

(Nearly) Optimal Private Linear Regression via Adaptive Clipping

Prateek Varshney

VPRATEEK@GOOGLE.COM

Abhradeep Thakurta

ATHAKURTA@GOOGLE.COM

Prateek Jain

PRAJAIN@GOOGLE.COM

Editors: Po-Ling Loh and Maxim Raginsky

Abstract

We study the problem of differentially private *linear regression* where each data point is sampled from a fixed sub-Gaussian style distribution. We propose and analyze a one-pass mini-batch stochastic gradient descent method (DP-AMBSSGD) where points in each iteration are sampled *without* replacement. Noise is added for DP but the noise standard deviation is estimated online. Compared to existing (ϵ, δ) -DP techniques which have sub-optimal error bounds, DP-AMBSSGD is able to provide nearly optimal error bounds in terms of key parameters like dimensionality d , number of points N , and the standard deviation σ of the noise in observations. For example, when the d -dimensional covariates are sampled i.i.d. from the normal distribution, then the excess error of DP-AMBSSGD due to *privacy* is $\frac{\sigma^2 d}{N} (1 + \frac{d}{\epsilon^2 N})$, i.e., the error is meaningful when number of samples $N = \Omega(d \log d)$ which is the standard operative regime for linear regression. In contrast, error bounds for existing efficient methods in this setting are: $\mathcal{O}(\frac{d^3}{\epsilon^2 N^2})$, even for $\sigma = 0$. That is, for constant ϵ , the existing techniques require $N = \Omega(d\sqrt{d})$ to provide a non-trivial result.

1. Introduction

Machine Learning (ML) models are known to be susceptible to leaks of sensitive private information of individual points in the training data. In fact, this risk is non-trivial even for problems as simple and canonical as linear regression (Hsu et al., 2012; Dieuleveut et al., 2016; Jain et al., 2018).

Several existing works have studied privacy preserving linear regression and more generally, private convex optimization with differential privacy as the privacy notion (Chaudhuri et al., 2011; Kifer et al., 2012; Bassily et al., 2014; Song et al., 2013; Bassily et al., 2019; McMahan et al., 2017; Wu et al., 2017; Andrew et al., 2021; Feldman et al., 2020; Bassily et al., 2020; Song et al., 2020; Iyengar et al., 2019). While tight upper and lower bounds are known for generic problem classes (e.g., convex Lipschitz losses, and strongly convex losses) (Bassily et al., 2014, 2019), surprisingly, the popular special instance of *linear regression* is much less mapped (Smith et al., 2017; Sheffet, 2019; Liu et al., 2021; Cai et al., 2021).

In this paper, we study the problem of DP linear regression (DP-LR) and provide (nearly) optimal excess population risk guarantees under standard assumptions, i.e., the data is sampled i.i.d. from a sub-Gaussian distribution. In particular, we study the following problem: Consider a dataset $D = \{(\mathbf{x}_0, y_0), \dots, (\mathbf{x}_{N-1}, y_{N-1})\}$ drawn i.i.d. from some fixed distribution \mathcal{D} with $\mathbf{H} = \mathbb{E}[\mathbf{x}\mathbf{x}^\top]$. Without loss of generality, $y = \langle \mathbf{x}, \mathbf{w}^* \rangle + z$ where $\mathbb{E}[z \cdot \mathbf{x}] = 0$ and $\mathbb{E}[z^2] = \sigma^2$. The goal in DP-LR is to output a model \mathbf{w}_{priv} that preserves privacy and also approximately minimizes the excess population risk i.e., $\mathcal{L}(\mathbf{w}_{\text{priv}}) - \mathcal{L}(\mathbf{w}^*)$, where,

$$\mathcal{L}(\mathbf{w}) = \frac{1}{2} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [(y - \langle \mathbf{x}, \mathbf{w} \rangle)^2].$$

Now, even for the above set of assumptions, the existing *polynomial time* methods require $N \geq d\sqrt{d}$ for them to be non-vacuous (for constant z). Recent work by Liu et al. (2021) indeed obtains strong error rates similar to our results, however their proposed method is exponential in d and N . In contrast, our methods are linear in both d and N .

The key issue in most of the existing works is that the global Lipschitz constant of \mathcal{L} scales as $\langle \mathbf{x}, \mathbf{w} - \mathbf{w}^* \rangle \leq \|\mathbf{x}\|_2 \|\mathbf{w} - \mathbf{w}^*\|_2$ which implies that sensitivity of gradient would already have $O(d/N)$ term that leads to rate of $d\sqrt{d}/N$ as d -dimensional noise vector with $O(d/N)$ standard deviation has norm of the order: $d\sqrt{d}/N$. This issue recurs in most of the existing works, as intuitively they try to ensure that *all* the directions of covariance are privacy preserving, while our goal is to ensure only direction along \mathbf{w}^* is differentially private.

In this work, we propose a method DP-SSGD (DP-Shuffled SGD) to get around this challenge by using a one-pass noisy SGD method that samples points *without* replacement and uses tail averaging. As we take only one-pass over the shuffled data, we can ensure that the t^{th} iterate \mathbf{w}_t is completely independent of the sampled \mathbf{x}_t and hence, \mathbf{w}_t are independent in each iteration. This implies that $\langle \mathbf{x}, \mathbf{w} \rangle$ can have a significantly tighter bounded for sub-Gaussian style distributions. In fact, we can separate the loss using standard bias-variance decomposition of the loss (Jain et al., 2018). This decomposition along with amplification by shuffling (Feldman et al., 2021), ensures that the excess risk $\mathcal{L}(\mathbf{w}_{\text{priv}}) - \mathcal{L}(\mathbf{w}^*)$ of the method is bounded by $\tilde{O}(\sigma^2 d/N + (1 + \sigma^2)d^2/\varepsilon^2 N^2)^1$. However, the proposed method has three issues: a) the sample complexity of the method is $N = \tilde{\Omega}(d^2)$, b) requires $\varepsilon \leq 1/\sqrt{N}$, and c) the second term is sub-optimal w.r.t. σ .

To address the above mentioned weaknesses with DP-SSGD’s analysis, we modify DP-SSGD to obtain a new method – DP-AMBSSGD (DP-Adaptive-Mini-Batch-Shuffled-SGD) – which divides the data into mini-batches and runs one pass of shuffled mini-batch SGD but where the noise is set adaptively according to the excess error in each iteration. For this method, we can obtain excess risk of the form $\tilde{O}(\sigma^2 d/N + \sigma^2 d^2/\varepsilon^2 N^2)$ with sample complexity $N \geq d \log^2 d$. That is, the proposed method is efficient with time complexity of only $O(Nd)$ while still ensuring nearly optimal error rate that matches the lower bound in Cai et al. (2021), up to constants and condition number factors.

Below we provide an informal version of this result in a simplified setting.

Theorem 1 (Informal result) *Let $\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $y_i = \langle \mathbf{x}_i, \mathbf{w}^* \rangle + z_i$, for $\mathbf{w}^* \in \mathbb{R}^d$ and $z_i \sim \sigma \mathcal{N}(0, 1)$. Then, there exists a method for DP-LR that guarantees (ε, δ) -DP, has time-complexity $O(Nd)$ and has excess risk bounded by (with probability $\geq 1 - 1/N^{100}$ over randomness in data and algorithm):*

$$\mathcal{L}(\mathbf{w}_{\text{priv}}) - \mathcal{L}(\mathbf{w}^*) \leq 8\sigma^2 \frac{d}{N} + \tilde{O}\left(\sigma^2 \frac{d^2}{\varepsilon^2 N^2}\right), \text{ if } N \geq d \log^3 d.$$

As mentioned above, the excess risk bound matches the lower bound by Cai et al. (2021). Furthermore, the algorithm is optimal in terms of time complexity. Finally, in the same setting, existing methods (Cai et al., 2021) have an excess risk bound of $\sigma^2 \frac{d}{N} + \frac{(d\|\mathbf{w}^*\|_{\mathbf{H}}^2 + \sigma^2) \cdot d^2}{\varepsilon^2 N^2}$ if $N \geq d^{3/2}$. Note that the above bound is significantly sub-optimal for $\sigma \leq \sqrt{d}$ which is a fairly practical setting. Also note that the bound in Cai et al. (2021) is stated only in the setting of $\sigma \geq \sqrt{d}$. Hence the claim of matching upper-lower bounds by Cai et al. (2021) hold only for specific setting of σ . In contrast, our result holds for all range of σ , especially for theoretically and empirically important settings of small

1. In $\tilde{\mathbf{w}}(\cdot)$, $\tilde{O}(\cdot)$, and $\tilde{\Omega}(\cdot)$ we hide polylog factors in N and $(1/\delta)$.

σ . Furthermore, while Wang et al. (2017) also claims optimal dimension dependence in excess risk bound, their results (and the ones cited within) hide factors corresponding to the Lipschitz constant (L) of the least squares function. As mentioned above, L can scale as $\|\mathbf{x}\|_2 \|\mathbf{w}^*\|_2$, and thus would have an additional factor of \sqrt{d} for the standard Gaussian distribution. Clearly, when $\varepsilon > \sqrt{d}/\sqrt{N}$, the first term dominates and hence one can assume that the privacy is for "free" asymptotically. Otherwise, when $\varepsilon = c\sqrt{d}/\sqrt{N}$ (with $c < 1$), the asymptotic bound translates to $\tilde{O}(d/cN)$, i.e., DP does a sample size reduction from $N \rightarrow cN$. This is common for DP-ERM. Finally, our results lead to optimal variance error, similar to non-private versions. Moreover, while the privacy related error term (second error term in Theorem 10) is optimal in d, N, ε , it has an additional dependence on κ —the conditional number of $\mathbf{H} = \mathbb{E}[\mathbf{xx}^\top]$. For first order methods, we do expect such a factor to be present due to composition theorems, but it is not clear what is the optimal dependence on κ for the privacy related error term. We leave further investigation into this factor for future work.

1.1. Our Algorithms

DP-Shuffled SGD (DP-SSGD): Our first algorithm is a fairly straightforward adaptation of standard DP-SGD (Song et al., 2013; Bassily et al., 2014), where we make a single pass over the dataset with single data sample mini-batches, along with appropriate clipping of the gradients² and addition of Gaussian noise. Finally, we output the average of last $N/2$ models produced by the algorithm. To have strong privacy guarantee, we randomly shuffle the dataset D prior to running the algorithm. Using the by-now standard tool of privacy amplification by shuffling (Úlfar Erlingsson et al., 2019; Feldman et al., 2021), we amplify the overall privacy guarantee (as compared to the analysis for the unshuffled dataset) by a factor of $\approx 1/\sqrt{N}$. With this tool in hand, it is not hard to show that one needs to add noise $\tilde{O}\left(\frac{\zeta}{\varepsilon\sqrt{N}}\right)$ to satisfy (ε, δ) -DP as long as $\varepsilon \leq \frac{1}{\sqrt{N}}$, where ζ is the clipping norm of the gradients (Abadi et al., 2016). Furthermore, using the bias-variance decomposition analysis of the linear regression loss (Jain et al., 2018), one can show that setting $\zeta = \tilde{O}(\sigma\sqrt{d})$ ensures a excess risk of $\tilde{O}(\sigma^2 d/N + (1 + \sigma^2)d^2/\varepsilon^2 N^2)$. For small values of the privacy parameter ε , this bound is optimal assuming the standard deviation of the inherent noise is $\Theta(1)$. However, the proposed method has three major drawbacks: a) requires $\varepsilon \leq 1/\sqrt{N}$, b) the sample complexity of the method is $N \geq \Omega(d^2)$ because of the upper bound on ε , and c) the second term in the error is sub-optimal w.r.t. σ . To that end, we propose Algorithm DP-AMBSSGD (DP-Adaptive-Mini-Batch-Shuffled-SGD) to address each of these issues.

DP-Adaptive-Mini-Batch-Shuffled-SGD (DP-AMBSSGD): This algorithm is based on the traditional DP-SGD framework (Differentially Private Stochastic Gradient Descent) (Song et al., 2013; Bassily et al., 2014; Abadi et al., 2016). For a mini-batch of data samples $\hat{D}_t = \{(\mathbf{x}_i, y_i)\}_{i=0}^{b-1}$, DP-AMBSSGD does a state update of the form $\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t - \eta \left(\frac{1}{b} \sum_{(\mathbf{x}, y) \in \hat{D}_t} \text{clip}_{\zeta_t}(\mathbf{x} \cdot (\langle \mathbf{x}, \mathbf{w}_t \rangle - y)) + \mathcal{N}\left(0, \frac{4\alpha^2 \zeta_t^2}{b^2}\right) \right)$, where α is called the noise multiplier (which is a function of only the number of iterations and the privacy parameters (ε, δ)), ζ_t is called the clipping norm, and $\text{clip}_{\zeta}(\boldsymbol{\nu}) = \boldsymbol{\nu} \cdot \min\left\{1, \frac{\zeta}{\|\boldsymbol{\nu}\|_2}\right\}$. The key idea behind the design of DP-AMBSSGD is to adaptively choose the clipping norm ζ_t s.t. a "overwhelming majority" of the gradients do not get clipped at any time step t . Observe that as one gets closer to the true model \mathbf{w}^* , the norm of the gradient of the linear regression

2. Clipping of a vector, corresponds to scaling down a vector to a fixed ℓ_2 -norm ζ , if its norm exceeds ζ .

loss at an data point (\mathbf{x}, y) approximately scales as $\mathbb{E} [\|\mathbf{x}\|_2] \cdot \sigma$, where σ is the standard deviation of the inherent noise. Since, σ and $\mathbb{E} [\|\mathbf{x}\|_2]$ are unknown, one cannot use this as a proxy for the clipping norm. We circumvent this issue as follows, at each time step t we estimate a threshold ζ_t (via a DP algorithm) s.t. the norms of 99% of the gradients in the mini-batch fall below ζ_t . We then use ζ_t as the clipping norm in the standard DP-SGD algorithm. We call the combined algorithm DP-AMBSSGD. As the training progresses and \mathbf{w}_t gets closer to the true model \mathbf{w}^* , concentration bound on ζ_t provably gets smaller and hence the noise added to ensure DP also progressively reduces. Using the same bias-variance analysis of DP-SSGD above will result in Theorem 10 that addresses all the three issues in Algorithm DP-SSGD. The privacy analysis follows from standard adaptive composition properties of DP (Bun and Steinke, 2016).

It is worth mentioning that such an iterative localization/clipping idea was also used in the context of mean estimation (Biswas et al., 2020). (Kamath et al., 2018) also previously used similar ideas for covariance estimation while (Feldman et al., 2020) used localization in a DP-SCO context.

It is also worth mentioning that the idea of *adaptive clipping* in DP-SGD has been empirically demonstrated to be beneficial (Andrew et al., 2021), but to our knowledge, DP-AMBSSGD is the first to formally demonstrate its advantage over fixed clipping.

1.2. Paper Organization

We present related works to our work in next subsection. We then define the problem and required notations in Section 2. Next, we present the one-pass SGD method and it’s analysis in Section 3.1. We then present our main result and the algorithm in Section 3.2. Finally, we conclude with a discussion of our results and several exciting future directions.

