'_{[0,t]}$.

Let $\Psi_{D,m,i}$ be a map from (distributions over) \mathbb{R}^p to (distributions over) \mathbb{R}^p that takes the point θ to the distribution $\Pi_{\mathcal{C}} \left(N \left(\theta - \left(\frac{\beta t}{m} dt \right) \nabla \mathcal{L}(\theta; D), 2 \frac{t}{m} \mathbb{I} \right) \right)$, where $\Pi_{\mathcal{C}}$ is the ℓ_2 -projection into \mathcal{C} . It is well known (see e.g. Lemma C.7) that:

$$R_\alpha(N(0, \sigma^2 \mathbb{I}), N(\mathbf{x}, \sigma^2 \mathbb{I})) \leq \frac{\alpha \|\mathbf{x}\|_2^2}{2\sigma^2}.$$

So by post-processing (Fact C.6) and the Lipschitzness assumption, $R_\alpha(\Psi_{D,m,i}(\theta), \Psi_{D',m,i}(\theta))$ is bounded by

$$\begin{aligned} & R_\alpha \left(N \left(\theta - \left(\frac{\beta t}{m} \right) \nabla \mathcal{L}(\theta; D), \frac{2t}{m} \mathbb{I} \right), N \left(\theta - \left(\frac{\beta t}{m} \right) \nabla \mathcal{L}(\theta; D'), \frac{2t}{m} \mathbb{I} \right) \right) \\ &= R_\alpha \left(N \left(\mathbf{0}, \frac{2t}{m} \mathbb{I} \right), N \left(\left(\frac{\beta t}{m} \right) (\nabla \mathcal{L}(\theta; D) - \nabla \mathcal{L}(\theta; D')), 2 \frac{t}{m} \mathbb{I} \right) \right) \\ &\leq \frac{\alpha \Delta^2}{4} \cdot \frac{\left(\frac{\beta t}{m} \right)^2}{t/m}. \end{aligned}$$

Let $\Psi_{D,m}$ denote the composition $\Psi_{D,m,m} \circ \Psi_{D,m,m-1} \circ \dots \circ \Psi_{D,m,1}$. By Lemma C.8, we have

$$R_\alpha(\Psi_{D,m}(\Theta_0), \Psi_{D',m}(\Theta_0)) \leq \sum_{i=1}^m \max_{\theta} \{ R_\alpha(\Psi_{D,m,i}(\theta), \Psi_{D',m,i}(\theta)) \}.$$

Plugging in the bound on $R_\alpha(\Psi_{D,m,i}(\theta), \Psi_{D',m,i}(\theta))$, we get

$$R_\alpha(\Psi_{D,m}(\Theta_0), \Psi_{D',m}(\Theta_0)) \leq \frac{\alpha \Delta^2}{4} \cdot \frac{m}{t} \sum_{i=1}^m \left(\frac{\beta t}{m} \right)^2 = \frac{\alpha \beta^2 \Delta^2 t}{4}$$

Note that $\Theta_t = \lim_{m \rightarrow \infty} \Psi_{D,m}(\Theta_0)$, and $\Theta'_t = \lim_{m \rightarrow \infty} \Psi_{D',m}(\Theta_0)$. Since $\exp((\alpha-1)R_\alpha(\mathcal{P}, \mathcal{Q}))$ is a monotone function of $R_\alpha(\mathcal{P}, \mathcal{Q})$ and is the expectation of a positive random variable, by Fatou's lemma we have:

$$\begin{aligned} R_\alpha(\Theta_t, \Theta'_t) &\leq \lim_{m \rightarrow \infty} R_\alpha(\Psi_{D,m}(\Theta_0), \Psi_{D',m}(\Theta_0)) \\ &\leq \frac{\alpha \beta^2 \Delta^2 t}{4}. \end{aligned}$$

This gives the bound on $R_\alpha(\Theta_t, \Theta'_t)$. To obtain the same bound for $R_\alpha(\Theta_{[0,t]}, \Theta'_{[0,t]})$, we modify $\Psi_{D,m,i}$ so that instead of receiving $\Theta_{(i-1)t/m}$ and outputting $\Theta_{it/m}$, it receives the joint distribution $\{\Theta_{jt/m}\}_{0 \leq j \leq i-1}$ and outputs $\{\Theta_{jt/m}\}_{0 \leq j \leq i}$ by appending the (also jointly distributed) variable

$$\Theta_{it/m} = \Pi_{\mathcal{C}} \left(N \left(\theta - \left(\frac{\beta t}{m} \right) \nabla \mathcal{L}(\Theta_{(i-1)t/m}; D), 2 \frac{t}{m} \mathbb{I} \right) \right)$$

That is, we update $\Psi_{D,m,i}$ so it outputs the distributions of all iterates seen so far instead of just the distribution of the last iterate; the limiting value of the joint distribution $\{\Theta_{jt/m}\}_{0 \leq j \leq i}$ is then $\Theta_{[0,t]}$ according to eq. (2), and the same divergence bound holds. \blacksquare

D.2. Proof of Theorem 2.3

We first need the following results, which let us analyze the LSI constant of the Gibbs distribution of interest easily. These results were originally stated for unconstrained domains defined over the space of real numbers, a proof for which can also be found in [Ledoux \(2001\)](#) using the theory of semigroup (see Corollary 1.4, 1.6 and Lemma 1.2). For the general convex set, \mathcal{C} , we refer the readers to Theorem 2.1 in [Kolesnikov and Milman \(2016\)](#).

Lemma D.1 (Proposition 3 and Corollaire 2 in [Bakry and Émery \(1985\)](#)) *Let P be the distribution with density proportional to $\exp(-f(x)) \cdot \mathbb{1}(x \in \mathcal{C})$ for convex \mathcal{C} . If f is m -strongly convex, then P satisfies LSI with constant m .*

Lemma D.2 (Page 1184 in [Holley and Stroock \(1987\)](#)) *Let f, f' be two functions such that*

$$\sup_{\theta \in \mathcal{C}} |f(\theta) - f'(\theta)| \leq \Delta.$$

Suppose the distribution with density proportional to $\mathbb{1}(\theta \in \mathcal{C}) \cdot \exp(-f)$ satisfies LSI with constant c . Then the distribution with density proportional to $\mathbb{1}(\theta \in \mathcal{C}) \cdot \exp(-f')$ satisfies LSI with constant $c \cdot \exp(-\Delta)$.

The following result shows convergence of (2) under LSI.

Lemma D.3 (Theorem 4 of [Vempala and Wibisono \(2019\)](#)) *Suppose Q satisfies LSI with constant c . Then let P_0 be any initial distribution over θ_0 , and P_t be the distribution over θ_t given by running (1) on $-\log q$, where q is the density of Q , for $\beta = 1$. Then $R_1(P_t, Q) \leq \exp(-2ct) \cdot R_1(P_0, Q)$.*

Using these results, we first prove an LSI constant for the Gibbs distribution.

Lemma D.4 *Suppose we have non-negative $\ell(\theta; d)$ such that ℓ is m -strongly convex wrt θ for all d , and let Q be the distribution with density proportional to*

$$\mathbb{1}(\theta \in \mathcal{C}) \cdot \exp \left(-\beta \max \left\{ \psi + \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D), \mathcal{L}(\theta; D) \right\} \right).$$

Then Q satisfies LSI with constant c , where

$$c := \begin{cases} \frac{m}{e\psi} & \text{for } \beta\psi > 1 \\ \beta m \exp(-\beta\psi) & \text{for } \beta\psi \leq 1 \end{cases}.$$

Proof For $a \in [0, 1]$, consider the distribution with density proportional to $e^{-f(\theta)} \cdot \mathbb{1}(\theta \in \mathcal{C})$ for

$$f(\theta) := \beta \left(a \cdot \mathcal{L}(\theta; D) + (1 - a) \cdot \max \left\{ \psi + \min_{\theta \in \mathcal{C}} \mathcal{L}(\theta; D), \mathcal{L}(\theta; D) \right\} \right). \quad (7)$$

Since \mathcal{L} is m -strongly convex, the function $\max\{\psi + \min_{\theta \in \mathcal{C}} \mathcal{L}(\theta; D), \mathcal{L}(\theta; D)\}$ is a convex function. Therefore, $f(\theta)$ defined in eq. (7) is an $(a\beta m)$ -strongly convex function. In other words, this distribution satisfies LSI with constant $a\beta m$ using Lemma D.1.