1.3. Other Related Works

Differentially private empirical risk minimization (DP-ERM) is probably one of the most well-studied subfields of differential privacy (Chaudhuri et al., 2011; Kifer et al., 2012; Bassily et al., 2014; Song et al., 2013; Abadi et al., 2016; Bassily et al., 2019; McMahan et al., 2017; Wu et al., 2017; Iyengar et al., 2019; Pichapati et al., 2019; Andrew et al., 2021; Feldman et al., 2020; Bassily et al., 2020; Song et al., 2020). The objective there is to solve the following problem while preserving

DP: $\arg \min_{\mathbf{w} \in \mathcal{C}} \frac{1}{N} \sum_{i=0}^{N-1} \ell(\mathbf{w}; d_i)$, where $D = \{d_0, \dots, d_{N-1}\}$ is the dataset. When the loss function

$\ell(\cdot; \cdot)$ is convex, tight excess population risk bound of $L \|\mathcal{C}\|_2 \cdot \tilde{\mathbf{w}} \left(\frac{1}{\sqrt{N}} + \frac{\sqrt{d}}{\varepsilon N} \right)$ is known under (ε, δ) -DP (Bassily et al., 2019, 2020), where L is the ℓ_2 -Lipschitz constant of the loss function $\ell(\cdot; \cdot)$ and \mathcal{C} is the diameter of the *convex* constraint set. The algorithm that achieves this bound is essentially a variant of differentially private stochastic gradient descent (DP-SGD) (Song et al., 2013; Bassily et al., 2014), or objective perturbation (Chaudhuri et al., 2011; Kifer et al., 2012). For the problem of linear regression considered in this paper, it is not hard to observe that $L \approx d$ when the constraint set $\|\mathcal{C}\|_2 = O(1)$. This means the sample complexity for the excess population bound above is $N = \tilde{\Omega}(d^2)$. One can take a direct approach, and perturb the sufficient statistics for linear regression (Smith et al., 2017; Sheffet, 2019) and achieve a sample complexity of $N = \tilde{\Omega}(d^{3/2})$. In this work we show that if the feature vectors are drawn from d -dimensional standard Normal distribution, we can get an excess population risk of $\tilde{O} \left(\frac{d}{N} + \frac{d^2}{\varepsilon^2 N^2} \right)$, which translates to a sample complexity of $N = \tilde{\Omega}(d)$. For DP sparse principal component analysis (PCA) (Chien et al., 2021),

and for mean estimation (Brown et al., 2021; Liu et al., 2021), similar improvements on the sample complexity were observed in prior works under ‘‘Gaussian’’ like assumptions. To the best of our knowledge, the techniques in these papers are complementary to ours, and cannot be used for the problem at hand.

A crucial component of our algorithms is that it makes a single pass over the data, after the dataset has been uniformly permuted. Our algorithms are thus an addition to the growing list of single pass algorithms (Feldman et al., 2020; Bassily et al., 2020) that achieve optimal population risk under (ε, δ) -DP.

2. Problem Definition and Preliminaries

Let $D = \{(\mathbf{x}_0, y_0), (\mathbf{x}_1, y_1), \dots, (\mathbf{x}_{N-1}, y_{N-1})\}$ be the input dataset with each point (\mathbf{x}_i, y_i) sampled i.i.d. from a distribution \mathcal{D} . Given D , the goal is to learn \mathbf{w}^* that minimizes the Least Squares Regression (LSR) problem:

$$(LR) : \quad \mathbf{w}^* = \arg \min_{\mathbf{w}} \mathcal{L}(\mathbf{w}) := \frac{1}{2} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [(y - \langle \mathbf{x}, \mathbf{w} \rangle)^2]. \quad (1)$$

Furthermore, the goal is to ensure that \mathbf{w}^* is privacy preserving (see definitions below). Without loss of generality, we have:

$$y_i = \langle \mathbf{x}_i, \mathbf{w}^* \rangle + z_i, \text{ where } \mathbb{E}[z_i \mathbf{x}_i] = 0, \text{ and } \mathbb{E}[z_i^2] = \sigma^2.$$

For simplicity, we assume that z is *independent* of \mathbf{x} . Note that our results hold for general case when z is dependent on \mathbf{x} with appropriately defined notation but for simplicity of exposition, we focus only the independent case.

We now discuss our notations and then later specify assumptions for the input distribution \mathcal{D} .

Notations: We use boldface lower letters such as \mathbf{x} , \mathbf{w} for vectors, boldface capital letters (e.g. \mathbf{A} , \mathbf{H}) for matrices and normal script font letters (e.g. \mathcal{M} , \mathcal{T}) for tensors. Let $\|\mathbf{A}\|_2$ denote the spectral norm of \mathbf{A} and $\|\mathbf{v}\|_2$ denote the Euclidean norm of \mathbf{v} . Unless specified otherwise, $\|\mathbf{A}\| = \|\mathbf{A}\|_2$ and $\|\mathbf{v}\| = \|\mathbf{v}\|_2$. Also, $\|\mathbf{x}\|_{\mathbf{A}} := \|\mathbf{A}^{1/2} \mathbf{x}\|$ and $\|\mathbf{M}\|_{\mathbf{A}} := \|\mathbf{A}^{-1/2} \mathbf{M} \mathbf{A}^{-1/2}\|_2$.

Let $\mathbf{H} := \mathbb{E}[\mathbf{x}\mathbf{x}^T]$ be the second-moment matrix and let $\mathcal{M} := \mathbb{E}[\mathbf{x} \otimes \mathbf{x} \otimes \mathbf{x} \otimes \mathbf{x}]$ be the fourth moment tensor. $\mathcal{T} \otimes \mathbf{a}$ denotes outer product of \mathcal{T} with \mathbf{a} . Let μ be the smallest eigenvalue of \mathbf{H} . Let R_x be the smallest R which satisfies $\mathbb{E}[\|\mathbf{x}\|_2^2 \mathbf{x}\mathbf{x}^T] \leq R^2 \mathbb{E}[\mathbf{x}\mathbf{x}^T]$. This implies $\mathbb{E}[\|\mathbf{x}\|_2^2] \leq R_x^2$. Finally, let $\Sigma = \mathbb{E}[(y - \langle \mathbf{x}, \mathbf{w}^* \rangle)^2 \mathbf{x}\mathbf{x}^T]$ be the error ‘‘covariance’’ matrix.

Definition 2 (($\mathbf{H}, K_2, a, \beta$)-Tail) *Let both \mathbf{H} and \mathcal{M} exist and are finite. Also, let $\mu > 0$ and $\kappa = \|\mathbf{H}\|_2 / \mu < \infty$. A random vector \mathbf{x} satisfies ($\mathbf{H}, K_2, a, \beta$)-Tail if the the following holds:*

- $\exists a > 0$ s.t. with probability $\geq 1 - \beta$,

$$\|\mathbf{x}\|_2^2 \leq \mathbb{E}[\|\mathbf{x}\|_2^2] \cdot \log^{2a}(1/\beta) \leq R_x^2 \cdot \log^{2a}(1/\beta), \quad (2)$$

- We have,

$$\max_{\mathbf{v}, \|\mathbf{v}\|=1} \mathbb{E} \left[\exp \left(\left(\frac{|\langle \mathbf{x}, \mathbf{v} \rangle|^2}{K_2^2 \|\mathbf{H}\|_2} \right)^{1/2a} \right) \right] \leq 1. \quad (3)$$

That is, from (2) and (3), for any fixed \mathbf{v} , w.p. $\geq 1 - \beta$:

$$(\langle \mathbf{x}, \mathbf{v} \rangle)^2 \leq K_2^2 \|\mathbf{H}\|_2 \|\mathbf{v}\|^2 \log^{2a}(1/\beta). \quad (4)$$

We assume the distribution on feature vectors $\mathbf{x} \sim \mathcal{D}$ satisfies $(\mathbf{H}, K_2, a, \beta)$ -Tail, and the inherent noise z satisfies $(\sigma^2 \mathbf{I}, K_2, a, \beta)$ -Tail. Note that if both \mathbf{x} and z are sampled from a normal distribution then the above requirements are satisfied for $a = 1/2$. Beyond sub-Gaussian, notice that Definition 2 captures a more general classes of distributions like subExponential (when we set $a = 1$).

Notation of Privacy: In this work, we operate in the standard setting of differential privacy (DP) (Dwork et al., 2006b,a) in the replacement model. In this paper, we assume $\varepsilon = O(1)$, and $\delta = O(1/\text{poly}(N))$, where N is the data set size. The formal definition is below.

Definition 3 (Differential privacy (Dwork et al., 2006b,a)) A randomized algorithm \mathcal{A} is (ε, δ) -differentially private if, for any pair of datasets D and D' differ in exactly one data point (i.e., one data point is present in one set and absent in another for both D and D'), and for all events \mathcal{S} in the output range of \mathcal{A} , we have

$$\Pr[\mathcal{A}(D) \in \mathcal{S}] \leq e^\varepsilon \cdot \Pr[\mathcal{A}(D') \in \mathcal{S}] + \delta,$$

where the probability is taken over the random coin flips of \mathcal{A} .

In the privacy analysis we use zCDP (zero-Concentrated Differential Privacy (Bun and Steinke, 2016)) for privacy accounting. However, we state the final privacy guarantee always in terms of (ε, δ) -DP. For completeness purposes we define zCDP below:

Definition 4 (zCDP Bun and Steinke (2016)) A randomized algorithm \mathcal{A} is ρ -zCDP if for any pair of data sets D and D' that differ in one record, we have $D_\alpha(\mathcal{A}(D) || \mathcal{A}(D')) \leq \rho\alpha$ for all $\alpha > 1$, where D_α is the Rényi divergence of order α .

3. Private Linear Regression

Recall that the goal is to find the optimal linear regression parameter vector while preserving differential privacy of each individual point (\mathbf{x}_i, y_i) . As mentioned earlier, existing techniques like DP-SGD (Abadi et al., 2016) suffer from error rates that require $N \geq d\sqrt{d}$ to be non-vacuous. Furthermore, it seems that the issue is fundamental to the method and the excess risk bounds are unlikely to improve by better analysis. For example, in DP-SGD, one needs to add noise proportional to the gradient of a randomly sampled point *with replacement*. Now, gradient of the loss term corresponding to the i^{th} data point is: $\nabla_{\mathbf{w}} L_i(\mathbf{w}) = (\langle \mathbf{x}_i, \mathbf{w} \rangle - y_i) \mathbf{x}_i$. As \mathbf{x}_i is sampled with replacement, \mathbf{w} can be dependent on \mathbf{x}_i (e.g. if the same index i is sampled twice in a row). This implies, that $\langle \mathbf{x}_i, \mathbf{w} \rangle$ might be as large as $\|\mathbf{x}_i\|_2 \|\mathbf{w}\|_2$. That is, for the simple setting of $\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\|\nabla_{\mathbf{w}} L_i(\mathbf{w})\|_2 \approx d$ which implies $O(d/N)$ noise would be required in each iteration of DP-SGD which leads to an additional error of $d\sqrt{d}/N$ in the estimation of \mathbf{w}^* .

The above observation provides a key motivation for our algorithm. In particular, it is clear that sampling with replacement can lead to challenging dependencies between the current iterate \mathbf{w} and the sampled point. Instead, if we sample points *without* replacement and take only *one* pass over the data, then at each step \mathbf{w} is *independent* of \mathbf{x}_i , which implies significantly better bound on the per step gradient. Naturally, the error bound can suffer if we take only *one* pass over the data. But by using some recent results (Jain et al., 2018, 2017) along with techniques like tail averaging, we provide nearly *optimal* excess risk bounds.

In the section below, we use the above intuition to propose the DP-Shuffled SGD (DP-SSGD) method and provide analysis of the algorithm. However, our result for the proposed method holds

only for $\varepsilon \leq 1/\sqrt{N}$, and has sub-optimal sample complexity. To alleviate these concerns, we modify the proposed method using mini-batches of data which provide a nearly optimal error bound without significant restrictions on privacy budget ε .

3.1. DP-Shuffled SGD

In this section, we present DP-SSGD (Algorithm 1) which is primarily based on DP-SGD with random shuffling i.e. sampling without replacement. That is, we first randomly permute the data points (Step 2). We then take one pass over the dataset, and update \mathbf{w} by gradient descent with the loss term corresponding to each data point. To ensure differential privacy, we clip the gradient norm to at most ζ and add Gaussian noise with standard deviation $2\eta\zeta\alpha$ (Step 6). Finally, we output the average of the last half of the iterates (Step 8).

Algorithm 1 DP-SSGD: DP-Shuffled SGD

- 1: **Input:** Samples: $\{(\mathbf{x}_i, y_i)\}_{i=0}^{N-1}$, Clipping Norm: ζ , DP Noise Multiplier: α , Learning Rate: η
 - 2: Randomly permute $\{(\mathbf{x}_i, y_i)\}_{i=0}^{N-1}$
 - 3: Initialize $\mathbf{w}_0 \leftarrow \mathbf{0}$
 - 4: **for** $t = 0 \dots N - 1$ **do**
 - 5: Sample $\mathbf{g}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{d \times d})$
 - 6: $\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t - \eta \left(\text{clip}_\zeta(\mathbf{x}_t(\langle \mathbf{x}_t, \mathbf{w}_t \rangle - y_t)) + \alpha \cdot 2\zeta \cdot \mathbf{g}_t \right)$ where $\text{clip}_\zeta(\boldsymbol{\nu}) = \boldsymbol{\nu} \cdot \min \left\{ 1, \frac{\zeta}{\|\boldsymbol{\nu}\|_2} \right\}$
 - 7: **end for**
 - 8: **Return:** $\bar{\mathbf{w}} := \frac{2}{N} \sum_{t=N/2+1}^N \mathbf{w}_t$
-

We first show that Algorithm 1 is (ε, δ) -DP. The proof follows using (Feldman et al., 2021, Theorem 3.8) which analyzes DP composition with "randomly shuffled" data.

Theorem 5 *Suppose we apply DP-SSGD (Algorithm 1) on N input samples with noise multiplier $\alpha = \Omega\left(\frac{\log(N/\delta)}{\varepsilon\sqrt{N}}\right)$. Then it satisfies (ε, δ) -differential privacy with $\varepsilon = \mathcal{O}\left(\sqrt{\frac{\log(N/\delta)}{N}}\right)$.*

Remarks: Note that the privacy budget ε is restricted to be $\varepsilon \leq 1/\sqrt{N}$ because ε has to be less than 1 in each iteration to satisfy technical requirements of Feldman et al. (2021). Also note that the privacy guarantee holds *irrespective* of data distribution \mathcal{D} . Finally, clipping parameter ζ as well as learning rate η requires knowledge of certain key quantities about the dataset and \mathcal{D} . In practice, we would need to estimate these quantities while preserving privacy. There are standard approaches for such hyperparameter tuning (Liu and Talwar, 2019; Papernot and Steinke, 2021). However, being consistent with prior literature on DP optimization, and to highlight the key points of our approach, throughout the paper we assume a priori knowledge of the hyperparameters. Further, the largest and smallest eigenvalue of Hessian can be estimated in a differentially private manner using standard sensitivity analysis with output perturbation, since the sensitivity of eigenvalues is constant if the feature vector have constant norm. Under our sample complexity assumption, the rates remain the same up to constants. Now, we present the error bounds for our method.

Theorem 6 *Let $D = \{(\mathbf{x}_i, y_i)\}_{i=0}^{N-1}$ be sampled i.i.d. with $\mathbf{x}_i \sim \mathcal{D}$ satisfying $(\mathbf{H}, K_2, a, \beta)$ -Tail, and the distribution of the inherent noise z satisfies $(\sigma^2, K_2, a, \beta)$ -Tail with $\beta = \frac{1}{N^{100}}$. Let \mathbf{w}^* be the optimal solution to the population least squares problem, κ be the condition number of the covariance*

matrix \mathbf{H} , and $\mathbb{E} \left[\|\mathbf{x}\|_2^2 \right] \leq R_x^2$ (refer Section 2 for details about the notations and assumptions).

Initialize parameters in DP-SSGD as follows: stepsize $\eta = \min \left\{ \frac{1}{2R_x^2}, \frac{c_1}{\log^{4a+2} N} \cdot \frac{1}{K_2^2 R_x^2 \kappa} \cdot \frac{1}{d\alpha_{\varepsilon, \delta, N}^2}, \frac{c_2}{\log^{2a+2} N} \cdot \frac{1}{R_x^2} \right\}$, where $c_1, c_2 > 0$ are global constants, noise multiplier $\alpha = \alpha_{\varepsilon, \delta, N} = \mathcal{O} \left(\frac{\log(N/\delta)}{\varepsilon\sqrt{N}} \right)$ and clipping threshold $\zeta = 4K_2 R_x \cdot \log^{2a} N \cdot (\sqrt{\|\mathbf{H}\|_2} \|\mathbf{w}^*\| + \sqrt{\kappa}\sigma)$. Then, the output $\bar{\mathbf{w}}$ of DP-SSGD achieves the following excess risk w.p. $\geq 1 - 1/N^{100}$ over randomness in data and algorithm:

$$\mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) \leq \tilde{\mathcal{O}} \left(\frac{\sigma^2 d}{N} + \frac{d^2 \log^2(N/\delta)}{\varepsilon^2 N^2} \cdot K_2^2 \kappa^2 (\|\mathbf{w}^*\|_{\mathbf{H}}^2 + \sigma^2) \right),$$

assuming $N \geq \Omega \left(\frac{\kappa d \log^{2a+1}(\kappa d) \log(1/\delta)}{\varepsilon} \right)$.

Remarks:

i) The error bound consists of two terms. The first term corresponds to the error due to noise in observations, which is minimax optimal up to a factor of 2. The second term corresponds to the error incurred due to the noise added for preserving privacy.

ii) Note that the error term corresponding to privacy budget for existing methods (Cai et al., 2021; Bassily et al., 2019; Wang et al., 2017) is $\approx \frac{d^2 R_x^2 \|\mathbf{w}^*\|^2}{\varepsilon^2 N^2}$. That is, the error term is $O(d)$ worse than our result. Furthermore, as analyzed in Corollary 7, for standard setting where each $\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $y_i = \langle \mathbf{x}_i, \mathbf{w}^* \rangle$, existing results imply error bound of $\frac{d^3}{\varepsilon^2 N^2}$ in contrast to our bound of $\frac{d^2}{\varepsilon^2 N^2}$.

iii) Theorem requires ε to scale as $1/\sqrt{N}$ which is restrictive in practice. Furthermore, under such restriction, the sample complexity of N is $N \geq \kappa^2 d^2$ which is also sub-optimal. Finally, note that the error term is independent of σ and hence even when $\sigma = 0$, the error term is non-trivial.

Corollary 7 Consider the setting of Theorem 6. Let $\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \in \mathbb{R}^d$, i.e., $R_x^2 = d$ and $\kappa = 1$. Let $y_i = \langle \mathbf{x}_i, \mathbf{w}^* \rangle + z_i$ where $z_i \sim \mathcal{N}(0, \sigma^2)$. Then, assuming $N \geq d^2 \log \frac{\varepsilon}{\log(1/\delta)}$, we have:

1. DP-SSGD with parameters $\eta = 1/4d$, $\zeta = \sqrt{d} \sqrt{\|\mathbf{w}^*\|^2 + \sigma^2} \log N$, $\alpha = \frac{\log(N/\delta)}{\varepsilon\sqrt{N}}$ is (ε, δ) -DP where $\varepsilon = \mathcal{O} \left(\sqrt{\frac{\log(1/\delta)}{N}} \right)$.
2. The output $\bar{\mathbf{w}}$ satisfies the following risk bound w.p. $\geq 1 - 1/N^{100}$: $\mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) = \tilde{\mathcal{O}} \left(\frac{d\sigma^2}{N} + \frac{d^2(\|\mathbf{w}^*\|^2 + \sigma^2)}{N^2 \varepsilon^2} \log^2(N/\delta) \right)$.

Below, we provide a proof of Theorem 6. See Appendix A for the supporting lemmas and their detailed proofs.

Proof [Theorem 6] Consider the event: $\mathcal{E} = \{\|\mathbf{x}_t(\langle \mathbf{x}_t, \mathbf{w}_t \rangle - y_t)\| \leq \zeta, \forall 0 \leq t \leq N-1\}$. Note that by Lemma 19, $Pr(\mathcal{E}) \geq 1 - \frac{1}{N^{100}}$. Now, if \mathcal{E} holds, then DP-SSGD does not have any clipping. Thus, under \mathcal{E} , by Lemma 8:

$$\mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) \leq \frac{R_x^2}{2} \exp \left(-\frac{\eta\mu}{2} N \right) \|\mathbf{w}^*\|^2 + \frac{1}{N} \text{Tr}(\mathbf{H}^{-1}(\Sigma + 4\zeta^2 \alpha^2 \mathbf{I})) + \frac{1}{N} \cdot \frac{\eta R_x^2}{1 - \eta R_x^2} d \|\Sigma\| + 4\zeta^2 \alpha^2 \|\mathbf{H}\|, \quad (5)$$

where $\Sigma = \mathbb{E}[z^2 \mathbf{x} \mathbf{x}^\top]$. Note that $\|\Sigma\|_2 \leq \sigma^2 \|\mathbf{H}\|_2$. Using ζ, α, η as defined in the theorem, we have (assuming event \mathcal{E}):

$$\begin{aligned} \mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) &\leq \exp\left(-\frac{\mu c}{2R_x^2} \cdot \frac{\varepsilon^2 N}{K_2^2 \kappa d \log^2(N/\delta) \log^{4a+2} N} \cdot N\right) \frac{d}{2} \|\mathbf{w}^*\|_{\mathbf{H}}^2 + \frac{2}{N} \sigma^2 d \\ &\quad + \frac{128dR_x^2 K_2^2 \kappa \log^{4a} N (\|\mathbf{w}^*\|_{\mathbf{H}}^2 + \sigma^2) \log^2(N/\delta)}{\mu \varepsilon^2 N^2}. \end{aligned} \quad (6)$$

Theorem now follows by using the sample complexity bound on N , the fact that $R_x^2 \leq d \|\mathbf{H}\|_2$, $\kappa = \|\mathbf{H}\|_2 / \mu$ and by combining the above inequality with $\Pr(\mathcal{E}) \geq 1 - 1/N^{100}$. \blacksquare

Lemma 8 *Let $D = \{(\mathbf{x}_i, y_i)\}_{i=0}^{N-1}$ be sampled from distribution \mathcal{D} ; see Section 2 for the notations and assumptions about \mathcal{D} . Let \mathbf{w}^* minimize the population squared loss. Let $\bar{\mathbf{w}}$ be the output of DP-SSGD (Algorithm 1) s.t. there is no gradient clipping at any step and $\eta \leq \frac{1}{2R_x^2}$. Then we have:*

$$\mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) \leq \frac{R_x^2}{2} e^{-\eta \mu (\frac{N}{2} + 1)} \|\mathbf{w}^*\|^2 + \frac{1}{N} \frac{\eta R_x^2}{1 - \eta R_x^2} d \|\Sigma + 4\zeta^2 \alpha^2 \mathbf{I}\|_{\mathbf{H}} + \frac{1}{N} \text{Tr}(\mathbf{H}^{-1}(\Sigma + 4\zeta^2 \alpha^2 \mathbf{I})).$$

See Appendix A for a detailed proof of the above lemma and the other supporting lemmas.

3.2. DP-Adaptive Mini-batch Shuffled SGD

Theorem 5 indicates that DP-SSGD only applies in the setting where ε is small which can be restrictive for practical usage. The key challenge is that shuffling amplification results along with privacy composition over T iterations require per-iteration privacy budget to be less than $1/\sqrt{T}$. As $T = N$ for DP-SSGD, it implies that $\varepsilon = O(1/\sqrt{N})$.

So the key challenge is to be able to reduce the number of iterations despite one-pass shuffling so that we do not have additional d factors. To this end, we process a batch of points in each update of shuffled SGD instead of using a single point per iteration as in DP-SSGD. As mentioned above, another weakness in the result of DP-SSGD is that the error bound does not vanish even for vanishing noise variance $\sigma \rightarrow 0$; recall that σ is the variance of noise in observations $\mathbb{E}[z^2]$. We fix this issue by adaptively estimating the clipping threshold for the gradients. Note that a similar technique was used by Biswas et al. (2020) but in a different context of mean estimation.

See Algorithm 2 for a pseudo code of our method. The algorithm randomly shuffles the data (Step 2) and takes one pass over it. The N data points are divided into $T \approx \kappa \log N$ batches of size $b + s$, where κ is the condition number of $\mathbf{H} = \mathbb{E}[\mathbf{x} \mathbf{x}^\top]$, i.e., $\kappa = \|\mathbf{H}\|_2 / \mu$. Each iteration first estimates the clipping threshold ζ_t while preserving differential privacy. Ideally we require ζ_t to be large enough so that it does not lead to any thresholding (with high probability). The gradient of a given point (\mathbf{x}, y) is $\mathbf{x}(\langle \mathbf{x}, \mathbf{w}_t \rangle - y)$. Furthermore, by using the sub-Gaussian style assumptions on \mathcal{D} (see (3), (2)) we have $(\langle \mathbf{x}, \mathbf{w}_t \rangle - y)^2 \lesssim \mathbb{E}[(\langle \mathbf{x}, \mathbf{w}_t \rangle - y)^2] \cdot \log^a N$ if \mathbf{x}, y are independent of \mathbf{w}_t . If $\mathbf{w}_t \approx \mathbf{w}^*$, then the above quantity decreases to around σ^2 which is desired to get stronger bounds.

To estimate ζ_t , we use DP-STAT method (see Algorithm 3) which employs a standard technique from the private statistics literature to approximately estimate the highest value in the dataset. At a high level, DP-STAT (Algorithm 3) exploits the fact that a single point cannot significantly perturb say 99th percentile of a given batch of points, and adds noise appropriately.

Algorithm 2 DP-AMBSSGD: DP-Adaptive-Mini-Batch-Shuffled-SGD

- 1: **Input:** Samples: $\{(\mathbf{x}_i, y_i)\}_{i=0}^{N-1}$, Learning Rate: η , DP Noise Multiplier: α , Batch Size: b , Expected \mathbf{x} Norm: R_x , Domain Size: B , Stat Sample Size: s
 - 2: Randomly permute $\{(\mathbf{x}_i, y_i)\}_{i=0}^{N-1}$
 - 3: $T = N/(b + s)$
 - 4: $\mathbf{w}_0 \leftarrow \mathbf{0}$
 - 5: **for** $t = 0 \dots T - 1$ **do**
 - 6: $\tau \leftarrow (b + s)t$
 - 7: $\gamma_t \leftarrow \text{DP-STAT}(\{(\mathbf{x}_i, y_i)\}_{i=\tau}^{\tau+s}, \mathbf{w}_t, B, s, \alpha, B/N^{200})$
 - 8: $\zeta_t \leftarrow R_x \cdot \gamma_t \cdot \log^a N$
 - 9: Sample $\mathbf{g}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{d \times d})$
 - 10: $\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t - \eta \left(\frac{1}{b} \sum_{i=0}^{b-1} \text{clip}_{\zeta_t}(\mathbf{x}_{\tau+s+i} \langle \mathbf{x}_{\tau+s+i}, \mathbf{w}_t \rangle - y_{\tau+s+i}) + \alpha \cdot \frac{2\zeta_t}{b} \cdot \mathbf{g}_t \right)$ where
 $\text{clip}_{\zeta}(\boldsymbol{\nu}) = \boldsymbol{\nu} \cdot \min \left\{ 1, \frac{\zeta}{\|\boldsymbol{\nu}\|_2} \right\}$
 - 11: **end for**
 - 12: **Result:** $\bar{\mathbf{w}} := \frac{2}{T} \sum_{t'=T/2+1}^T \mathbf{w}_{t'}$
-

Algorithm 3 DP-STAT: Private Estimation of Approximately Maximum

- 1: **Input:** Samples: $\{(\mathbf{x}_i, y_i)\}_{i=0}^s$, Parameter: \mathbf{w} , Domain Size: B , Stat Sample Size: s , Noise Multiplier: α , Discretization Width: Δ
 - 2: $\gamma_0 \leftarrow \Delta$.
 - 3: **for** $i \in \{0, \dots, \lceil \log_2(B/\Delta) \rceil\}$ **do**
 - 4: $c \leftarrow |\{(\mathbf{x}_j, \mathbf{w}) - y_j \leq \gamma_i : j \in \{0, \dots, s\}\}|$. $c_{\text{priv}} \leftarrow c + \mathcal{N}(0, \lceil \log_2(B/\Delta) \rceil \alpha^2)$.
 - 5: **if** $c_{\text{priv}} < s$ **then** $\gamma_{i+1} \leftarrow 2 \cdot \gamma_i$, **else break.**
 - 6: **end for**
 - 7: **return** $\gamma_{\text{priv}} \leftarrow \gamma_i$.
-

After estimating ζ_t (Step 7, Algorithm 2), we select the next batch of points and apply the standard mini-batch SGD update but with added Gaussian noise of standard deviation $2\eta\zeta_t\alpha/b$ where α is the standard noise multiplier to ensure differential privacy (Step 10). Finally, we output $\bar{\mathbf{w}}$ which is the average of the last $T/2$ iterates (Step 12).

Theorem 9 *Algorithm DP – AMBSSGD with noise multiplier α satisfies $\frac{1}{\alpha^2}$ -zCDP, and correspondingly satisfies (ε, δ) -differential privacy when we set the noise multiplier $\alpha \geq \frac{2\sqrt{\log(1/\delta)+\varepsilon}}{\varepsilon}$. Furthermore, if $\varepsilon \leq \log(1/\delta)$, then $\alpha \geq \frac{\sqrt{8\log(1/\delta)}}{\varepsilon}$ suffices to ensure (ε, δ) -differential privacy.*

Remarks: The privacy proof follows by combining privacy analysis of clipped SGD with that of approximate maximum estimation (DP-STAT). Since each data-sample is observed in only one mini-batch and the algorithm DP-AMBSSGD takes only a single pass over the entire data (and by implication only a single pass each data point), we use parallel composition property of DP to avoid the \sqrt{T} factor which would have manifested when employing sequential composition. Further, our algorithm and result holds for any ε , in contrast with the $\varepsilon = O(1/\sqrt{N})$ requirement of DP-SSGD. We now present the excess risk bound for DP-AMBSSGD.