Now the distribution Q has density proportional to $\exp(-f)$, where

$$f' := \max\{\psi + \min_{\theta \in \mathcal{C}} \mathcal{L}(\theta; D), \mathcal{L}(\theta; D)\}.$$

When $\mathcal{L}(\theta; D) > \psi$, $f = f'$. Otherwise, since ℓ is non-negative, we have that f differs from f' by at most $a\beta\psi$ everywhere. Then by lemma D.2, Q satisfies LSI with constant $a\beta m \exp(-a\beta\psi)$.

Now, we maximize this bound over $a \in [0, 1]$. If $\beta\psi \leq 1$, then this bound is maximized at $a = 1$ and we get an LSI constant of $\beta m \exp(-\beta\psi)$. Otherwise, this bound is maximized at $a = 1/\beta\psi$ and we get an LSI constant of $\frac{m}{e\psi}$. This completes the proof of Lemma D.4. \blacksquare

Given the LSI constant of Rashomon sampler, we can use Lemma D.3 to upper bound the time we need to run DP-LD.

Lemma D.5 *Under the assumptions of Theorem 2.3, let*

$$\mathcal{L}'(\theta; D) := \max \left\{ \psi + \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D), \mathcal{L}(\theta; D) \right\}.$$

Let Q be the distribution with density proportional to $\exp(-\beta\mathcal{L}')$. Let P_0 be the distribution over θ_0 that is uniform over \mathcal{C} . Let P_t be the resulting distribution over θ_t given by running (2) on \mathcal{L}' . Then assuming $\psi \leq L \|\mathcal{C}\|_2$:

- If $\beta\psi > 1$, for $t = \frac{e\psi}{2m} \log(\frac{2\beta L \|\mathcal{C}\|_2}{\delta^2})$, we have $\text{TVD}(P_t, Q) \leq \delta/2$.
- If $\beta\psi \leq 1$, for $t = \frac{e}{2\beta m} \log(\frac{2\beta L \|\mathcal{C}\|_2}{\delta^2})$, we have $\text{TVD}(P_t, Q) \leq \delta/2$.

Proof We wish to apply Lemma D.3, which was originally stated in the unconstrained case for (1). One can see that it applies to running (2) on \mathcal{L}' by the following argument: Consider extending \mathcal{L}' to \mathbb{R}^p by defining $\mathcal{L}'(\theta') = \mathcal{L}'(\Pi_{\mathcal{C}}(\theta')) + c \|\theta' - \Pi_{\mathcal{C}}(\theta')\|_2$ for $\theta' \notin \mathcal{C}$. As c goes to infinity, the Gibbs distribution induced by the extended loss in the unconstrained setting approaches the Gibbs distribution induced by the loss in the constrained setting, and (1) approaches (2).

We have:

$$\begin{aligned} R_1(P_0, Q) &\leq \max_{\theta \in \text{supp}(P_0)} \log(P_0(\theta)/Q(\theta)) \\ &\leq \log \left(\frac{1}{\int_{\theta \in \mathcal{C}} 1 d\theta} \cdot \frac{\int_{\theta \in \mathcal{C}} \exp(-\beta \max\{\mathcal{L}(\theta^*; D) + \psi, \mathcal{L}(\theta; D)\}) d\theta}{\min_{\theta \in \mathcal{C}} \exp(-\beta \max\{\mathcal{L}(\theta^*; D) + \psi, \mathcal{L}(\theta; D)\})} \right) \\ &\leq \log \left(\frac{\exp(-\beta(\mathcal{L}(\theta^*; D) + \psi))}{\exp(-\beta(L \|\mathcal{C}\|_2 + \psi))} \right) \\ &\leq 2\beta L \|\mathcal{C}\|_2. \end{aligned} \quad (8)$$

Here, we use the fact that each ℓ has minimum 0. Recall that

$$\mathcal{L}'(\theta; D) := \max \left\{ \psi + \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D), \mathcal{L}(\theta; D) \right\}.$$

Now running (2) on $\beta \mathcal{L}'$, replacing β in (2) with 1, is equivalent to running (2) on \mathcal{L}' . Therefore, applying Lemma D.3 and Lemma D.4:

- If $\beta\psi > 1$, for $t = \frac{e\psi}{2m} \log \left(\frac{4\beta L \|\mathcal{C}\|_2}{\delta^2} \right)$, we have:

$$R_1(P_t, Q) \leq R_1(P_0, Q) \cdot \exp \left(-\frac{2m}{e\psi} t \right) \leq 2\beta L \|\mathcal{C}\|_2 \cdot \exp \left(-\frac{2m}{e\psi} t \right) = \delta^2/2,$$

where the first inequality is using Lemma D.3 and Lemma D.4 and the second inequality is due to eq. (8). The first bullet now follows using Pinsker's inequality (Lemma C.18).

- If $\beta\psi \leq 1$, note that $\beta m \exp(-\beta\psi) \geq \beta m/e$. So for $t = \frac{e}{2\beta m} \log \left(\frac{4\beta L \|\mathcal{C}\|_2}{\delta^2} \right)$, we have:

$$R_1(P_t, Q) \leq R_1(P_0, Q) \cdot \exp \left(-\frac{2m}{e\psi} t \right) \leq 2\beta L \|\mathcal{C}\|_2 \cdot \exp \left(-\frac{2\beta m}{e} t \right) = \delta^2/2,$$

where the first inequality is using Lemma D.3 and Lemma D.4 and the second inequality is due to eq. (8). The second bullet now follows using Pinsker's inequality (Lemma C.18).

This completes the proof of Lemma D.5. ■

In Lemma D.5, we used a uniform distribution on \mathcal{C} as our initialization (i.e., P_0 was uniform distribution over \mathcal{C}). In the case where \mathcal{C} is *unconstrained* and the losses satisfy $\|\nabla \ell(\theta; d) - \nabla \ell(\theta; d')\|_2 \leq L$ instead of Lipschitzness within \mathcal{C} , we can obtain the same bound by letting \mathcal{C} be the convex hull of the minimizers of all per-example loss functions $\ell(\theta; d)$ and using a uniform distribution on \mathcal{C} as our initial distribution.

We now complete the proof of Theorem 2.3.

Proof Let $\mathbf{v}_D(\theta) = \partial_\theta \max \{ \mathcal{L}(\theta; D), \mathcal{L}(\theta^*; D) + \psi \}$ be the sub-differential of the score function used in the Gibbs distribution. Using the Lipschitz property and the smoothness assumption on $\ell(\theta; d)$, for any two neighboring data sets D and D' , we have

$$\|\mathbf{v}_D(\theta) - \mathbf{v}_{D'}(\theta)\|_2 \leq 3 \cdot \max \left\{ \frac{L}{n}, \sqrt{\frac{M\psi}{2}} \right\}. \quad (9)$$

In particular, the Lipschitzness bounds the sensitivity between the gradients of $\mathcal{L}(\theta; D)$ and $\mathcal{L}(\theta; D')$ by L/n . Using smoothness, for D , the sensitivity between the gradients of $\mathcal{L}(\theta; D)$ and $\max \left\{ \psi + \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D), \mathcal{L}(\theta; D) \right\}$ is bounded by $\sqrt{M\psi/2}$. Similarly, we have the sensitivity between the gradients of $\mathcal{L}(\theta; D')$ and $\max \left\{ \psi + \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D'), \mathcal{L}(\theta; D') \right\}$ bounded by $\sqrt{M\psi/2}$.