Theorem 10 Let $D = \{(\mathbf{x}_i, y_i)\}_{i=0}^{N-1}$ be sampled i.i.d. with $\mathbf{x}_i \sim \mathcal{D}$ satisfying $(\mathbf{H}, K_2, a, \beta)$ -Tail, and the distribution of the inherent noise z_i satisfies $(\sigma^2, K_2, a, \beta)$ -Tail with $\beta = \frac{1}{N^{200}}$. Let \mathbf{w}^* be the optimal solution to the population least squares problem, κ be the condition number of the covariance matrix \mathbf{H} , and $\mathbb{E} \left[\|\mathbf{x}\|_2^2 \right] \leq R_x^2$ (see Section 2 for notations and assumptions).

Initialize parameters in DP-AMBSSGD as follows: batch size $b = \frac{N}{T} - s$, stat sample size $s = \frac{b}{10}$, stepsize $\eta = \frac{b}{R_x^2 + (b-1)\|\mathbf{H}\|_2}$, number of iterations $T = c_1 \kappa \log(N)$ where $c_1 > 0$ is a global constant, domain size $B = K_2 R_x (\|\mathbf{w}^*\|_{\mathbf{H}} + \sigma) \log^{2a} N$, and noise multiplier $\alpha = \frac{\sqrt{8 \log(1/\delta)}}{\varepsilon}$. Then, the output $\bar{\mathbf{w}}$ achieves the following excess risk w.p. $\geq 1 - 1/N^{100}$ over the randomness in data and algorithm:

$$\mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) \leq \frac{\|\mathbf{w}^*\|_{\mathbf{H}}^2}{N^{100}} + \frac{8\sigma^2 d}{N} + \tilde{O} \left(\frac{\sigma^2 \kappa^2 d^2 \log(1/\delta)}{\varepsilon^2 N^2} \right),$$

assuming $N \geq \tilde{\Omega} \left(\kappa^2 d \left(1 + \frac{\sqrt{\log(1/\delta)}}{\varepsilon} \right) \right)$.

Remarks: i) Note that the excess risk bound has three terms. The first term depends on $\|\mathbf{w}^*\|_{\mathbf{H}}$ but is polynomially small in N . The second term matches the information theoretically optimum rate (up to constant factor) for non-private linear regression. Finally, the third term is approximately $\sigma^2 d^2 / (\varepsilon N)^2$.

ii) Up to the first term which is $1/N^{100}$ and up to $\kappa, \log N$ factors, the bounds match the lower bound of $\sigma^2 d/N + \sigma^2 d^2 / (\varepsilon^2 N^2)$ by Cai et al. (2021). Note that the first term is present as DP-AMBSSGD can take only pass over the data. However, by using DP-SGD with initialization $\mathbf{w}_0 \leftarrow \bar{\mathbf{w}}$, and running the algorithm for $\log(N/\sigma)$ iterations, we can get rid of the first term. One key observation here is that, as the error term dependent on $\|\mathbf{w}^*\|_{\mathbf{H}}$ is already small, the *additional* d term introduced by DP-SGD is inconsequential to the overall bound. As the DP-SGD part is relatively standard, we do not present the above mentioned argument in this paper.

iii) The result holds for nearly optimal sample complexity of $N = \tilde{\Omega}(d)$, ignoring κ and $\log N$ factors. This matches the sample complexity bound for *non-private* linear regression as well. Furthermore, time complexity of our method is $O(Nd)$ which is linear in the input size.

iv) The leading term in the error bound corresponding to the MLE error of $\sigma^2 d/N$ is independent of κ and $\log N$ factors. But the last term as well as sample complexity depends on κ . The exact dependence can be optimized by more careful analysis. Furthermore, using a slightly more complicated algorithm that forms multiple batch sizes and tail averaging in each phase can perhaps further reduce dependence on κ . However, it is not clear if the dependence can be removed completely; we leave further investigation into dependence on κ for future work.

v) Consider the standard setting of linear regression with $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $z \sim \sigma \mathcal{N}(0, 1)$. Here, our bound is given by the below corollary which is nearly optimal up to the first term and $\log N$ factors. Our second term is smaller than existing results for the same setting by $O(d)$ and a σ^2 factor.

vi) We want to clarify that the DP guarantees should hold over worst-case data sets whereas the high-probability that gradient clipping does not happen is over the randomness of the data. In our privacy analysis, one cannot exploit this randomness and get rid of the clipping. However, if clipping indeed happens, one can exactly quantify the clipped loss as a Huberized version of the least squared error loss on individual data points. (For a detailed discussion, See (Song et al., 2020, Section 5.1))

Corollary 11 Consider the setting of Theorem 10. Let $\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \in \mathbb{R}^d$, i.e., $R_x^2 = d$ and $\kappa = 1$. Let $y_i = \langle \mathbf{x}_i, \mathbf{w}^* \rangle + z_i$ where $z_i \sim \mathcal{N}(0, \sigma^2)$. Then, assuming $N \geq \tilde{\Omega} \left(d \left(1 + \frac{\sqrt{\log(1/\delta)}}{\varepsilon} \right) \right)$, we have:

1. Algorithm DP-AMBSSGD with parameters $\eta = 1/4d$, $\alpha = \frac{\sqrt{8 \log(1/\delta)}}{\varepsilon}$ is (ε, δ) -DP.
2. The output $\bar{\mathbf{w}}$ satisfies the following risk bound:

$$\mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) \leq \frac{\|\mathbf{w}^*\|_{\mathbf{H}}^2}{N^{100}} + \frac{8d\sigma^2}{N} \left(1 + \tilde{\mathcal{O}} \left(\frac{d \log(1/\delta)}{N\varepsilon^2} \right) \right).$$

Proof [Theorem 10] Let $\Gamma = \Omega \left(\frac{\sqrt{\log(1/\delta)}}{\varepsilon} \cdot \log N \right)$. Then, using Lemma 12, with probability $\geq 1 - 1/N^{200}$, at least $s - \Gamma$ points satisfies $|\langle \mathbf{x}_i, \mathbf{w}_t \rangle - y_i| \leq \gamma_t$. Hence, using $s = b/10$ and sample complexity assumption, we have: $\kappa \|\mathbf{w}^* - \mathbf{w}_t\|_{\mathbf{H}}^2 + \sigma^2 \leq \gamma_t^2$ with probability $\geq 1 - 1/N^{200}$. Hence, by setting $\zeta_t = R_x \gamma_t \log^a N$ and using assumptions (2), (3), along with the fact that \mathbf{w}_t is independent of $\mathbf{x}_{\tau+s+i}$ for $0 \leq i \leq b - 1$, we have (w.p. $\geq 1 - 1/N^{200}$):

$$\|\mathbf{x}_{\tau+s+i}(\langle \mathbf{x}_{\tau+s+i}, \mathbf{w}_t \rangle - y_{\tau+s+i})\| \leq \zeta_t,$$

for all iterations t . That is, with probability $\geq 1 - 1/N^{199}$, the following event holds: $\mathcal{E} := \{\text{thresholding is not required for any point}\}$. Now, using the second part of the utility lemma, $s = b/10$, and the sample complexity as above, we also know that $\gamma_t \leq K_2 \log^a N \left(\sqrt{\kappa} \|\mathbf{w}_t - \mathbf{w}^*\|_{\mathbf{H}} + \sigma + \Delta \right)$ and $\Delta = \frac{\|\mathbf{w}^*\|_{\mathbf{H}} + \sigma}{N^{100}}$. Hence, the requirement for ζ_t given in Lemma 13 is satisfied with probability $\geq 1 - 1/N^{199}$.

That is, with probability $\geq 1 - 2/N^{199}$, both the above mentioned events hold, i.e., thresholding is not required for any point, and ζ_t requirement is satisfied for each iteration. Theorem now follows by applying Lemma 13 conditioned over the above two high probability events. The final expression in the theorem is obtained by applying sample complexity bound assumption along with fact that $R_x^2 \leq d \|\mathbf{H}\|_2$ and $\kappa = \|\mathbf{H}\|_2 / \mu$. ■

Lemma 12 (DP-STAT) In the following we provide the privacy and utility guarantees:

(Privacy) Algorithm DP-STAT satisfies $\frac{1}{2\alpha^2}$ -zCDP.

(Utility) Let $\Gamma = \alpha \sqrt{2 \log(B/\Delta) \log(\log(B/\Delta)/\beta)}$. Now, w.p. at least $1 - \beta$, Algorithm DP-STAT outputs γ_{priv} s.t. $|\{|\langle \mathbf{x}_i, \mathbf{w} \rangle - y_i| \leq \gamma_{\text{priv}} : i \in \{0, \dots, s\}\}| \geq s - \Gamma$, and $|\{|\langle \mathbf{x}_i, \mathbf{w} \rangle - y_i| \leq \max\{\frac{\gamma_{\text{priv}}}{2}, \Delta\} : i \in \{0, \dots, s\}\}| < s - \Gamma$.

Lemma 13 Let $D = \{(\mathbf{x}_i, y_i)\}_{i=0}^{N-1}$ be sampled from distribution \mathcal{D} ; see Section 2 for the notations and assumptions about \mathcal{D} . Let \mathbf{w}^* be the optimal solution to the population least squares problem. Let $\bar{\mathbf{w}}$ be the output of DP-AMBSSGD (Algorithm 2) s.t. there is no gradient clipping at any step, $\eta \leq \frac{b}{R_x^2 + (b-1)\|\mathbf{H}\|}$, batch size b , total number of iterations T s.t. $T \cdot (b + s) = N$ and $T = c_1 \kappa \log N$

where $c_1 > 0$ is a global constant and $\left(\frac{N}{T} - s\right)^2 \geq \frac{24\eta\alpha^2 R_x^2 K_2^2 \kappa \log^{4a} N \text{Tr}(\mathbf{H})}{\mu}$. Then we have:

$$\begin{aligned} \mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) &\leq \left(\frac{4}{\eta^2 \mu^2 T^2} e^{-\eta\mu} + \frac{96\alpha^2}{Tb^2} K_2^2 R_x^2 \kappa \log^{4a} N \text{Tr}(\mathbf{H}^{-1}) \right) \\ &\quad \left(\frac{\|\mathbf{w}^*\|_{\mathbf{H}}^2}{N^{\frac{\eta\mu}{4} c_1 \kappa}} + \frac{2\eta}{\mu b} \text{Tr}(\mathbf{H}\Sigma) + \frac{24\eta\alpha^2}{\mu b^2} K_2^2 R_x^2 \log^{4a} N (\sigma^2 + \Delta^2) \text{Tr}(\mathbf{H}) \right) \\ &\quad + \frac{8}{Tb} \left(\text{Tr}(\mathbf{H}^{-1}\Sigma) + \frac{12\alpha^2}{b} K_2^2 R_x^2 \log^{4a} N (\sigma^2 + \Delta^2) \text{Tr}(\mathbf{H}^{-1}) \right), \end{aligned}$$

where $\zeta_t \leq R_x K_2 \log^{2a} N \left(\sqrt{\kappa} \|\mathbf{w}_t - \mathbf{w}^*\|_{\mathbf{H}} + \sigma + \Delta \right)$ and $\Delta = \frac{\|\mathbf{w}^*\|_{\mathbf{H} + \sigma}}{N^{100}}$.

See Appendix B for a detailed proof of the above lemma and the other supporting lemmas.

4. Conclusions

We studied the problem of differentially private linear regression, and proposed a novel algorithm DP-AMBSSGD based on mini-batch shuffled SGD with adaptive clipping. The method can guarantee nearly optimal error rate in terms of d , N , σ , and ε , for sub-Gaussian style distributions, a significant improvement over existing risk bounds for polynomial time methods. However, our results depend critically on the condition number κ of Hessian $\mathbb{E}[\mathbf{xx}^T]$. An investigation of both the lower bound as well as the upper bound w.r.t. κ is an important future direction. Furthermore, removing log dependencies in the analysis is interesting. Finally, generalizing our techniques to obtain optimal error rates for generalized linear models like logistic regression should be of wide interest.

References

- Martin 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.
- Galen Andrew, Om Thakkar, Brendan McMahan, and Swaroop Ramaswamy. Differentially private learning with adaptive clipping. *Advances in Neural Information Processing Systems*, 34, 2021.
- 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. *CoRR*, abs/1908.09970, 2019. URL <http://arxiv.org/abs/1908.09970>.
- Raef Bassily, Vitaly Feldman, Cristóbal Guzmán, and Kunal Talwar. Stability of stochastic gradient descent on nonsmooth convex losses. *Advances in Neural Information Processing Systems*, 33, 2020.
- Sourav Biswas, Yihe Dong, Gautam Kamath, and Jonathan Ullman. Coinpress: Practical private mean and covariance estimation, 2020.
- Gavin Brown, Marco Gaboardi, Adam Smith, Jonathan Ullman, and Lydia Zakyntinou. Covariance-aware private mean estimation without private covariance estimation. *Advances in Neural Information Processing Systems*, 34, 2021.
- Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016.
- T. Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *The Annals of Statistics*, 49(5):2825 – 2850, 2021. doi: 10.1214/21-AOS2058. URL <https://doi.org/10.1214/21-AOS2058>.
- Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(Mar):1069–1109, 2011.
- Steve Chien, Prateek Jain, Walid Krichene, Steffen Rendle, Shuang Song, Abhradeep Thakurta, and Li Zhang. Private alternating least squares: Practical private matrix completion with tighter rates. In *International Conference on Machine Learning*, pages 1877–1887. PMLR, 2021.
- Aymeric Dieuleveut, Nicolas Flammarion, and Francis Bach. Harder, better, faster, stronger convergence rates for least-squares regression, 2016.
- 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.
- 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.
- Vitaly Feldman, Audra McMillan, and Kunal Talwar. Hiding among the clones: A simple and nearly optimal analysis of privacy amplification by shuffling, 2021.
- Daniel Hsu, Sham M. Kakade, and Tong Zhang. Random design analysis of ridge regression. In Shie Mannor, Nathan Srebro, and Robert C. Williamson, editors, *Proceedings of the 25th Annual Conference on Learning Theory*, volume 23 of *Proceedings of Machine Learning Research*, pages 9.1–9.24, Edinburgh, Scotland, 25–27 Jun 2012. PMLR.
- 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.
- Prateek Jain, Sham M. Kakade, Rahul Kidambi, Praneeth Netrapalli, Krishna Pillutla, and Aaron Sidford. A markov chain theory approach to characterizing the minimax optimality of stochastic gradient descent (for least squares). In *FSTTCS*, 2017.
- Prateek Jain, Sham M. Kakade, Rahul Kidambi, Praneeth Netrapalli, and Aaron Sidford. Parallelizing stochastic gradient descent for least squares regression: mini-batching, averaging, and model misspecification, 2018.
- Gautam Kamath, Jerry Li, Vikrant Singhal, and Jonathan R. Ullman. Privately learning high-dimensional distributions. *CoRR*, abs/1805.00216, 2018. URL <http://arxiv.org/abs/1805.00216>.
- 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.
- Jingcheng Liu and Kunal Talwar. Private selection from private candidates. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 298–309, 2019.
- Xiyang Liu, Weihao Kong, and Sewoong Oh. Differential privacy and robust statistics in high dimensions. *arXiv preprint arXiv:2111.06578*, 2021.
- H Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. *arXiv preprint arXiv:1710.06963*, 2017.
- Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.
- Nicolas Papernot and Thomas Steinke. Hyperparameter tuning with renyi differential privacy. *arXiv preprint arXiv:2110.03620*, 2021.

- Venkatadheeraj Pichapati, Ananda Theertha Suresh, Felix X Yu, Sashank J Reddi, and Sanjiv Kumar. Adacclip: Adaptive clipping for private sgd. *arXiv preprint arXiv:1908.07643*, 2019.
- Mark Rudelson and Roman Vershynin. Hanson-wright inequality and sub-gaussian concentration, 2013.
- Or Sheffet. Old techniques in differentially private linear regression. In *Algorithmic Learning Theory*, pages 789–827. PMLR, 2019.
- 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.
- 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.
- Tianhao Wang, Jeremiah Blocki, Ninghui Li, and Somesh Jha. Locally differentially private protocols for frequency estimation. In *26th USENIX Security Symp.*, pages 729–745, 2017.
- Xi Wu, Fengan 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.
- Úlfar Erlingsson, Ilya Mironov, Ananth Raghunathan, and Shuang Song. That which we call private, 2019.

Appendix A. Missing Proofs from Section 3.1

A.1. Proof of Privacy Guarantees

A.1.1. AUXILIARY LEMMAS

Here we present a few auxiliary results which will be used in proving our DP-SSGD guarantees.

Lemma 14 *Each update step of DP-SSGD (Algorithm 1) is $(\varepsilon_0, \delta_0)$ private if $\alpha = \frac{c}{\varepsilon_0}$ where $c \geq \sqrt{2 \log(1.25/\delta_0)}$ and L is the clipping norm.*

Proof Each update step (excluding the DP-noise addition) is of the form:

$$\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t - \eta \text{clip}_\zeta(\mathbf{x}_t(\langle \mathbf{x}_t, \mathbf{w}_t \rangle - y_t)),$$

where $\text{clip}_\zeta(\boldsymbol{\nu}) = \boldsymbol{\nu} \cdot \max\left\{1, \frac{\zeta}{\|\boldsymbol{\nu}\|_2}\right\}$.

Therefore, the local L_2 sensitivity of the \mathbf{w}_{t+1} can be computed by considering a difference in the t^{th} iteration data sample as follows:

$$\begin{aligned} \Delta_2 &= \|\mathbf{w}'_{t+1} - \mathbf{w}_{t+1}\| \\ &= \|\eta \text{clip}_\zeta(\mathbf{x}'_t(\langle \mathbf{x}'_t, \mathbf{w}_t \rangle - y'_t)) - \eta \text{clip}_\zeta(\mathbf{x}_t(\langle \mathbf{x}_t, \mathbf{w}_t \rangle - y_t))\| \\ &\leq 2\eta \|\text{clip}_\zeta(\mathbf{x}_t(\langle \mathbf{x}_t, \mathbf{w}_t \rangle - y_t))\| \\ &= 2\eta\zeta. \end{aligned}$$

Using Gaussian mechanism to induce Local Differential Privacy with the above L_2 sensitivity proves the Lemma. \blacksquare

Theorem 15 (From Feldman et al. (2021)) *For a domain \mathcal{D} , let $\mathcal{R}^{(i)} : \mathcal{f} \times \mathcal{D} \rightarrow \mathcal{S}^{(i)}$ for $i \in [n]$ (where $\mathcal{S}^{(i)}$ is the range space of $\mathcal{R}^{(i)}$) be a sequence of algorithms such that $\mathcal{R}^{(i)}(z_{1:i-1}, \cdot)$ is a $(\varepsilon_0, \delta_0)$ -DP local randomizer for all values of auxiliary inputs $z_{1:i-1} \in \mathcal{S}^{(1)} \times \dots \times \mathcal{S}^{(i-1)}$. Let $\mathcal{A}_s : \mathcal{D}^n \rightarrow \mathcal{S}^{(1)} \times \dots \times \mathcal{S}^{(n)}$ be the algorithm that given a dataset $x_{1:n} \in \mathcal{D}^n$, samples a uniformly random permutation π , then sequentially computes $z_i = \mathcal{R}^{(i)}(z_{1:i-1}, x_{\pi(i)})$ for $i \in [n]$, and outputs $z_{1:n}$. Then for any $\delta \in [0, 1]$ such that $\varepsilon_0 \leq \log\left(\frac{n}{16 \log(2/\delta)}\right)$, \mathcal{A}_s is $(\varepsilon, \delta + \mathcal{O}(e^\varepsilon \delta_0 n))$ -DP where ε is:*

$$\varepsilon = \mathcal{O}\left((1 - e^{-\varepsilon_0})\left(\frac{\sqrt{e^{\varepsilon_0} \log(1/\delta)}}{\sqrt{n}} + \frac{e^{\varepsilon_0}}{n}\right)\right).$$

A.1.2. PROOF OF THEOREM 5

Proof We will use Theorem 15 to obtain the stated result. Denoting auxiliary inputs using u , we can rewrite the update steps of DP-SSGD as a sequence of one step algorithms:

$$\mathcal{R}^{(t+1)}(u_{0:t}, (\mathbf{x}, y)) := \mathbf{w}_{t+1} \leftarrow \mathbf{w}_t(u_{0:t}) - \eta \text{clip}_\zeta(\mathbf{x}_{\pi(t)}(\langle \mathbf{x}_{\pi(t)}, \mathbf{w}_t(u_{0:t}) \rangle - y_{\pi(t)})) - 2\eta\zeta\alpha g_t,$$

where $\pi(t)$ denotes the t^{th} iteration sample after randomly permuting the input data.

From Lemma 14, each $\mathcal{R}^{(t+1)}(u_{0:t}, \cdot)$ is a $(\varepsilon_0, \delta_0)$ -DP local randomizer algorithm with $\varepsilon_0 \leq \log\left(\frac{N}{16 \log(2/\widehat{\delta})}\right)$. The output of DP-SSGD is obtained by post-processing of the shuffled outputs $u_{t+1} = \mathcal{R}^{(t+1)}(u_{0:t}, (\mathbf{x}, y))$ for $t \in \{0, \dots, N-1\}$.

Therefore, Theorem 15 implies that Algorithm DP-SSGD is $(\widehat{\varepsilon}, \widehat{\delta} + \mathcal{O}(e^{\widehat{\varepsilon}} \delta_0 N))$ -DP such that:

$$\widehat{\varepsilon} = \mathcal{O}\left((1 - e^{-\varepsilon_0})\left(\frac{\sqrt{e^{\varepsilon_0} \log(1/\widehat{\delta})}}{\sqrt{N}} + \frac{e^{\varepsilon_0}}{N}\right)\right)$$

Now, assume that $\varepsilon_0 \leq \frac{1}{2}$. The above implies $\exists c_1 > 0$ s.t.

$$\begin{aligned} \widehat{\varepsilon} &\leq c_1 \cdot (1 - e^{-\varepsilon_0}) \left(\frac{\sqrt{e^{\varepsilon_0} \log(1/\widehat{\delta})}}{\sqrt{N}} + \frac{e^{\varepsilon_0}}{N} \right) \\ &\leq c_1 \cdot \left((e^{\varepsilon_0/2} - e^{-\varepsilon_0/2}) \sqrt{\frac{\log(1/\widehat{\delta})}{N}} + (e^{\varepsilon_0} - 1) \frac{1}{N} \right) \\ &\leq c_1 \cdot \left(((1 + \varepsilon_0) - (1 - \varepsilon_0/2)) \sqrt{\frac{\log(1/\widehat{\delta})}{N}} + ((1 + 2\varepsilon_0) - 1) \frac{1}{N} \right) \\ &= c_1 \cdot \varepsilon_0 \left(\frac{1}{2} \sqrt{\frac{\log(1/\widehat{\delta})}{N}} + \frac{2}{N} \right). \end{aligned} \quad (7)$$

From Lemma 14, we set $\alpha = \frac{\sqrt{2 \log(1.25/\delta_0)}}{\varepsilon_0}$, then each update step of DP-SSGD independently satisfies $(\varepsilon_0, \delta_0)$ -DP. This follows from standard Gaussian mechanism arguments (Mironov, 2017; Dwork et al., 2006a). In (7) we replace $\varepsilon_0 = \frac{\sqrt{2 \log(1.25/\delta_0)}}{\alpha}$ to obtain the following:

$$\widehat{\varepsilon} \leq c_1 \cdot \frac{\sqrt{2 \log(1.25/\delta_0)}}{\alpha} \cdot \left(\frac{1}{2} \sqrt{\frac{\log(1/\widehat{\delta})}{N}} + \frac{2}{N} \right) \leq c_1 \cdot \frac{\sqrt{2 \log(1.25/\delta_0) \log(1/\widehat{\delta})}}{\alpha \sqrt{N}}. \quad (8)$$

Now to ensure that we satisfy overall (ε, δ) -DP, we will set $\widehat{\delta} = \frac{\delta}{2}$, and $\delta_0 = c_2 \cdot \frac{\delta}{e^{\widehat{\varepsilon}} N}$ for some universal constant $c_2 > 0$. From (8) we have the following:

$$\widehat{\varepsilon} \leq c_1 \cdot \frac{\sqrt{2 \log(c_2 \cdot 1.25 \cdot e^{\widehat{\varepsilon}} N / \delta) \cdot \log(2/\delta)}}{\alpha \sqrt{N}} \quad (9)$$

For any $\varepsilon \leq 1$, if we set $\alpha = c_3 \cdot \frac{\log(N/\delta)}{\varepsilon \sqrt{N}} \geq c_3 \frac{\sqrt{\log(N/\delta) \log(1/\delta)}}{\varepsilon \sqrt{N}}$ for sufficiently large $c_3 > 0$, then from (9) we have $\widehat{\varepsilon} \leq \varepsilon$. Furthermore, we need $\varepsilon_0 = \frac{\sqrt{2 \log(1.25/\delta_0)}}{\alpha} < \frac{1}{2}$ due to the assumption in Theorem 15, which will be ensured if we set $\varepsilon = \mathcal{O}\left(\sqrt{\frac{\log(N/\delta)}{N}}\right)$.

This immediately implies for $\alpha = \Omega\left(\frac{\log(N/\delta)}{\varepsilon \sqrt{N}}\right)$, DP-SSGD satisfies (ε, δ) -DP as long as $\varepsilon = \mathcal{O}\left(\sqrt{\frac{\log(N/\delta)}{N}}\right)$, which completes the proof. ■

A.2. Proof of Risk Bounds

A.2.1. AUXILIARY TECHNICAL LEMMAS

Here we present a few auxiliary results which will be used in proving our DP-SSGD risk bounds.

Lemma 16 (From Rudelson and Vershynin (2013)) *Hanson-Wright Inequality: For any $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \Sigma)$, the following holds for $t \geq 0$:*

$$\mathbb{P}(\|\mathbf{X}\|^2 \geq \text{Tr}(\Sigma) + 2\sqrt{t}\|\Sigma\|_F + 2t\|\Sigma\|_{\text{op}}) \leq e^{-t}.$$

Lemma 17 *Let $\mathbf{B}_t := \mathbf{I} - \eta \mathbf{x}_t \mathbf{x}_t^\top$ s.t. each $\mathbf{x}_t \in \mathbb{R}^d \forall nt \in \{j, \dots, T-1\}$ has been sampled i.i.d. from \mathcal{D} . Let $\mathbf{z} \in \mathbb{R}^d$ be a vector independent of all \mathbf{B}_t 's. Then for $\beta > 0$ and $\eta < \frac{1}{R_x^2}$, we have with probability $\geq 1 - \beta$:*

$$\|\mathbf{B}_{T-1} \mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{z}\|^2 \leq \frac{1}{\beta} e^{-\eta \mu (T-j)} \|\mathbf{z}\|^2.$$

Proof Note that

$$\begin{aligned} & \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{B}_{T-1} \mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{z}\|^2] \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [(\mathbf{B}_{T-1} \mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{z})^\top \mathbf{B}_{T-1} \mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{z}] \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\mathbf{z}^\top \mathbf{B}_j^\top \dots \mathbf{B}_{T-2}^\top \mathbf{B}_{T-1}^\top \mathbf{B}_{T-1} \mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{z}] \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\mathbf{z}^\top \mathbf{B}_j^\top \dots \mathbf{B}_{T-2}^\top \mathbb{E}_{(\mathbf{x}_{T-1}, y_{T-1}) \sim \mathcal{D}} [\mathbf{B}_{T-1}^\top \mathbf{B}_{T-1}] \mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{z}] \\ &\leq \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\mathbf{z}^\top \mathbf{B}_j^\top \dots \mathbf{B}_{T-2}^\top (\mathbf{I} - 2\eta \mathbf{H} + \eta^2 R_x^2 \mathbf{H}) \mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{z}] \\ &\leq \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\lambda_{\max}(\mathbf{I} - 2\eta \mathbf{H} + \eta^2 R_x^2 \mathbf{H}) \|\mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{z}\|^2] \\ &= \lambda_{\max}(\mathbf{I} - 2\eta \mathbf{H} + \eta^2 R_x^2 \mathbf{H}) \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{z}\|^2] \\ &\leq (1 - \eta(2 - \eta R_x^2) \mu) \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{z}\|^2] \\ &\leq (1 - \eta \mu) \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{z}\|^2] \tag{10} \\ &\leq \dots \\ &\leq (1 - \eta \mu)^{(T-j)} \|\mathbf{z}\|^2 \\ &\leq e^{-\eta \mu (T-j)} \|\mathbf{z}\|^2, \end{aligned}$$

where in (10), we have used the fact that $\eta < \frac{1}{R_x^2}$.