By Lemma 2.2, Lemma C.4, and sensitivity bound in eq. (9), in order for $\{\Theta_{t'}\}_{t' \in [0, t]}$ to satisfy (ε, δ) -DP, it suffices if

$$\beta \leq \frac{\varepsilon}{6\sqrt{2} \cdot \max \left\{ \frac{L}{n}, \sqrt{\frac{M\psi}{2}} \right\} \sqrt{t \log(1/\delta)}}. \quad (10)$$

Suppose $\beta\psi \leq 1$. Recall that, in Lemma D.5, $t = \frac{e}{2\beta m} \log(\frac{4\beta L\|\mathcal{C}\|_2}{\delta^2})$. Plugging in this value of t in eq. (10) and observing that $\psi \leq \frac{L\|\mathcal{C}\|_2}{2}$, it suffices to ensure (ε, δ) -DP if

$$\beta = \tilde{\Theta} \left(\frac{\varepsilon^2 n^2 m}{\max \left\{ L, n\sqrt{M\psi/2} \right\}^2 \log(L\|\mathcal{C}\|_2/\delta^2) \log(1/\delta)} \right). \quad (11)$$

We now consider two cases (on top of the constraint that $\psi \leq L\|\mathcal{C}\|_2/2$):

- (i) $\psi > \frac{2L^2}{Mn^2}$, and
- (ii) $\psi \leq \frac{2L^2}{Mn^2}$.

In the case (i), in order to satisfy $\beta\psi \leq 1$ it is sufficient to set $\beta = \tilde{\Theta} \left(\frac{\varepsilon^2(m/M)}{\log((L\|\mathcal{C}\|_2 - \psi)/\delta^2) \log(1/\delta)} \cdot \frac{1}{\psi} \right)$. In the second, case it sufficient to set $\beta = \tilde{\Theta} \left(\frac{\varepsilon^2 n^2 m}{L^2 \log((L\|\mathcal{C}\|_2 - \psi)/\delta^2) \log(1/\delta)} \right)$.

We will not analyze the setting when $\beta\psi > 1$. From Lemma D.5, it is not hard to observe that it will not provide any better conditions on ψ and β to ensure (ε, δ) -DP.

So, we have that for the choices of ψ, β, t in Theorem 2.3, outputting $\{\Theta_{t'}\}_{t' \in [0, t]}$ satisfies (ε, δ) -DP. Furthermore, by Lemma D.5, this gives the desired bound on total variation distance completing the proof of Theorem 2.3. \blacksquare

D.3. Probability of Hitting the Rashomon Set

Theorem D.6 *Let θ^{priv} be the model output by a Rashomon sampler. Let $\Pr[\mathcal{L}(\theta^{\text{priv}}; D) \leq \min_{\theta \in \mathcal{C}} \mathcal{L}(\theta; D) + \psi + \gamma] \geq 1 - \phi$. Then for Rashomon set \mathcal{G} , $\Pr[\theta^{\text{priv}} \in \mathcal{G}] \geq 1 - \phi - \frac{p\gamma}{\psi}$.*

Proof Consider the differential conic region Ω centered at the true minimizer $\theta^* = \arg \min_{\theta \in \mathcal{C}} \mathcal{L}(\theta; D)$. Let A be the cone from this region with height r , and B be the one with height $(r + \Delta)$, with r and Δ chosen such that θ in the boundaries of the cones A and B satisfy $\mathcal{L}(\theta; D) = \psi + \min_{\theta} \mathcal{L}(\theta; D)$ and $\mathcal{L}(\theta; D) = \psi + \gamma + \min_{\theta} \mathcal{L}(\theta; D)$ respectively. By law of total probability, it suffices to show that conditioned on sampling θ in any such cone, the desired probability bound holds. By the maximality condition on the Rashomon sampler,

$$\Pr[\theta^{\text{priv}} \in \mathcal{G} | \theta \in \Omega] \geq 1 - \phi - \frac{\text{Vol}(A)}{\text{Vol}(B)} = 1 - \phi - \left(1 - \frac{\Delta}{r + \Delta}\right)^p \geq 1 - \phi - \frac{p\Delta}{r}.$$

Note that within Ω , we have $\mathcal{L}(\theta; D) = f(\|\theta - \theta^*\|_2)$ where f is a convex function. By convexity, we have:

$$f(r) - f(0) \leq f'(r) \cdot r \rightarrow f'(r) \geq \frac{f(r) - f(0)}{r} = \psi/r.$$

$$\gamma = f(r + \Delta) - f(r) \geq f'(r) \cdot \Delta \geq \frac{\psi \Delta}{r} \implies \Delta/r \leq \gamma/\psi.$$

This completes the proof of Theorem D.6. ■

D.4. Proof of Theorem 3.2

We start by showing some properties of the convolved loss function:

Theorem D.7 *Let \mathcal{L} be a convex and M -smooth loss, let $\mathcal{L}'(\theta) = \max\{\mathcal{L}(\theta), \mathcal{L}(\theta^*) + \psi\}$ and let $\tilde{\mathcal{L}}'(\theta) := \mathbb{E}_{\xi \sim N(0, \lambda^2 \mathbb{I}_p)} [\mathcal{L}'(\theta + \xi)]$. Then:*

1. $\max_{\theta} \left| \tilde{\mathcal{L}}'(\theta) - \mathcal{L}'(\theta) \right| \leq 2\psi + Mp\lambda^2$.
2. $\tilde{\mathcal{L}}'(\theta)$ is $M + \frac{\sqrt{\psi M/2}}{\lambda}$ -smooth.

Proof Let $\tilde{\mathcal{L}}(\theta) := \mathbb{E}_{\xi \sim N(0, \lambda^2 \mathbb{I}_p)} [\mathcal{L}(\theta + \xi)]$. Then:

$$\max_{\theta} \left| \tilde{\mathcal{L}}'(\theta) - \mathcal{L}'(\theta) \right| \leq 2\psi + \max_{\theta} \left| \tilde{\mathcal{L}}(\theta) - \mathcal{L}(\theta) \right|.$$

We bound the second term for all θ simultaneously as follows:

$$\begin{aligned} \left| \tilde{\mathcal{L}}(\theta) - \mathcal{L}(\theta) \right| &= \left| \mathbb{E}_{\xi \sim N(0, \lambda^2 \mathbb{I}_p)} [\mathcal{L}(\theta + \xi) - \mathcal{L}(\theta)] \right| \\ &\stackrel{(*)}{=} \left| \mathbb{E}_{\xi \sim N(0, \lambda^2 \mathbb{I}_p)} [\mathcal{L}(\theta + \xi) - \mathcal{L}(\theta) - \langle \nabla \mathcal{L}(\theta), \xi \rangle] \right| \\ &\stackrel{(**)}{\leq} \mathbb{E}_{\xi \sim N(0, \lambda^2 \mathbb{I}_p)} \left[M \|\xi\|_2^2 \right] = Mp\lambda^2. \end{aligned}$$

In equality (*), we use the fact that ξ is mean 0, and in inequality (**), we use the fact that $0 \leq \mathcal{L}(\theta + \xi) - \mathcal{L}(\theta) - \langle \nabla \mathcal{L}(\theta), \xi \rangle \leq M \|\xi\|_2^2$ by convexity and M -smoothness of \mathcal{L} . This gives the first part of the theorem.

For the second part, fix any θ_1, θ_2 . The smoothness parameter of $\tilde{\mathcal{L}}'$ is at most the smoothness parameter of $\tilde{\mathcal{L}}$ (which is at most M , since \mathcal{L} is M -smooth), plus the smoothness of $\tilde{\mathcal{L}}' - \tilde{\mathcal{L}}$.

Now $\tilde{\mathcal{L}}' - \tilde{\mathcal{L}}$ is the convolution of the Gaussian kernel with the function $\mathcal{L}' - \mathcal{L}$, which is 0 outside of Rashomon set, \mathcal{G} and equal to $\mathcal{L}(\theta^*) + \psi - \mathcal{L}(\theta)$ inside the Rashomon set, \mathcal{G} . By M -smoothness of \mathcal{L} , $\mathcal{L}' - \mathcal{L}$ is globally $\sqrt{\frac{\psi M}{2}}$ -Lipschitz. Then, by Theorem 33 of Feldman et al. (2018), $\tilde{\mathcal{L}}' - \tilde{\mathcal{L}}$ is $\frac{1}{\lambda} \sqrt{\frac{\psi M}{2}}$ -smooth, giving the second part of the theorem.

This completes the proof of Theorem D.7. ■

Similarly to Lemma D.4 we have the following:

Lemma D.8 Let $\tilde{\mathcal{L}}'$ be defined as in Theorem D.7, for some $\lambda \geq 0$. Then the distribution with density proportional to $\exp(-\beta \cdot \tilde{\mathcal{L}}')$ satisfies LSI with constant c for

$$c := \beta m \exp(-\beta(3\psi + Mp\lambda^2)).$$

Recall the SGLD equation:

$$\theta_{t+1} = \theta_t - \eta\beta\nabla\mathcal{L}(\theta_t; D) + \xi_t, \quad \xi_t \sim \mathcal{N}(0, 2\eta I_p). \quad (12)$$

We can now apply the following result of Chewi et al. (2021):

Theorem D.9 (Theorem 4 of Chewi et al. (2021)) For any T and η , let Θ_T be the distribution of $\theta_{T\eta}$ given by the solution to (1), and let Θ'_T be the distribution of θ_T given by (12). Suppose \mathcal{L} is M -smooth and satisfies LSI (Definition C.19) with constant c . Then for any $\kappa > 0, \alpha \geq 2$,

$$\eta = O\left(\frac{c\kappa}{p\alpha M^2} \cdot \min\left\{\frac{1}{\log(\alpha)}, \frac{p}{\alpha\kappa}\right\}\right), \quad \text{and} \quad T = \Omega\left(\frac{p\alpha^2 M^2}{c^2 \kappa} r \cdot \max\left\{\log(\alpha), \frac{\alpha\kappa}{p}\right\}\right),$$

we have $R_\alpha(\Theta'_T, \Theta_\infty) \leq \kappa$.

We can now prove a general result in terms of Rényi divergence:

Theorem D.10 Suppose $\|\nabla\mathcal{L}(\theta; D) - \nabla\mathcal{L}(\theta; D')\|_2 \leq \Delta$, and that each individual loss function ℓ is m -strongly convex and M -smooth. Let $0 \leq \lambda \leq \sqrt{\frac{\psi}{Mp}}$ and let

$$\tilde{\mathcal{L}} := \mathbb{E}_{\xi \sim N(0, \lambda^2 I_p)} \left[\min\{\mathcal{L}(\theta + \xi; D), \min_{\theta^* \in \mathcal{C}} \mathcal{L}(\theta^*; D) + \psi\} \right].$$

Let Q be the (unconstrained) Gibbs distribution of $\beta \cdot \tilde{\mathcal{L}}$. Let Θ'_T be the solution to (12) starting from the distribution Θ'_0 , run on the loss $\tilde{\mathcal{L}}$. Fix any $r > 0$. Then for

$$\beta = O\left(\frac{\varepsilon^2 m}{\max\{\Delta^2, M\psi\} \alpha \log(R_\alpha(\Theta'_0, Q)/\kappa) \log(1/\delta)}\right),$$

$$T = \Omega\left(\frac{p\alpha^4 \log^3(R_\alpha(\Theta'_0, Q)/\kappa) (M^2 + \frac{M\psi}{\lambda^2}) \max\{\Delta^4, M^2\psi^2\} \log^2(1/\delta)}{\varepsilon^4 m^4 \kappa} \cdot \max\left\{\log(\alpha), \frac{\alpha\kappa}{p}\right\}\right)$$

and an appropriate choice of η , outputting a sample from Θ'_T is (ε, δ) -DP. Furthermore,

$$R_\alpha(\Theta'_T, Q) \leq \kappa.$$

Proof In Theorem D.9 we need:

$$T\eta = \Theta\left(\frac{\alpha}{c} \log\left(\frac{R_\alpha(\Theta'_0, Q)}{\kappa}\right)\right)$$

Plugging in Lemmas 3.1 and D.8 for outputting Θ'_T from (12) to be (ε, δ) -DP it suffices if:

$$\beta = O\left(\frac{\varepsilon^2 m \exp(-\beta(3\psi + Mp\lambda^2))}{\max\{\Delta^2, M\psi\} \alpha \log(R_\alpha(\Theta'_0, Q)/\kappa) \log(1/\delta)}\right).$$

If we assume $\lambda \leq \sqrt{\frac{\psi}{Mp}}$, for $\varepsilon \leq 1$ this condition implies $\beta = O(\frac{1}{\psi + Mp\lambda^2})$. So this can be simplified to:

$$\beta = O\left(\frac{\varepsilon^2 m}{\max\{\Delta^2, M\psi\} \alpha \log(R_\alpha(\Theta'_0, Q)/\kappa) \log(1/\delta)}\right).$$

Now we can plug in this value of β , the smoothness bound from Theorem D.7, Lemma D.8, and our assumption on λ into the number of iterations T to get an iteration complexity requirement of:

$$T = \Omega\left(\frac{p\alpha^4(M^2 + \frac{M\psi}{\lambda^2}) \max\{\Delta^4, M^2\psi^2\} \log^2(1/\delta)}{\varepsilon^4 m^4 \kappa} \log^3(R_\alpha(\Theta'_0, Q)/\kappa) \max\{\log(\alpha), \frac{\alpha\kappa}{p}\}\right).$$

This completes the proof of Theorem D.10. \blacksquare

In order to bound the initial divergence, we use Θ'_0 that is a normal distribution centered at a point in the convex hull of minimizers of ℓ , i.e. the convex hull of $\{\arg \min_\theta \ell(\theta; d) : d \in \tau\}$.

Lemma D.11 *Suppose $\|\nabla \mathcal{L}(\theta; D) - \nabla \mathcal{L}(\theta; D')\|_2 \leq \Delta$, and let θ_0 be an arbitrary point in the convex hull of $\{\arg \min_\theta \ell(\theta; d) | d \in \tau\}$. Let P be the Gibbs distribution on $\beta \cdot \tilde{\mathcal{L}}'$ as defined in Theorem D.7. Then for $\alpha \geq 2$:*

$$R_\alpha\left(N\left(\theta_0, \frac{1}{\beta M} \mathbb{I}_p\right), P\right) = O\left(\frac{\alpha \Delta \beta^2 M^2}{m^2} + p \ln\left(\frac{M}{m}\right) + \beta \psi\right).$$

Proof By e.g. Lemma 16 in Ganesh and Talwar (2020), if θ^* is the true minimizer of \mathcal{L} and P_1 is the Gibbs distribution on $\beta \cdot \mathcal{L}$, for all $\alpha \geq 1$:

$$R_\alpha\left(N\left(\theta^*, \frac{1}{\beta M} \mathbb{I}_p\right), P_1\right) \leq \frac{p \ln\left(\frac{M}{m}\right)}{2}.$$

By Theorem D.7, $\beta \cdot \mathcal{L}$ and $\beta \cdot \tilde{\mathcal{L}}'$ differ by at most $4\beta\psi$ everywhere if $\lambda \leq \sqrt{\frac{\psi}{Mp}}$. Then:

$$R_\infty(P_1, P) \leq 8\beta\psi.$$

So by the approximate triangle inequality for Rényi divergences (Fact C.9):

$$R_\alpha\left(N\left(\theta^*, \frac{1}{\beta M} \mathbb{I}_p\right), P\right) \leq \frac{p \ln\left(\frac{M}{m}\right)}{2} + 8\beta\psi$$

Finally, by the assumption on $\nabla \mathcal{L}$, all points in $\{\arg \min_\theta \ell(\theta; d) | d \in \tau\}$ are distance at most Δ/m apart by strong convexity. Then by Lemma C.7:

$$R_\alpha\left(N\left(\theta_0, \frac{1}{\beta M} \mathbb{I}_p\right), N\left(\theta^*, \frac{1}{\beta M} \mathbb{I}_p\right)\right) \leq \frac{\alpha \Delta \beta^2 M^2}{2m^2}$$

Applying Fact C.9 again gives Lemma D.11. \blacksquare

Now by Pinsker’s inequality and monotonicity of Rényi divergences, we can use Lemma D.11, $\alpha = 2$, and $\kappa = \sqrt{\delta/2}$ to get a total variation distance bound of δ to the stationary distribution in Theorem D.10, proving Theorem 3.2.

Appendix E. Optimal DP-ERM/SCO Bounds from Rashomon Samplers

In this section we show that just using the Rashomon sampler with $\psi = 0$ as a primitive is enough to derive near-optimal bounds for DP-ERM/SCO in all settings.

E.1. Pure DP, Convex Losses

For L -Lipschitz convex losses on \mathcal{C} , the best possible excess empirical risk under ε -DP is $O\left(\frac{L\|\mathcal{C}\|_{2^p}}{\varepsilon n}\right)$. This bound is achieved by the exponential mechanism as shown in Bassily et al. (2014), which is exactly what the Gibbs distribution is for $\psi = 0$.