Using Markov Inequality, for $\beta > 0$:

$$\begin{aligned} & \Pr\left\{\mathbf{z} \geq \frac{\mathbb{E}[\mathbf{Z}]}{\beta}\right\} \leq \beta \\ \implies & \Pr\left\{\|\mathbf{B}_{T-1} \dots \mathbf{B}_j \mathbf{z}\|^2 \leq \frac{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{B}_{T-1} \dots \mathbf{B}_j \mathbf{z}\|^2]}{\beta}\right\} \geq 1 - \beta. \end{aligned}$$

Therefore with probability at least $1 - \beta$:

$$\|\mathbf{B}_{T-1} \dots \mathbf{B}_0 \mathbf{z}\|^2 \leq \frac{1}{\beta} e^{-\eta \mu (T-j)} \|\mathbf{z}\|^2.$$

■

Lemma 18 (From Jain et al. (2017)) *If the SGD update step is given by:*

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta(\mathbf{x}_t(\langle \mathbf{x}_t, \mathbf{w}_t \rangle - y_t)),$$

then,

$$\frac{1}{2} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\bar{\mathbf{w}}_{t:T} - \mathbf{w}^*\|_{\mathbf{H}}^2 | \mathbf{w}_0 = \mathbf{w}^*] \leq \frac{\text{Tr}(\mathbf{C}_\infty)}{\eta(T-t+1)},$$

where

$$\begin{aligned} \bar{\mathbf{w}}_{t:T} &:= \frac{1}{T-t+1} \sum_{t'=t}^T \mathbf{w}_{t'}, \\ \text{Tr}(\mathbf{C}_\infty) &\leq \frac{1}{2} \frac{\eta^2 R_x^2}{1 - \eta R_x^2} d \|\Sigma\|_{\mathbf{H}} + \frac{\eta}{2} \text{Tr}(\mathbf{H}^{-1} \Sigma). \end{aligned}$$

Lemma 19

Let η be such that $\eta \leq \min \left\{ \frac{c_1}{\log^{4a+2} N} \cdot \frac{1}{K_2^2 R_x^2 \kappa} \cdot \frac{1}{d \alpha_{\varepsilon, \delta, N}^2}, \frac{c_2}{\log^{2a+2} N} \cdot \frac{1}{R_x^2} \right\}$, where $c_1, c_2 > 0$ are global constants and $\zeta = 4K_2 R_x \cdot \log^{2a} N \cdot \left(\sqrt{\|\mathbf{H}\|_2} \|\mathbf{w}^*\| + \sqrt{\kappa} \sigma \right)$. Furthermore, let $\alpha = \alpha_{\varepsilon, \delta, N}$ be a function of ε, δ, N . Then, with probability $\geq 1 - \frac{1}{N^{100}}$, $\|\mathbf{x}_t(\langle \mathbf{x}_t, \mathbf{w}_t \rangle - y_t)\| \leq \zeta$ for all $0 \leq t \leq N-1$; \mathbf{w}_t is the t^{th} iterate of Algorithm DP-SSGD.

Proof We will prove the Lemma using Principle of Complete Induction.

Base Case: Check for $t = 0$.

The norm of the gradient is given by:

$$\begin{aligned} \|\mathbf{x}_0(\langle \mathbf{x}_0, \mathbf{w}_0 \rangle - y_0)\| &= \|\mathbf{x}_0(\langle \mathbf{x}_0, \mathbf{0} \rangle - y_0)\| \\ &= \|\mathbf{x}_0 y_0\| \\ &\leq \|\mathbf{x}_0\| |y_0| \\ &\leq R_x \log^a(1/\beta_x) \left(K_2 \sqrt{\|\mathbf{H}\|_2} \|\mathbf{w}^*\| \log^a(1/\beta_{\mathbf{w}^*}) + \sigma K_2 \log^a(1/\beta_\sigma) \right) \\ &= K_2 R_x \log^a(1/\beta_x) \left(\sqrt{\|\mathbf{H}\|_2} \|\mathbf{w}^*\| \log^a(1/\beta_{\mathbf{w}^*}) + \sigma \log^a(1/\beta_\sigma) \right). \end{aligned}$$

\therefore from Section 2, we have $\|\mathbf{x}_0\| \leq R_x \log^a(1/\beta_x)$ w.p. at least $1 - \beta_x$ and $|y_0| = |\langle \mathbf{x}_0, \mathbf{w}^* \rangle + z_0| \leq |\langle \mathbf{x}_0, \mathbf{w}^* \rangle| + |z_0| \leq K_2 \sqrt{\|\mathbf{H}\|_2} \|\mathbf{w}^*\| \log^a(1/\beta_{\mathbf{w}^*}) + \sigma K_2 \log^a(1/\beta_\sigma)$ w.p. at least $1 - \beta_{\mathbf{w}^*} - \beta_\sigma$. As each β is $1/\text{poly}(N)$, the Lemma holds.

Induction hypothesis: The Lemma holds for iterations $t = 0, \dots, T-1$.

Inductive case: Check at $t = T$.

The norm of the gradient at iteration $t = T$ is $\|\mathbf{x}_T(\langle \mathbf{x}_T, \mathbf{w}_T \rangle - y_T)\|$

$$\begin{aligned}
&= \|\mathbf{x}_T \mathbf{x}_T^\top (\mathbf{w}_T - \mathbf{w}^*) + \mathbf{x}_T z_T\| \\
&\leq \|\mathbf{x}\| (\|\mathbf{x}_T^\top (\mathbf{w}_T - \mathbf{w}^*)\| + \|z_T\|) \\
&\leq R_x \log^a(1/\beta_x) \left(K_2 \sqrt{\|\mathbf{H}\|_2} \|\mathbf{w}_T - \mathbf{w}^*\| \log^a(1/\beta_{\mathbf{w}_T}) + \sigma K_2 \log^a(1/\beta_\sigma) \right) \\
&= K_2 R_x \log^{2a} N \left(\sqrt{\|\mathbf{H}\|_2} \|\mathbf{w}_T - \mathbf{w}^*\| + \sigma \right), \tag{11}
\end{aligned}$$

since each $\beta_x, \beta_{\mathbf{w}_T}, \beta_\sigma$ is $1/\text{poly}(N)$. Therefore, we need to find a bound on $\|\mathbf{w}_T - \mathbf{w}^*\|$.

As ζ has not been exceeded in iterations $t = 0, \dots, T-1$, we have the following decomposition at iteration $T-1$:

$$\begin{aligned}
\mathbf{w}_T - \mathbf{w}^* &= \mathbf{w}_{T-1} - \mathbf{w}^* - \eta \left(\text{clip}_\zeta(\mathbf{x}_{T-1}(\langle \mathbf{x}_{T-1}, \mathbf{w}_{T-1} \rangle - y_{T-1})) + 2\zeta \alpha \mathbf{g}_{T-1} \right) \\
&= \mathbf{w}_{T-1} - \mathbf{w}^* - \eta \left(\mathbf{x}_{T-1}(\langle \mathbf{x}_{T-1}, \mathbf{w}_{T-1} \rangle - y_{T-1}) + 2\zeta \alpha \mathbf{g}_{T-1} \right) \\
&= (\mathbf{I} - \eta \mathbf{x}_{T-1} \mathbf{x}_{T-1}^\top) (\mathbf{w}_{T-1} - \mathbf{w}^*) + \eta \boldsymbol{\xi}_{T-1} - 2\eta \zeta \alpha \mathbf{g}_{T-1} \\
&= \mathbf{B}_{T-1} (\mathbf{w}_{T-1} - \mathbf{w}^*) + \eta (\boldsymbol{\xi}_{T-1} - 2\zeta \alpha \mathbf{g}_{T-1}) \\
&= \mathbf{B}_{T-1} (\mathbf{w}_{T-1} - \mathbf{w}^*) + \eta \boldsymbol{\nu}_{T-1} \\
&= \mathbf{B}_{T-1} \dots \mathbf{B}_0 (\mathbf{w}_0 - \mathbf{w}^*) + \eta (\boldsymbol{\nu}_{T-1} + \mathbf{B}_{T-1} \boldsymbol{\nu}_{T-2} + \dots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \boldsymbol{\nu}_0),
\end{aligned}$$

where,

$$\mathbf{B}_t := \mathbf{I} - \eta \mathbf{x}_t \mathbf{x}_t^\top, \quad \boldsymbol{\xi}_t := \mathbf{x}_t (y_t - \langle \mathbf{x}_t, \mathbf{w}^* \rangle) = \mathbf{x}_t z_t, \quad \boldsymbol{\nu}_t := \boldsymbol{\xi}_t - 2\zeta \alpha \mathbf{g}_t.$$

Therefore,

$$\begin{aligned}
\|\mathbf{w}_T - \mathbf{w}^*\|^2 &= \|\mathbf{B}_{T-1} \dots \mathbf{B}_0 (\mathbf{w}_0 - \mathbf{w}^*) + \eta (\boldsymbol{\nu}_{T-1} + \mathbf{B}_{T-1} \boldsymbol{\nu}_{T-2} + \dots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \boldsymbol{\nu}_0)\|^2 \\
&\leq 2 \left(\underbrace{\|\mathbf{B}_{T-1} \dots \mathbf{B}_0 (\mathbf{w}_0 - \mathbf{w}^*)\|^2}_{\text{Bias Term}} + \underbrace{\|\eta (\boldsymbol{\nu}_{T-1} + \mathbf{B}_{T-1} \boldsymbol{\nu}_{T-2} + \dots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \boldsymbol{\nu}_0)\|^2}_{\text{Variance Term}} \right). \tag{12}
\end{aligned}$$

Bound on Bias Term:

Since $\mathbf{w}_0 - \mathbf{w}^*$ is independent of all \mathbf{x}_i 's, we have (w.p. $\geq 1 - T\beta_{\mathbf{B}}$),

$$\begin{aligned}
\|\mathbf{B}_{T-1} \dots \mathbf{B}_0 (\mathbf{w}_0 - \mathbf{w}^*)\|^2 &\leq \|\mathbf{B}_{T-1} \dots \mathbf{B}_0\|^2 \|\mathbf{w}_0 - \mathbf{w}^*\|^2 = \left\| \prod_{j=0}^{T-1} \mathbf{B}_j \right\|^2 \|\mathbf{w}_0 - \mathbf{w}^*\|^2 \\
&\stackrel{\zeta_1}{\leq} \left(\prod_{j=0}^{T-1} \|\mathbf{B}_j\|^2 \right) \|\mathbf{w}_0 - \mathbf{w}^*\|^2 \stackrel{\zeta_2}{\leq} \|\mathbf{w}_0 - \mathbf{w}^*\|^2 \stackrel{\zeta_3}{\leq} \|\mathbf{w}^*\|^2, \tag{13}
\end{aligned}$$

where ζ_1 follows by $\|\mathbf{A}\mathbf{B}\| \leq \|\mathbf{A}\| \|\mathbf{B}\|$, ζ_2 follows from using $\eta \leq \frac{1}{R_x^2 \log^{2a}(1/\beta_{\mathbf{B}})}$ and the fact that with probability $\geq 1 - \beta_{\mathbf{B}}$, we have $\|\mathbf{x}_j\|^2 \leq 1/\eta$, i.e.,

$$\mathbf{B}_j = \mathbf{I} - \eta \mathbf{x}_j \mathbf{x}_j^\top \preceq \mathbf{I} \implies \|\mathbf{B}_j\| \leq \|\mathbf{I}\| = 1.$$

Finally, ζ_3 follows from $\mathbf{w}_0 = \mathbf{0}$.

Bound on Variance Term:

$$\begin{aligned} & \|\eta(\boldsymbol{\nu}_{T-1} + \mathbf{B}_{T-1}\boldsymbol{\nu}_{T-2} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \boldsymbol{\nu}_0)\|^2 \\ &= \eta^2 \|(\boldsymbol{\xi}_{T-1} - 2\zeta\alpha\mathbf{g}_{T-1}) + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1(\boldsymbol{\xi}_0 - 2\zeta\alpha\mathbf{g}_0)\|^2 \\ &\leq 2\eta^2 \left(\|\boldsymbol{\xi}_{T-1} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \boldsymbol{\xi}_0\|^2 + 4\zeta^2\alpha^2 \|\mathbf{g}_{T-1} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \mathbf{g}_0\|^2 \right). \end{aligned} \quad (14)$$

We will compute the bound on each of the two terms separately and add them up.

Bound on $\|\mathbf{g}_{T-1} + \mathbf{B}_{T-1}\mathbf{g}_{T-2} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \mathbf{g}_0\|$:

As \mathbf{g}_τ is a Gaussian vector for all $0 \leq \tau \leq T-1$, and \mathbf{B}_τ 's are all independent of \mathbf{g}_τ , we have:

$$\mathbf{g}_{T-1} + \mathbf{B}_{T-1}\mathbf{g}_{T-2} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \mathbf{g}_0 = \tilde{\mathbf{g}} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}),$$

where the covariance \mathbf{Q} is given by:

$$\begin{aligned} \mathbf{Q} &= \mathbf{I} + \mathbf{B}_{T-1}\mathbf{B}_{T-1}^\top + \mathbf{B}_{T-1}\mathbf{B}_{T-2}\mathbf{B}_{T-2}^\top\mathbf{B}_{T-1}^\top + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1\mathbf{B}_1^\top \dots \mathbf{B}_{T-1}^\top \\ &= \mathbf{I} + \sum_{\tau=1}^{T-1} \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right)^\top \\ &= \mathbf{I} + \sum_{\tau=1}^{\tilde{\tau}} \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right)^\top + \sum_{\tau=\tilde{\tau}+1}^{T-1} \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right)^\top. \end{aligned}$$

Using Lemma 17, we have w.p. $\geq 1 - T\beta_{var1,max}$,

$$\begin{aligned} \text{Tr} \left(\sum_{\tau=1}^{\tilde{\tau}} \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right)^\top \right) &= \sum_{\tau=1}^{\tilde{\tau}} \text{Tr} \left(\left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right)^\top \right) \\ &\leq \sum_{\tau=1}^{\tilde{\tau}} \frac{d}{\beta_{var1,min}} e^{-\eta\mu \cdot (T-\tau)}, \end{aligned} \quad (15)$$

where $\beta_{var1,min}$ and $\beta_{var1,max}$ are the minimum and maximum β 's respectively across all the Lemma 17 invocations.

Furthermore, for any τ , we have:

$$\left\| \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right)^\top \right\| \leq \prod_{j=\tau}^{T-1} \|\mathbf{B}_j\|^2 \leq 1 \quad (16)$$

w.p. $\geq 1 - (T-\tau)\beta_{\mathbf{B}}$, where we use the fact that since $\eta \leq \frac{1}{R_x^2 \log^{2a}(1/\beta_{\mathbf{B}})}$ w.p. $\geq 1 - \beta_{\mathbf{B}}$,

$$\mathbf{B}_j = \mathbf{I} - \eta \mathbf{x}_j \mathbf{x}_j^\top \preceq \mathbf{I} \implies \|\mathbf{B}_j\| \leq \|\mathbf{I}\| = 1.$$

Thus using (15) and (16), we have:

$$\begin{aligned}
\text{Tr}(\mathbf{Q}) &= \text{Tr}(\mathbf{I}) + \text{Tr} \left(\sum_{\tau=1}^{\tilde{\tau}} \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right)^\top \right) + \text{Tr} \left(\sum_{\tau=\tilde{\tau}+1}^{T-1} \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right)^\top \right) \\
&\leq d + \sum_{\tau=1}^{\tilde{\tau}} \frac{d}{\beta_{var1,\min}} e^{-\eta\mu(T-\tau)} + \sum_{\tau=\tilde{\tau}+1}^{T-1} d \left\| \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right)^\top \right\| \\
&\leq d + \frac{d}{\beta_{var1,\min}} e^{-\eta\mu(T-\tilde{\tau})} \frac{1}{1 - e^{-\eta\mu}} + d(T - \tilde{\tau} - 1) \cdot 1 \\
&\leq d \cdot (T - \tilde{\tau}) + \frac{3d}{\beta_{var1,\min}} e^{-\eta\mu(T-\tilde{\tau})}.
\end{aligned}$$

Hence, selecting $T - \tilde{\tau} = \frac{101 \log(1/\beta_{var1,\min})}{\eta\mu}$, we have w.p. $\geq 1 - T(\beta_{var1,\max} + \beta_{\mathbf{B}})$,

$$\text{Tr}(\mathbf{Q}) \leq d \cdot \left(\frac{101}{\eta\mu} \log(1/\beta_{var1,\min}) + 3\beta_{var1,\min}^{100} \right). \quad (17)$$

Now, using standard Gaussian property along with the bound on $\text{Tr}(\mathbf{Q})$ given above we have w.p. $\geq 1 - T(\beta_{var1,\max} + \beta_{\mathbf{B}}) - \beta_{tr}$,

$$\begin{aligned}
\|\mathbf{g}_{T-1} + \mathbf{B}_{T-1}\mathbf{g}_{T-2} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \mathbf{g}_0\| &\leq \sqrt{\text{Tr}(\mathbf{Q})} \sqrt{\log(1/\beta_{tr})} \\
&\leq \sqrt{d \log(1/\beta_{tr})} \cdot \sqrt{\frac{101}{\eta\mu} \log(1/\beta_{var1,\min}) + 3\beta_{var1,\min}^{100}}. \quad (18)
\end{aligned}$$

Bound on $\|\xi_{T-1} + \mathbf{B}_{T-1}\xi_{T-2} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \xi_0\|$

Note that

$$\begin{aligned}
&\|\xi_{T-1} + \mathbf{B}_{T-1}\xi_{T-2} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \xi_0\| \\
&= \|\mathbf{x}_{T-1} z_{T-1} + \mathbf{B}_{T-1} \mathbf{x}_{T-2} z_{T-2} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \mathbf{x}_0 z_0\| := \|\mathbf{Vz}\|,
\end{aligned}$$

where \mathbf{z} is a vector having entries $[z_0, \dots, z_{T-1}]$ and \mathbf{V} is a $d \times T$ matrix with the j^{th} column being the vector $\mathbf{B}_{T-1} \dots \mathbf{B}_j \mathbf{x}_{j-1}$ and the T^{th} column being the vector \mathbf{x}_{T-1} .

Note that since each z_j is independent and $z_j \sim \mathcal{N}(0, \sigma^2) \implies \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}) \implies \mathbf{Vz} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{V}^\top \mathbf{V})$.

Therefore, using Lemma 16, we have with probability $\geq 1 - \beta_f$,

$$\begin{aligned}
\|\mathbf{Vz}\|^2 &\leq \text{Tr}(\sigma^2 \mathbf{V}^\top \mathbf{V}) + 2\|\sigma^2 \mathbf{V}^\top \mathbf{V}\|_{\text{F}} \sqrt{\log(1/\beta_f)} + 2\|\sigma^2 \mathbf{V}^\top \mathbf{V}\| \log(1/\beta_f) \\
&\leq \sigma^2 \text{Tr}(\mathbf{V}^\top \mathbf{V}) + 2\sigma^2 \|\mathbf{V}^\top \mathbf{V}\|_{\text{F}} \sqrt{\log(1/\beta_f)} + 2\sigma^2 \text{Tr}(\mathbf{V}^\top \mathbf{V}) \log(1/\beta_f) \quad (19) \\
&= \sigma^2 \text{Tr}(\mathbf{V}^\top \mathbf{V}) + 2\sigma^2 \sqrt{\text{Tr}(\mathbf{V}^\top \mathbf{V} \mathbf{V}^\top \mathbf{V})} \log(1/\beta_f) + 2\sigma^2 \text{Tr}(\mathbf{V}^\top \mathbf{V}) \log(1/\beta_f) \\
&\leq \sigma^2 \text{Tr}(\mathbf{V}^\top \mathbf{V}) + 2\sigma^2 \sqrt{\|\mathbf{V}^\top \mathbf{V}\| \text{Tr}(\mathbf{V}^\top \mathbf{V})} \log(1/\beta_f) + 2\sigma^2 \text{Tr}(\mathbf{V}^\top \mathbf{V}) \log(1/\beta_f) \\
&= \sigma^2 \text{Tr}(\mathbf{V}^\top \mathbf{V}) + 2\sigma^2 \text{Tr}(\mathbf{V}^\top \mathbf{V}) \sqrt{\log(1/\beta_f)} + 2\sigma^2 \text{Tr}(\mathbf{V}^\top \mathbf{V}) \log(1/\beta_f) \\
&= \sigma^2 \text{Tr}(\mathbf{V}^\top \mathbf{V}) (1 + 2\sqrt{\log(1/\beta_f)} + 2\log(1/\beta_f)) \\
&\leq 5\sigma^2 \text{Tr}(\mathbf{V}^\top \mathbf{V}) \log(1/\beta_f), \quad (20)
\end{aligned}$$

where in (19), we have used the fact that $\|\mathbf{M}\| \leq \text{Tr}(\mathbf{M})$ and $\|\mathbf{M}\|_F = \sqrt{\text{Tr}(\mathbf{M}^T \mathbf{M})}$ for a symmetric matrix \mathbf{M} .

Our aim is now to compute the value:

$$\begin{aligned}
 \text{Tr}(\mathbf{V}^T \mathbf{V}) &= \|\mathbf{x}_{T-1}\|^2 + \sum_{j=1}^{T-1} \|\mathbf{B}_{T-1} \mathbf{B}_{T-2} \dots \mathbf{B}_j \mathbf{x}_{j-1}\|^2 \\
 &= \|\mathbf{x}_{T-1}\|^2 + \sum_{\tau=1}^{T-1} \left\| \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \mathbf{x}_{j-1} \right\|^2 \\
 &= \|\mathbf{x}_{T-1}\|^2 + \sum_{\tau=1}^{\tilde{\tau}} \left\| \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \mathbf{x}_{j-1} \right\|^2 + \sum_{\tau=\tilde{\tau}+1}^{T-1} \left\| \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \mathbf{x}_{j-1} \right\|^2. \tag{21}
 \end{aligned}$$

Using Lemma 17, we have w.p. $\geq 1 - T\beta_{var2,max}$,

$$\sum_{\tau=1}^{\tilde{\tau}} \left\| \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \mathbf{x}_{j-1} \right\|^2 \leq \sum_{\tau=1}^{\tilde{\tau}} \frac{R_x^2}{\beta_{var2,min}} e^{-\eta\mu \cdot (T-\tau)}, \tag{22}$$

where $\beta_{var2,min}$ and $\beta_{var2,max}$ are the minimum and maximum β 's respectively across all the Lemma 17 invocations.

Also, for any τ , we have w.p. $\geq 1 - (T - \tau)\beta_{\mathbf{B}}$,

$$\left\| \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \mathbf{x}_{j-1} \right\|^2 \leq \left\| \prod_{j=\tau}^{T-1} \mathbf{B}_j \right\|^2 \|\mathbf{x}_{j-1}\|^2 \leq \|\mathbf{x}_{j-1}\|^2 \leq R_x^2 \log^{2a}(1/\beta_{\mathbf{B}}), \tag{23}$$

where we again use the fact that since $\eta \leq \frac{1}{R_x^2 \log^{2a}(1/\beta_{\mathbf{B}})}$ w.p. $\geq 1 - \beta_{\mathbf{B}}$,

$$\mathbf{B}_j = \mathbf{I} - \eta \mathbf{x}_j \mathbf{x}_j^T \preceq \mathbf{I} \implies \|\mathbf{B}_j\| \leq \|\mathbf{I}\| = 1.$$

Using (22) and (23) in (21), we have w.p. at least $1 - T\beta_{var2,max} - (T - \tilde{\tau})\beta_{\mathbf{x},max} - T\beta_{\mathbf{B}} \geq 1 - T(\beta_{var2,max} + \beta_{\mathbf{x},max} + \beta_{\mathbf{B}})$,

$$\begin{aligned}
 \text{Tr}(\mathbf{V}^T \mathbf{V}) &= \|\mathbf{x}_{T-1}\|^2 + \sum_{\tau=1}^{\tilde{\tau}} \left\| \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \mathbf{x}_{j-1} \right\|^2 + \sum_{\tau=\tilde{\tau}+1}^{T-1} \left\| \left(\prod_{j=\tau}^{T-1} \mathbf{B}_j \right) \mathbf{x}_{j-1} \right\|^2 \\
 &\leq R_x^2 \log^{2a}(1/\beta_{\mathbf{x},min}) + \sum_{\tau=1}^{\tilde{\tau}} \frac{R_x^2}{\beta_{var2,min}} e^{-\eta\mu \cdot (T-\tau)} + \sum_{\tau=\tilde{\tau}+1}^{T-1} R_x^2 \log^{2a}(1/\beta_{\mathbf{x},min}) \\
 &\leq R_x^2 \log^{2a}(1/\beta_{\mathbf{x},min}) \\
 &\quad + \frac{R_x^2}{\beta_{var2,min}} e^{-\eta\mu(T-\tilde{\tau})} \frac{1}{1 - e^{-\eta\mu}} + R_x^2 \log^{2a}(1/\beta_{\mathbf{x},min})(T - \tilde{\tau} - 1) \\
 &\leq R_x^2 \log^{2a}(1/\beta_{\mathbf{x},min})(T - \tilde{\tau}) + \frac{3R_x^2}{\beta_{var2,min}} e^{-\eta\mu(T-\tilde{\tau})}.
 \end{aligned}$$

Again, selecting $T - \tilde{\tau} = \frac{101 \log(1/\beta_{var2,\min})}{\eta\mu}$, we have w.p. $\geq 1 - T(\beta_{var2,\max} + \beta_{\mathbf{x},\max} + \beta_{\mathbf{B}})$,

$$\text{Tr}(\mathbf{V}^T \mathbf{V}) \leq R_x^2 \cdot \left(\log^{2a}(1/\beta_{\mathbf{x},\min}) \frac{101}{\eta\mu} \log(1/\beta_{var2,\min}) + 3\beta_{var2,\min}^{100} \right). \quad (24)$$

Using this in (18), we have w.p. $\geq 1 - T(\beta_{var2,\max} + \beta_{\mathbf{x},\max} + \beta_{\mathbf{B}}) - \beta_f$,

$$\|\mathbf{V}\mathbf{z}\|^2 \leq 5\sigma^2 R_x^2 \cdot \left(\log^{2a}(1/\beta_{\mathbf{x},\min}) \frac{101}{\eta\mu} \log(1/\beta_{var2,\min}) + 3\beta_{var2,\min}^{100} \right) \log(1/\beta_f). \quad (25)$$

Combining (18) and (25) in (14), we obtain the overall variance bound w.p. at least $1 - T(2\beta_{var,\max} + \beta_{\mathbf{x},\max} + \beta_{\mathbf{B}}) - \beta_f - \beta_{tr}$,

$$\begin{aligned} & \|\eta(\boldsymbol{\nu}_{T-1} + \mathbf{B}_{T-1}\boldsymbol{\nu}_{T-2} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \boldsymbol{\nu}_0)\|^2 \\ & \leq 2\eta^2 \left(\|\boldsymbol{\xi}_{T-1} + \mathbf{B}_{T-1}\boldsymbol{\xi}_{T-2} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \boldsymbol{\xi}_0\|^2 \right. \\ & \quad \left. + 4\zeta^2 \alpha^2 \|\mathbf{g}_{T-1} + \mathbf{B}_{T-1}\mathbf{g}_{T-2} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \mathbf{g}_0\|^2 \right) \\ & \leq 2\eta^2 \left(\left(5\sigma^2 R_x^2 \cdot \left(\log^{2a}(1/\beta_{\mathbf{x},\min}) \frac{101}{\eta\mu} \log(1/\beta_{var,\min}) + 3\beta_{var,\min}^{100} \right) \log(1/\beta_f) \right) \right. \\ & \quad \left. + 4\zeta^2 \alpha^2 d \log(1/\beta_{tr}) \left(\frac{101}{\eta\mu} \log(1/\beta_{var,\min}) + 3\beta_{var,\max}^{100} \right) \right) \\ & = 2\eta^2 \left(\left(5\sigma^2 R_x^2 \log(1/\beta_{var,\min}) \log^{2a}(1/\beta_{\mathbf{x},\min}) \log(1/\beta_f) \right) \right. \\ & \quad \left. + 4\zeta^2 \alpha^2 d \log(1/\beta_{tr}) \log(1/\beta_{var,\min}) \right) \frac{101}{\eta\mu} \\ & \quad + \left(5\sigma^2 R_x^2 \log(1/\beta_f) + 4\zeta^2 \alpha^2 d \log(1/\beta_{tr}) \right) 3\beta_{var,\max}^{100}, \end{aligned} \quad (26)$$

where $\beta_{var,\max} = \max(\beta_{var1,\max}, \beta_{var2,\max})$ and $\beta_{var,\min} = \max(\beta_{var1,\max}, \beta_{var2,\max})$.

Setting all the above β 's as $\tilde{\beta}$, and for $\tilde{\beta} \leq 1/\eta\mu$, we get w.p. at least $1 - (4T + 2)\tilde{\beta}$,

$$\|\eta(\boldsymbol{\nu}_{T-1} + \mathbf{B}_{T-1}\boldsymbol{\nu}_{T-2} + \cdots + \mathbf{B}_{T-1} \dots \mathbf{B}_1 \boldsymbol{\nu}_0)\|^2 \leq \frac{210\eta}{\mu} \left(5\sigma^2 R_x^2 \log^{2a+2}(1/\tilde{\beta}) + 4\zeta^2 \alpha^2 d \log^2(1/\tilde{\beta}) \right). \quad (27)$$

Using the Bias bound (13) and Variance bound (27) in (12) and setting the additional $\beta_{\mathbf{B}}$ arising from the Bias bound as well to $\tilde{\beta}$, we have w.p. $\geq 1 - (5T + 2)\tilde{\beta}$,

$$\|\mathbf{w}_T - \mathbf{w}^*\|^2 \leq 2 \left(\|\mathbf{w}^*\|^2 + \frac{210\eta}{\mu} \left(5\sigma^2 R_x^2 \log^{2a+2}(1/\tilde{\beta}) + 4\zeta^2 \alpha^2 d \log^2(1/\tilde{\beta}) \right) \right),$$

Setting $\tilde{\beta} = \frac{1}{\text{Poly}(N)}$, we get w.p. $\geq 1 - \frac{1}{\text{Poly}(N)}$,

$$\|\mathbf{w}_T - \mathbf{w}^*\| \leq 2\|\mathbf{w}^*\| + 2\sqrt{\frac{1050\eta}{\mu}} R_x \sigma \log^{a+1} N + 2\sqrt{\frac{210\eta \cdot d}{\mu}} \cdot 2\zeta \alpha \log N. \quad (28)$$

Using (28) in (11), we get w.p. $\geq 1 - \frac{1}{\text{Poly}(N)}$,

$$\begin{aligned} \|\mathbf{x}_T(\langle \mathbf{x}_T, \mathbf{w}_T \rangle - y_T)\| &\leq K_2 R_x \log^{2a} N \left(\sqrt{\|\mathbf{H}\|_2} \|\mathbf{w}_T - \mathbf{w}^*\| + \sigma \right) \\ &\leq K_2 R_x \log^{2a} N \left(2\sqrt{\|\mathbf{H}\|_2} \left(\|\mathbf{w}^*\| + \sqrt{\frac{1050\eta}{\mu}} R_x \sigma \log^{a+1} N + 2\sqrt{\frac{210\eta \cdot d}{\mu}} \cdot \zeta \alpha \log N \right) + \sigma \right). \end{aligned}$$