The best possible excess population risk under ε -DP is $O\left(\frac{L\|\mathcal{C}\|_{2^p}}{\varepsilon n} + \frac{L\|\mathcal{C}\|_2}{\sqrt{n}}\right)$. We can achieve this via Theorem 2.6 by setting $\psi = 0$ and adding the regularizer $\frac{L}{2\|\mathcal{C}\|_2\sqrt{n}}\|\theta - \theta_0\|_2^2$, where θ_0 is an arbitrary point in \mathcal{C} , to the loss function to give the algorithm uniform stability. The excess population risk of the Gibbs distribution with respect to the regularized loss is $O\left(\frac{L\|\mathcal{C}\|_{2^p}}{\varepsilon n} + \frac{L\|\mathcal{C}\|_2}{\sqrt{n}}\right)$ by Theorem 2.6, and the regularized and unregularized loss differ by at most $O\left(\frac{L\|\mathcal{C}\|_2}{\sqrt{n}}\right)$ everywhere in \mathcal{C} . Putting it all together, we get the following:

Theorem E.1 *For convex, L -Lipschitz losses over \mathcal{C} , there exists an ε -DP algorithm with excess population risk $O\left(\frac{L\|\mathcal{C}\|_{2^p}}{\varepsilon n} + \frac{L\|\mathcal{C}\|_2}{\sqrt{n}}\right)$.*

E.2. Pure DP and Strongly Convex Losses

For L -Lipschitz, m -strongly convex losses on \mathcal{C} , the best possible excess empirical risk under ε -DP is $O\left(\frac{L^2 p^2}{\varepsilon^2 n^2 m}\right)$. Unfortunately, the guarantee given by Theorem 2.4 is worse by a quadratic factor. In Bassily et al. (2014), the optimal excess empirical risk is achieved (up to log factors) by first choosing a smaller ball using the Laplace mechanism, and then running the Gibbs sampler on this smaller ball. We show the best-known bound can be achieved using the Gibbs sampler only as a primitive, which does not allow us to use the Laplace mechanism. We propose the iterated exponential mechanism, given as Algorithm 2, with the following guarantee:

Algorithm 2 Iterated Exponential Mechanism

Require: Loss function \mathcal{L} , constraint set \mathcal{C}_0 , Lipschitz constant L , strong convexity parameter m , number of iterations k , privacy parameter sequence $\{\varepsilon_i\}_{i=1}^k$, and data set D of n samples.

- 1: **for** $i = 1$ to k **do**
 - 2: Sample θ_i from \mathcal{C}_{i-1} with probability proportional to $\exp\left(-\frac{\varepsilon_i n}{2L\|\mathcal{C}_{i-1}\|_2}\mathcal{L}(\theta; D)\right)$.
 - 3: $\mathcal{C}_i \leftarrow \left\{ \theta \in \mathcal{C}_{i-1} : \|\theta - \theta_i\|_2 \leq \sqrt{\frac{cL(p+3\log n)\|\mathcal{C}_{i-1}\|_2}{m\varepsilon_i n}} \right\}$.
 - 4: **end**
 - 5: **return** θ_k
-

Theorem E.2 For any ε , suppose we run Algorithm 2 with $k = 1 + \lceil \log \log(\frac{\varepsilon n}{p + \log n}) \rceil$, and $\varepsilon_i = \varepsilon/2^{k-i+1}$. Assume each of the individual loss function in $\mathcal{L}(\theta; D)$ is L -Lipschitz within the constraint set \mathcal{C}_0 . Then Algorithm 2 is ε -differentially private. Additionally, if the loss function $\mathcal{L}(\theta; D)$ is m -strongly convex, and the constraint set \mathcal{C}_0 is convex, then over the randomness of the algorithm, the output θ_k of Algorithm 2 has excess empirical risk:

$$O\left(\frac{L^2(p^2 + p \log n)}{\varepsilon^2 n^2 m}\right).$$

The theorem follows by solving a recurrence for $\|\mathcal{C}_i\|_2$ to bound the diameter of the final set \mathcal{C}_{k-1} . Then, we show that the minimizer over \mathcal{C}_0 is also in \mathcal{C}_{k-1} with high probability, so the analysis of the exponential mechanism run on \mathcal{C}_{k-1} gives the theorem. We note that in addition to only using the Gibbs sampler as a primitive, we improve the $p^2 \log n$ in Bassily et al. (2014) result to $p^2 + p \log n$. To bound the excess loss, we first need the following lemma, which shows that with high probability we choose a series of \mathcal{C}_i that all contain the optimal θ for \mathcal{C}_0 .

The following lemma follows from a tail bound on the excess loss of the exponential mechanism, and using m -strong convexity to translate this into a distance bound.

Lemma E.3 Suppose we sample θ from the convex constraint set \mathcal{C} with probability proportional to $\exp\left(-\frac{\varepsilon n}{2L\|\mathcal{C}\|_2}\mathcal{L}(\theta; D)\right)$, where $\mathcal{L}(\cdot; D)$ is an m -strongly convex function. Let $\theta^* = \arg \min_{\theta \in \mathcal{C}} \mathcal{L}(\theta; D)$. Then for any $t \geq 0$ and for some sufficiently large constant c we have

$$\Pr\left[\|\theta - \theta^*\|_2 \leq \sqrt{\frac{cL(p+t)\|\mathcal{C}\|_2}{m\varepsilon n}}\right] \geq 1 - 2^{-t}.$$

Proof By e.g. the proof of (Bassily et al., 2014, Theorem III.2), we know that for some sufficiently large constant c :

$$\Pr\left[\mathcal{L}(\theta; D) - \mathcal{L}(\theta^*; D) \leq \frac{cL\|\mathcal{C}\|_2}{2\varepsilon n}(p+t)\right] \geq 1 - 2^{-t}. \quad (13)$$

We now show that the claim holds conditioned on this event. By optimality of θ^* and convexity of \mathcal{C} , we know

$$\langle \nabla \mathcal{L}(\theta^*; D), \theta - \theta^* \rangle \geq 0. \quad (14)$$

So, by m -strong convexity, we have

$$\begin{aligned} \frac{cL\|\mathcal{C}\|_2}{2\varepsilon n}(p+t) &\stackrel{(13)}{\geq} \mathcal{L}(\theta; D) - \mathcal{L}(\theta^*; D) \\ &\geq \langle \nabla \mathcal{L}(\theta^*; D), \theta - \theta^* \rangle + \frac{m}{2}\|\theta - \theta^*\|_2^2 \\ &\stackrel{(14)}{\geq} \frac{m}{2}\|\theta - \theta^*\|_2^2. \end{aligned}$$

Rearranging gives the claim in Lemma E.3. ■

Given Lemma E.3, we can now prove Theorem E.2.

Proof [Proof of Theorem E.2] The privacy guarantee is immediate from the privacy guarantee of the exponential mechanism, composition, and the fact that for this choice of ε_i, k , we have $\sum_{i=1}^k \varepsilon_i < \varepsilon$.

Setting $t = 3 \log n$ in Lemma E.3, in iteration i , letting $\theta_i^* = \arg \min_{\theta \in \mathcal{C}_{i-1}} \mathcal{L}(\theta; D)$, we have that with probability $1 - 2^{-t} = 1 - \frac{1}{n^3}$, $\theta_i^* \in \mathcal{C}_i$, and thus $\theta_i^* = \theta_{i+1}^*$. Then by a union bound, we have that with probability $1 - \frac{k}{n^3} \geq 1 - \frac{\log \log(\varepsilon n)}{n^3}$, $\theta_1^* \in \mathcal{C}_{k-1}$ (equivalently, $\theta_1^* = \theta_2^* = \dots = \theta_k^*$). When this event fails to happen, our excess loss is at most $L \|\mathcal{C}_0\|_2$, and in turn the contribution of this event failing to hold to the expected excess loss is $O\left(\frac{L \|\mathcal{C}_0\|_2 \log \log(\varepsilon n)}{n^3}\right)$, which is asymptotically less than our desired excess loss bound. So it suffices to provide the desired expected excess loss bound conditioned on this event. By the analysis of the exponential mechanism, conditioned on this event, we have that

$$\mathbb{E}_{\theta_k} [\mathcal{L}(\theta_k; D)] - \mathcal{L}(\theta_1^*; D) = O\left(\frac{Lp \|\mathcal{C}_{k-1}\|_2}{\varepsilon_k n}\right) = O\left(\frac{Lp \|\mathcal{C}_{k-1}\|_2}{\varepsilon n}\right). \quad (15)$$

Note that $\mathcal{L}(\theta_1^*; D) = \min_{\theta \in \mathcal{C}_0} \mathcal{L}(\theta; D)$ by definition, so it now suffices to bound $\|\mathcal{C}_{k-1}\|_2$ by $O\left(\frac{L(p+\log n)}{m\varepsilon n}\right)$. To do this, we have the recurrence relation:

$$\|\mathcal{C}_i\|_2 \leq 2 \sqrt{\frac{cL(p+3\log n) \|\mathcal{C}_{i-1}\|_2}{m\varepsilon_i n}}.$$

Solving the recurrence relation for \mathcal{C}_{k-1} , we get:

$$\begin{aligned} \|\mathcal{C}_{k-1}\|_2 &\leq \left(\frac{4cL(p+3\log n)}{mn}\right)^{1-2^{-(k-1)}} \cdot (\|\mathcal{C}_0\|_2)^{2^{-(k-1)}} \cdot \prod_{i=1}^{k-1} \varepsilon_i^{-2^{-(k-i)}} \\ &= \left(\frac{4cL(p+3\log n)}{m\varepsilon n}\right)^{1-2^{-(k-1)}} \cdot (\|\mathcal{C}_0\|_2)^{2^{-(k-1)}} \cdot \prod_{i=1}^{k-1} (2^{(k-i+1)})^{2^{-(k-i)}}. \end{aligned} \quad (16)$$

We claim the following:

$$\|\mathcal{C}_0\|_2 \leq \frac{2L}{m}. \quad (17)$$

Let θ_{global} be the minimizer of $\mathcal{L}(\theta; D)$ over all of \mathbb{R}^p . By triangle inequality, there exists a point θ in \mathcal{C}_0 which is at distance at least $\|\mathcal{C}_0\|_2/2$ far from θ_{global} . By m -strong convexity, this implies that the gradient at θ has ℓ_2 -norm at least $m \|\mathcal{C}_0\|_2/2$. Now, by Lipschitzness over \mathcal{C}_0 , we know that the gradient at θ has ℓ_2 -norm at most L . This gives us eq. (17).

Using eq. (17), we can simplify eq. (16) to

$$\|\mathcal{C}_{k-1}\|_2 \leq \frac{2L}{m} \cdot \left(\frac{2c(p+3\log n)}{\varepsilon n}\right)^{1-2^{-(k-1)}} \cdot \prod_{i=1}^{k-1} (2^{(k-i+1)})^{2^{-(k-i)}}.$$

We have:

$$\log_2 \left(\prod_{i=1}^{k-1} (2^{(k-i+1)})^{2^{-(k-i)}} \right) = \sum_{i=1}^{k-1} (k-i+1)2^{-(k-i)} \leq \sum_{j=1}^{\infty} (j+1)2^{-j} = 3.$$

In other words, $\prod_{i=1}^{k-1} (2^{(k-i+1)})^{2^{-(k-i)}}$ is at most 8, regardless of the value of k . Now, using the fact that $m^{1/\log m} = O(1)$ is a constant, our final upper bound on $\|\mathcal{C}_{k-1}\|_2$ is:

$$\|\mathcal{C}_{k-1}\|_2 = O \left(\frac{L}{m} \cdot \left(\frac{(p + \log n)}{\varepsilon n} \right)^{1-2^{-(k-1)}} \right) = O \left(\frac{L(p + \log n)}{m\varepsilon n} \right).$$

Plugging in eq. (15) gives us Lemma E.2. ■

The best possible population risk bound is $O\left(\frac{L^2 p^2}{\varepsilon^2 n^2 m} + \frac{L^2}{mn}\right)$. In order for Algorithm 2 to achieve this bound (up to log factors) we make a slight modification: we choose the radius of each ball defined by the algorithm such that the population minimizer, rather than the empirical minimizer, is in \mathcal{C}_{k-1} with high probability. Then, we can apply uniform stability of the Gibbs sampler on strongly convex losses to the exponential mechanism run on \mathcal{C}_{k-1} to get the following DP-SCO bound:

Theorem E.4 *Let θ_k be the output of Algorithm 2 except we let the radius of \mathcal{C}_i be $\sqrt{\frac{cL(p+3\log n)\|\mathcal{C}_{i-1}\|_2}{m\varepsilon_i n}} + \frac{cL\sqrt{\log n}}{m\sqrt{n}}$ instead of $\sqrt{\frac{cL(p+3\log n)\|\mathcal{C}_{i-1}\|_2}{m\varepsilon_i n}}$. Then θ_k has excess population risk*

$$O \left(\frac{L^2 p^2 \log n}{m\varepsilon^2 n^2} + \frac{L^2}{mn} \right).$$

We first show that the empirical minimizer is close to the population minimizer with high probability:

Lemma E.5 *Let ℓ be a m -strongly convex function and \mathcal{C} be a convex set such that for any d, θ ,*

$$\|\nabla \ell(\theta; d) - \mathbb{E}_{d \sim \mathcal{D}} [\nabla \ell(\theta; d)]\|_2 \leq \Delta,$$

and let $\theta^ := \arg \min_{\theta \in \mathcal{C}} \mathbb{E}_{d \sim \mathcal{D}} [\ell(\theta; d)]$ and $\theta^{\text{emp}} := \arg \min_{\theta \in \mathcal{C}} \ell(\theta; D)$. Then for $D \sim \mathcal{D}^n$, with probability $1 - \gamma$, we have:*

$$\|\theta^{\text{emp}} - \theta^*\|_2 = O \left(\frac{\Delta \sqrt{\log(1/\gamma)}}{m\sqrt{n}} \right)$$

Proof Consider a function $\tilde{\ell}$ which has gradient $\nabla \tilde{\ell}(\theta) = \nabla \ell(\Pi_{\mathcal{C}}(\theta)) + m(\theta - \Pi_{\mathcal{C}}(\theta))$. For any D , the empirical minimizer of $\tilde{\ell}$ over \mathbb{R}^p is equal to

$$\tilde{\theta}^{\text{emp}} := \theta^{\text{emp}} - \frac{1}{m} \cdot \nabla \ell(\theta^{\text{emp}}; D),$$

and the population minimizer of $\mathbb{E}_{d \sim \mathcal{D}} [\tilde{\ell}(\theta; d)]$ is

$$\tilde{\theta}^* := \theta^* - \frac{1}{m} \cdot \mathbb{E}_{d \sim \mathcal{D}} [\nabla \ell(\theta^*; d)].$$

By optimality of $\theta^{\text{emp}}, \theta^*$ and convexity of \mathcal{C} ,

$$\langle \nabla \ell(\theta^{\text{emp}}; D), \theta^* - \theta^{\text{emp}} \rangle \geq 0 \quad \text{and} \quad \langle \mathbb{E}_{d \sim \mathcal{D}} [\nabla \ell(\theta^*; D)], \theta^{\text{emp}} - \theta^* \rangle \geq 0,$$

This implies that $\left\| \tilde{\theta}^{\text{emp}} - \tilde{\theta}^* \right\|_2 \geq \|\theta^{\text{emp}} - \theta^*\|_2$. In addition, since projection to convex sets is a non-expansive map, $\tilde{\ell}$ is m -strongly convex if ℓ is, and for any d, θ we have

$$\|\nabla \ell(\theta; d) - \mathbb{E}_{d \sim \mathcal{D}} [\nabla \ell(\theta; d)]\|_2 = \left\| \nabla \tilde{\ell}(\theta; d) - \mathbb{E}_{d \sim \mathcal{D}} [\nabla \tilde{\ell}(\theta; d)] \right\|_2 \leq \Delta.$$

This holds for any D . Therefore, if we prove the lemma for $\tilde{\ell}$ and \mathbb{R}^p , then this would imply that the lemma holds for ℓ and \mathcal{C} . So it suffices to show the lemma for $\mathcal{C} = \mathbb{R}^p$.

If $\mathcal{C} = \mathbb{R}^p$ then $\mathbb{E}_{d \sim \mathcal{D}} [\nabla \ell(\theta; d)] = \mathbf{0}$. Now, by the assumptions in the lemma and a vector Azuma inequality (Hayes, 2003), we have $\|\nabla \ell(\theta^*; D)\|_2 = O\left(\frac{\Delta \sqrt{\log(1/\gamma)}}{\sqrt{n}}\right)$ with probability $1 - \gamma$ over D . Furthermore, we know $\nabla \ell(\theta^{\text{emp}}; D) = \mathbf{0}$ by strong convexity and since $\mathcal{C} = \mathbb{R}^p$. Then by strong convexity, we have

$$\|\theta^* - \theta^{\text{emp}}\|_2 \leq \frac{\|\nabla \ell(\theta^*; D) - \nabla \ell(\theta^{\text{emp}}; D)\|_2}{m} = \frac{\|\nabla \ell(\theta^*; D)\|_2}{m} = O\left(\frac{\Delta \sqrt{\log(1/\gamma)}}{m\sqrt{n}}\right)$$

with probability $1 - \gamma$ as desired in the statement of Lemma E.5. \blacksquare

Given Lemma E.5, if we want to ensure the *population* minimizer rather than empirical minimizer remains in the sets we choose in Algorithm 2, we just need to choose a slightly larger ball. From this modification and uniform stability, we get our DP-SCO bound:

Proof [Proof of Theorem E.4] Note that by L -Lipschitzness of ℓ in \mathcal{C} , we have

$$\|\nabla \ell(\theta; d) - \mathbb{E}_{d \sim \mathcal{D}} [\nabla \ell(\theta; d)]\|_2 \leq 2L.$$

By Lemma E.5, Lemma E.3, and a triangle inequality, we have that the population minimizer of ℓ in \mathcal{C}_i is in \mathcal{C}_{i+1} for each i with probability $1 - 2/n^3$. Then by a union bound, we have that the population minimizer is in \mathcal{C}_{k-1} . When this event fails to hold, our excess population risk is $O(L \|\mathcal{C}_0\|_2)$ and so the contribution of this event to the expected excess loss is $O\left(\frac{L \|\mathcal{C}_0\|_2 \log \log(\varepsilon n)}{n^3}\right) = O\left(\frac{L^2 \log \log(\varepsilon n)}{mn^3}\right)$, which is asymptotically less than our desired bound. So it suffices to provide the desired expected excess loss bound conditioned on this event. We can bound the radius of \mathcal{C}_{k-1} similarly to the proof of Theorem E.2, by noting that:

$$\|\mathcal{C}_i\|_2 \leq 2 \cdot \max \left\{ \sqrt{\frac{cL(p + 3 \log n) \|\mathcal{C}_{i-1}\|_2}{m\varepsilon_i n}}, \frac{cL\sqrt{\log n}}{m\sqrt{n}} \right\}$$

Then, rolling out the recursion, we have similarly to the proof of Theorem E.2:

$$\|\mathcal{C}_{k-1}\|_2 = O\left(\frac{L(p + \log n)}{m\varepsilon n} + \frac{L\sqrt{\log n}}{m\sqrt{n}}\right).$$

Now, combining Theorem C.11 and the uniform stability bound of Theorem 2.5 (for $\psi = 0$), we get that the expected excess population risk of θ_k compared to the population minimizer over \mathcal{C}_{k-1} is:

$$\begin{aligned} O\left(\frac{Lp\|\mathcal{C}_{k-1}\|_2}{\varepsilon n} + \frac{L^2}{mn}\right) &= O\left(\frac{L^2}{mn} \cdot \left(\frac{p^2 + \log^2 n}{\varepsilon^2 n} + \frac{p\sqrt{\log n}}{\varepsilon\sqrt{n}} + 1\right)\right) \\ &= O\left(\frac{L^2(p^2 \log n + \log^2 n)}{m\varepsilon^2 n^2} + \frac{L^2}{mn}\right). \end{aligned}$$

In the final equality, we use the fact that $\frac{p\sqrt{\log n}}{\varepsilon\sqrt{n}} \leq \max\{\frac{p^2 \log n}{\varepsilon^2 n}, 1\}$. We conclude by noting that conditioned on the event the population minimizer over \mathcal{C}_0 is contained in \mathcal{C}_{k-1} , θ_k has this same excess population risk bound compared to the population minimizer over \mathcal{C}_0 , completing the proof of Theorem E.4. \blacksquare

E.3. Approximate DP and Strongly Convex Losses

The results in Theorem 2.4, 2.7 and 2.6 with $\psi = 0$ combined immediately give that the Gibbs sampler achieves the optimal bounds of $O(\frac{L^2 p \log(1/\delta)}{\varepsilon^2 n^2 m})$ and $O(\frac{L^2 p \log(1/\delta)}{\varepsilon^2 n^2 m} + \frac{L^2}{mn})$ for excess empirical and population risk respectively in this setting. We note that one could also use Theorem 2.3 instead of Theorem 2.7 and obtain bounds that are within logarithmic factors of the optimal bounds, with a finite-time object.

E.4. Approximate DP and Convex Losses

Similarly to pure-DP, we can adapt our results in the strongly convex setting to the convex setting by adding a regularizer. For a near-optimal empirical guarantee, we use the regularizer $\frac{L\sqrt{p \log(1/\delta)}}{2\|\mathcal{C}\|_2 \varepsilon n} \|\theta - \theta_0\|_2^2$.

The regularized and unregularized losses differ by at most $O(\frac{L\|\mathcal{C}\|_2 \sqrt{p \log(1/\delta)}}{\varepsilon n})$ everywhere in \mathcal{C} , and the empirical excess loss bound we get by plugging in the strong convexity parameter $\frac{L\sqrt{p \log(1/\delta)}}{\|\mathcal{C}\|_2 \varepsilon n}$ into the bound for strongly convex losses is $O(\frac{L\|\mathcal{C}\|_2 \sqrt{p \log(1/\delta)}}{\varepsilon n})$, giving nearly the optimal bound of $O(\frac{L\|\mathcal{C}\|_2 \sqrt{p \log(1/\delta)}}{\varepsilon n})$.

For population risk, we use the regularizer $\frac{L}{2\|\mathcal{C}\|_2} \cdot \max\{\frac{\sqrt{p \log(1/\delta)}}{\varepsilon n}, \frac{1}{\sqrt{n}}\} \|\theta - \theta_0\|_2^2$. By a similar argument, this gives the optimal bound of $O(\frac{L\|\mathcal{C}\|_2 \sqrt{p \log(1/\delta)}}{\varepsilon n} + \frac{L\|\mathcal{C}\|_2}{n})$.

Appendix F. Lower Bound on DP-ERM for Non-Convex Losses

In this section, we show the following lower bound on the excess empirical risk for 1-Lipschitz non-convex loss functions. The lower bound implies that that there is no advantage, in terms of the dependence on dimensions (p), to move from ε -DP to (ε, δ) -DP.

Theorem 2 Let $\varepsilon \leq 1$, $2^{-\Omega(n)} \leq \delta \leq 1/n^{1+\Omega(1)}$, and $B(\mathbf{0}, 1)$ be a unit Euclidean ball centered at origin. Then there exists 1-Lipschitz non-convex function $\mathcal{L} : B(\mathbf{0}, 1) \times \mathcal{X} \rightarrow \mathbb{R}$ and a dataset¹¹ $D = \{d_1, \dots, d_n\}$ such that for every $p \in \mathbb{N}$, there is no (ε, δ) -differentially private algorithm \mathcal{A} that outputs θ^{priv} such that

$$\mathbb{E} \left[\mathcal{L}(\theta^{\text{priv}}; D) - \min_{\theta \in \mathbb{B}_p(1)} \mathcal{L}(\theta; D) \right] = o \left(\frac{p \log(1/\delta)}{n\varepsilon} \right), \quad (18)$$

Proof We first perform two translations of Lemma C.13: first from $(1, \frac{1}{ns})$ to (ε, δ) from Steinke and Ullman (2015) and then from sample complexity to a result stated in the terms of accuracy bound. A direct corollary of Lemma C.13 with $k = 1$ is as follows: for every $s \in \mathbb{N}$, no (ε, δ) -differentially private algorithm on input X satisfying the premise of Lemma C.13 outputs an index $j \in [s]$ such that

$$\mathbb{E}_{\mathcal{M}} \left[\frac{1}{n} \sum_{i=1}^n \mathbf{x}_{i,j} \right] - \max_{u \in [s]} \frac{1}{n} \sum_{i=1}^n \mathbf{x}_{i,u} = o \left(\frac{1}{n\varepsilon} \log(s) \log(1/\delta) \right), \quad (19)$$

where $\varepsilon \leq 1$ and $2^{-\Omega(n)} \leq \delta \leq 1/n^{1+\Omega(1)}$.

Using this lower bound on top-selection, we give our lower bound by defining an appropriate non-convex loss function. In particular, we define a packing over the p -dimensional Euclidean ball such that there is a bijective mapping between the centers of the packing and $[s]$. Then the function attains the minimum at the center of packing which corresponds to the coordinate $j \in [s]$ with maximum frequency. Since the size of the α -net is $\approx 1/\alpha^p$ and there is a bijective mapping, this gives a lower bound using eq. (19).

Let $B(\mathbf{0}, 1)$ be the p -dimensional Euclidean ball centered at origin and let $\alpha \in (0, 1/2)$ be a constant. Consider an α -packing with centers $C = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \}$. It is known that the size of such packing, $N(\alpha)$ is $(\frac{1}{\alpha})^p \leq N(\alpha) \leq (\frac{3}{\alpha})^p$. Let $s = N(\alpha)$. Further, let $f : B(\mathbf{0}, 1) \rightarrow \{1, \dots, s\}$ be an injective function defined as follows:

$$f(\theta) = \left\{ j : \mathbf{c}_j = \arg \min_{\mathbf{c} \in C} \|\theta - \mathbf{c}\|_2 \right\}.$$

In particular, f is the function that maps a point on the unit ball to its closest point in C .

We now define our loss function as follows:

$$\mathcal{L}(\theta; D) := \frac{1}{n} \sum_{d_i \in D} \ell(\theta; d_i) \text{ where } \ell(\theta; d_i) = \min_{\mathbf{c}_j \in C} \left(\frac{\|\theta - \mathbf{c}_j\|}{\alpha} - 1 \right) d_{i,j}. \quad (20)$$

For Lipschitz property, note that each loss function is $1/\alpha$ -Lipschitz because the gradient when it is defined is just $\frac{\theta - \mathbf{c}_j}{\alpha \|\theta - \mathbf{c}_j\|_2}$. We prove it formally.

Consider any θ, θ' in $B(\mathbf{0}, 1)$ and a data point $d_i \in D$. We wish to show $|\ell(\theta; d_i) - \ell(\theta'; d_i)| \leq \frac{1}{\alpha} \|\theta - \theta'\|_2$. We can split the line segment from θ to θ' into a sequence of line segments

$$(\theta_0, \theta_1), (\theta_1, \theta_2), \dots, (\theta_{k-1}, \theta_k),$$

11. The dataset, $D = \{d_1, \dots, d_n\}$ is such that $d_i \in \{0, 1\}^s$ for all $i \in [n]$, $d_{i,j}$ is independent (conditioned on P) and $\mathbb{E}[d_{i,j}] = P^j$ for all $i \in [n]$ and $j \in [s]$. Here \mathcal{P} is the distribution that is defined in Lemma C.13.

where $\theta_0 = \theta, \theta_k = \theta'$, such that for any line segment (θ_m, θ_{m+1}) , θ_m and θ_{m+1} share a minimizer in C of $\left(\frac{\|\theta - \mathbf{c}_j\|_2}{\alpha}\right) d_{i,j}$.¹²

It now suffices to show $|\ell(\theta_m; d_i) - \ell(\theta_{m+1}; d_i)| \leq \frac{1}{\alpha} \|\theta_m - \theta_{m+1}\|_2$ for each m , since we then have:

$$|\ell(\theta; d_i) - \ell(\theta'; d_i)| \leq \sum_{m=0}^{k-1} |\ell(\theta_m; d_i) - \ell(\theta_{m+1}; d_i)| \leq \frac{1}{\alpha} \sum_{m=0}^{k-1} \|\theta_m - \theta_{m+1}\|_2 = \frac{1}{\alpha} \|\theta - \theta'\|_2.$$

Let \mathbf{c}_j be a shared minimizer of $\left(\frac{\|\theta - \mathbf{c}_j\|_2}{\alpha}\right) d_{i,j}$ for θ_m and θ_{m+1} . If $d_{i,j} = 0$, then trivially $|\ell(\theta_m; d_i) - \ell(\theta_{m+1}; d_i)| \leq \frac{1}{\alpha} \|\theta_m - \theta_{m+1}\|_2$. Otherwise $d_{i,j} = 1$ and by triangle inequality, we have:

$$|\ell(\theta_m; d_i) - \ell(\theta_{m+1}; d_i)| = \left| \frac{\|\theta_m - \mathbf{c}_j\|_2}{\alpha} - \frac{\|\theta_{m+1} - \mathbf{c}_j\|_2}{\alpha} \right| \leq \frac{1}{\alpha} \|\theta_m - \theta_{m+1}\|_2.$$

Now let us suppose there is an (ε, δ) -differentially private algorithm \mathcal{A} that on input a non-convex function \mathcal{L} and n data points $\{d_1, \dots, d_n\}$, outputs a θ^{priv} such that

$$\mathbf{E}_{\mathcal{A}} [\mathcal{L}(\theta^{\text{priv}}; D)] - \min_{\theta \in B(1)} \mathcal{L}(\theta; D) = o\left(\frac{p \log(1/\delta)}{n\varepsilon}\right), \quad (21)$$

where $D = \{d_1, \dots, d_n\}$.

We will construct an algorithm that uses \mathcal{A} as subroutine and solve top-selection problem with an error $o(\log(s))$, contracting the lower bound of Lemma C.13.

Algorithm \mathcal{B} :

- On input $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, invokes \mathcal{A} on the function defined by eq. (20) and data points X to get θ^{priv} as output.
- Output $f(\theta^{\text{priv}})$.

Since the last step is post-processing, \mathcal{B} is (ε, δ) -differentially private. We now show that if \mathcal{A} outputs a θ^{priv} satisfying eq. (21), then $j := f(\theta^{\text{priv}})$ satisfies eq. (19) leading to a contradiction.

First note that, for any $\mathbf{c} \in C$ and all $\theta \in \mathbb{B}_p(\mathbf{c}, \alpha)$ such that $\|\theta - \mathbf{c}\|_2 \leq \frac{\alpha}{2}$,

$$\mathcal{L}(\mathbf{c}; D) = -\frac{1}{n} \sum_{i=1}^n \mathbf{x}_{i,f(\mathbf{c})} \leq \mathcal{L}(\theta; D).$$

Therefore,

$$\mathcal{L}(\theta^*; X) := \min_{\mathbf{c} \in C} \mathcal{L}(\mathbf{c}, X) = \min_{\mathbf{c} \in C} \left(-\frac{1}{n} \sum_{i=1}^n \mathbf{x}_{i,f(\mathbf{c})} \right)$$

12. In particular, for each \mathbf{c}_j let B_j be the set of points in $B(\mathbf{0}, 1)$ such that \mathbf{c}_j is a minimizer of $\left(\frac{\|\theta - \mathbf{c}_j\|_2}{\alpha}\right) d_{i,j}$. We can split the line segment from θ to θ' at each point where it enters or leaves some B_j to get this sequence of line segments, and by this construction each line segment's endpoints are both in B_j for some j .

This implies that

$$f(\theta^*) = \arg \max_{1 \leq j \leq s} \frac{1}{n} \sum_{i=1}^n \mathbf{x}_{i,j},$$

which is exactly the top-selection problem. Therefore, eq. (21) implies eq. (19) because $p \log\left(\frac{1}{\alpha}\right) \leq \log(s) \leq p \log\left(\frac{3}{\alpha}\right)$ and $\alpha \in (0, 1/2)$ is a constant. ■