Now, $\alpha = \alpha_{\varepsilon, \delta, N}$. Hence, setting η , s.t.,

$$1 - 4\sqrt{\frac{210\eta \cdot d}{\mu}} \alpha_{\varepsilon, \delta, N} \cdot K_2 \sqrt{\|\mathbf{H}\|_2} \cdot R_x \log^{2a+1} N \geq 1/2, \quad (29)$$

i.e., if

$$\eta \leq \min \left\{ \frac{1}{64 \cdot 210 \cdot \log^{4a+2} N} \cdot \frac{\mu}{K_2^2 R_x^2 \|\mathbf{H}\|} \cdot \frac{1}{d \alpha_{\varepsilon, \delta, N}^2}, \frac{1}{4 \cdot 1050 \cdot \log^{2a+2} N} \cdot \frac{1}{R_x^2} \right\}, \quad (30)$$

then w.p. $\geq 1 - \frac{1}{\text{Poly}(N)}$,

$$\|\mathbf{x}_T(\langle \mathbf{x}_T, \mathbf{w}_T \rangle - y_T)\| \leq \zeta,$$

and

$$\zeta = 4K_2 R_x \cdot \log^{2a} N \cdot \left(\sqrt{\|\mathbf{H}\|_2} \|\mathbf{w}^*\| + \sqrt{\kappa} \sigma \right). \quad (31)$$

■

A.2.2. PROOF OF LEMMA 8

Proof As in the proof of Lemma 19, let us define:

$$\mathbf{B}_t := \mathbf{I} - \eta \mathbf{x}_t \mathbf{x}_t^\top, \quad \boldsymbol{\xi}_t := (y_t - \langle \mathbf{x}_t, \mathbf{w}^* \rangle) \mathbf{x}_t = z_t \mathbf{x}_t, \quad \boldsymbol{\nu}_t := \boldsymbol{\xi}_t - 2\zeta \alpha \mathbf{g}_t.$$

Since there is no clipping, the update rule can be written as:

$$\begin{aligned} \mathbf{w}_{t+1} - \mathbf{w}^* &= \mathbf{w}_t - \mathbf{w}^* - \eta \text{clip}_\zeta(\mathbf{x}_t(\langle \mathbf{x}_t, \mathbf{w}_t \rangle - y_t)) - 2\eta \zeta \alpha \mathbf{g}_t \\ &= \mathbf{w}_t - \mathbf{w}^* - \eta \left(\mathbf{x}_t(\langle \mathbf{x}_t, \mathbf{w}_t \rangle - \langle \mathbf{x}_t, \mathbf{w}^* \rangle - z_t) + 2\zeta \alpha \mathbf{g}_t \right) \\ &= (\mathbf{I} - \eta \mathbf{x}_t \mathbf{x}_t^\top) (\mathbf{w}_t - \mathbf{w}^*) + \eta_t \boldsymbol{\xi}_t - 2\eta \zeta \alpha \mathbf{g}_t \\ &= \mathbf{B}_t (\mathbf{w}_t - \mathbf{w}^*) + \eta (\boldsymbol{\xi}_t - 2\zeta \alpha \mathbf{g}_t) = \mathbf{B}_t (\mathbf{w}_t - \mathbf{w}^*) + \eta \boldsymbol{\nu}_t. \end{aligned} \quad (32)$$

From recursion,

$$\mathbf{w}_t - \mathbf{w}^* = \mathbf{B}_{t-1} \dots \mathbf{B}_0 (\mathbf{w}_0 - \mathbf{w}^*) + \eta (\boldsymbol{\nu}_{t-1} + \mathbf{B}_{t-1} \boldsymbol{\nu}_{t-2} + \dots + \mathbf{B}_{t-1} \dots \mathbf{B}_1 \boldsymbol{\nu}_0). \quad (33)$$

Note that both the recursive form (32) and the expanded form (33) are similar to that of equation (1) in Jain et al. (2017) which analyzed a non-private version of SGD, except that we now have the term $\boldsymbol{\nu}_t = \boldsymbol{\xi}_t - 2\zeta \alpha \mathbf{g}_t$ in place of $\boldsymbol{\xi}_t$.

The above equation implies:

$$\begin{aligned}
\bar{\mathbf{w}} - \mathbf{w}^* &= \left(\frac{2}{N} \sum_{\tau=N/2+1}^N \mathbf{w}_\tau \right) - \mathbf{w}^* = \frac{2}{N} \sum_{\tau=N/2+1}^N (\mathbf{w}_\tau - \mathbf{w}^*) \\
&= \frac{2}{N} \sum_{\tau=N/2+1}^N \left(\mathbf{B}_{\tau-1} \dots \mathbf{B}_0 (\mathbf{w}_0 - \mathbf{w}^*) + \eta (\boldsymbol{\nu}_{\tau-1} + \dots + \mathbf{B}_{\tau-1} \dots \mathbf{B}_1 \boldsymbol{\nu}_0) \right) \\
&= \frac{2}{N} \sum_{\tau=N/2+1}^N \mathbf{B}_{\tau-1} \dots \mathbf{B}_0 (\mathbf{w}_0 - \mathbf{w}^*) \\
&\quad + \frac{2\eta}{N} \sum_{\tau=N/2+1}^N \boldsymbol{\nu}_{\tau-1} + \mathbf{B}_{\tau-1} \boldsymbol{\nu}_{\tau-2} + \dots + \mathbf{B}_{\tau-1} \dots \mathbf{B}_1 \boldsymbol{\nu}_0.
\end{aligned}$$

Therefore,

$$\begin{aligned}
\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\bar{\mathbf{w}} - \mathbf{w}^*\|_{\mathbf{H}}^2] &\leq \left(\sqrt{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left\| \frac{2}{N} \sum_{\tau=N/2+1}^N \mathbf{B}_{\tau-1} \dots \mathbf{B}_0 (\mathbf{w}_0 - \mathbf{w}^*) \right\|_{\mathbf{H}}^2 \right]} \right. \\
&\quad \left. + \sqrt{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left\| \frac{2\eta}{N} \sum_{\tau=N/2+1}^N \boldsymbol{\nu}_{\tau-1} + \mathbf{B}_{\tau-1} \boldsymbol{\nu}_{\tau-2} + \dots + \mathbf{B}_{\tau-1} \dots \mathbf{B}_1 \boldsymbol{\nu}_0 \right\|_{\mathbf{H}}^2 \right]} \right)^2, \\
\implies \mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) &\leq \frac{1}{2} \left(\underbrace{\sqrt{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\bar{\mathbf{w}} - \mathbf{w}^*\|_{\mathbf{H}}^2 | \boldsymbol{\nu}_0 = \dots = \boldsymbol{\nu}_{N-1} = \mathbf{0}]} }_{\text{Bias Term}} \right. \\
&\quad \left. + \underbrace{\sqrt{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\bar{\mathbf{w}} - \mathbf{w}^*\|_{\mathbf{H}}^2 | \mathbf{w}_0 = \mathbf{w}^*]} }_{\text{Variance Term}} \right)^2 \tag{34}
\end{aligned}$$

where,

$$\begin{aligned}
\underbrace{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\bar{\mathbf{w}} - \mathbf{w}^*\|_{\mathbf{H}}^2 | \boldsymbol{\nu}_0 = \dots = \boldsymbol{\nu}_{N-1} = \mathbf{0}]}_{\text{Bias Term}} &:= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left\| \frac{2}{N} \sum_{\tau=N/2+1}^N \mathbf{B}_{\tau-1} \dots \mathbf{B}_0 \mathbf{w}^* \right\|_{\mathbf{H}}^2 \right], \\
\underbrace{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\bar{\mathbf{w}} - \mathbf{w}^*\|_{\mathbf{H}}^2 | \mathbf{w}_0 = \mathbf{w}^*]}_{\text{Variance Term}} &:= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left\| \frac{2\eta}{N} \sum_{\tau=N/2+1}^N \boldsymbol{\nu}_{\tau-1} + \dots + \mathbf{B}_{\tau-1} \dots \mathbf{B}_1 \boldsymbol{\nu}_0 \right\|_{\mathbf{H}}^2 \right].
\end{aligned}$$

This is because,

$$\mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) = \frac{1}{2} \|\bar{\mathbf{w}} - \mathbf{w}^*\|_{\mathbf{H}}^2.$$

Following along the lines of [Jain et al. \(2017\)](#), the Bias-Variance Analysis is as follows:

Bias Term Analysis

$$\begin{aligned}
 & \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_t - \mathbf{w}^*\|^2 | \boldsymbol{\nu}_0 = \dots = \boldsymbol{\nu}_{N-1} = \mathbf{0}] = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_t - \mathbf{w}^*\|^2 | \mathbf{g}_0 \dots \mathbf{g}_{N-1} = \mathbf{0}, \boldsymbol{\xi}_0 \dots \boldsymbol{\xi}_{N-1} = \mathbf{0}] \\
 & = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{B}_{t-1}(\mathbf{w}_{t-1} - \mathbf{w}^*)\|^2] = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|(\mathbf{I} - \eta \mathbf{x}_{t-1} \mathbf{x}_{t-1}^\top)(\mathbf{w}_{t-1} - \mathbf{w}^*)\|^2] \\
 & = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{t-1} - \mathbf{w}^*\|^2] - \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [2\eta \langle (\mathbf{w}_{t-1} - \mathbf{w}^*), \mathbf{x}_{t-1} \mathbf{x}_{t-1}^\top (\mathbf{w}_{t-1} - \mathbf{w}^*) \rangle] \\
 & \quad + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\eta^2 \langle (\mathbf{w}_{t-1} - \mathbf{w}^*), \|\mathbf{x}_{t-1}\|^2 \mathbf{x}_{t-1} \mathbf{x}_{t-1}^\top (\mathbf{w}_{t-1} - \mathbf{w}^*) \rangle] \\
 & \leq \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{t-1} - \mathbf{w}^*\|^2] - 2\eta \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\langle (\mathbf{w}_{t-1} - \mathbf{w}^*), \mathbf{H}(\mathbf{w}_{t-1} - \mathbf{w}^*) \rangle] \\
 & \quad + \eta^2 R_x^2 \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\langle (\mathbf{w}_{t-1} - \mathbf{w}^*), \mathbf{H}(\mathbf{w}_{t-1} - \mathbf{w}^*) \rangle] \\
 & = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{t-1} - \mathbf{w}^*\|^2] - \eta(2 - \eta R_x^2) \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\langle (\mathbf{w}_{t-1} - \mathbf{w}^*), \mathbf{H}(\mathbf{w}_{t-1}, \mathbf{w}^*) \rangle] \\
 & \leq \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{t-1} - \mathbf{w}^*\|^2] - \eta \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\langle (\mathbf{w}_{t-1} - \mathbf{w}^*), \mathbf{H}(\mathbf{w}_{t-1}, \mathbf{w}^*) \rangle] \tag{35} \\
 & \leq (1 - \eta\mu) \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{t-1} - \mathbf{w}^*\|^2] \leq (1 - \eta\mu)^t \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_0 - \mathbf{w}^*\|^2] = e^{-\eta\mu t} \|\mathbf{w}^*\|^2,
 \end{aligned}$$

where in (35), we have used the fact that $\eta < \frac{1}{R_x^2}$.

The Bias Term's contribution to the overall risk is therefore:

$$\begin{aligned}
 & \frac{1}{2} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\bar{\mathbf{w}} - \mathbf{w}^*\|_{\mathbf{H}}^2 | \boldsymbol{\nu}_0 \dots \boldsymbol{\nu}_{N-1} = \mathbf{0}] \leq \frac{R_x^2}{2} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\bar{\mathbf{w}} - \mathbf{w}^*\|^2 | \boldsymbol{\nu}_0 \dots \boldsymbol{\nu}_{N-1} = \mathbf{0}] \\
 & = \frac{R_x^2}{2} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left\| \frac{2}{N} \sum_{\tau=N/2+1}^N (\mathbf{w}_\tau - \mathbf{w}^*) \right\|^2 | \boldsymbol{\nu}_0 \dots \boldsymbol{\nu}_{N-1} = \mathbf{0} \right] \\
 & \leq \frac{R_x^2}{N} \sum_{\tau=N/2+1}^N \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_\tau - \mathbf{w}^*\|^2 | \mathbf{g}_0 \dots \mathbf{g}_{N-1} = \mathbf{0}, \boldsymbol{\xi}_0 \dots \boldsymbol{\xi}_{N-1} = \mathbf{0}] \\
 & \leq \frac{R_x^2}{N} \sum_{\tau=N/2+1}^N e^{-\eta\mu\tau} \|\mathbf{w}^*\|^2 \leq \frac{R_x^2}{N} \cdot \frac{N}{2} e^{-\eta\mu(N/2+1)} \|\mathbf{w}^*\|^2 = \frac{R_x^2}{2} e^{-\eta\mu(N/2+1)} \|\mathbf{w}^*\|^2. \tag{36}
 \end{aligned}$$

Variance Term Analysis

Now suppose $\mathbf{w}_0 = \mathbf{w}^*$. Define the covariance matrix:

$$\mathbf{C}_t := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [(\mathbf{w}_t - \mathbf{w}^*)(\mathbf{w}_t - \mathbf{w}^*)^\top | \mathbf{w}_0 = \mathbf{w}^*].$$

Using the recursion $\mathbf{w}_{t+1} - \mathbf{w}^* = \mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) + \eta\boldsymbol{\nu}_t$ from (32), we have \mathbf{C}_{t+1} :

$$\begin{aligned}
& := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [(\mathbf{w}_{t+1} - \mathbf{w}^*)(\mathbf{w}_{t+1} - \mathbf{w}^*)^\top] = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) + \eta\boldsymbol{\nu}_t \right) \left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) + \eta\boldsymbol{\nu}_t \right)^\top \right] \\
& = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right) \left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right)^\top \right] + \eta \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right) \boldsymbol{\nu}_t^\top \right] \\
& \quad + \eta \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\boldsymbol{\nu}_t \left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right)^\top \right] + \eta^2 \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\boldsymbol{\nu}_t \boldsymbol{\nu}_t^\top \right] \\
& = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right) \left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right)^\top \right] + \eta \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) (\boldsymbol{\xi}_t^\top - 2\zeta\alpha\mathbf{g}_t^\top) \right] \\
& \quad + \eta \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\boldsymbol{\xi}_t - 2\zeta\alpha\mathbf{g}_t) \left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right)^\top \right] + \eta^2 \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\boldsymbol{\xi}_t - 2\zeta\alpha\mathbf{g}_t) (\boldsymbol{\xi}_t - 2\zeta\alpha\mathbf{g}_t)^\top \right] \\
& = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right) \left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right)^\top \right] + \eta \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right) \boldsymbol{\xi}_t^\top \right] \\
& \quad - 2\eta\zeta\alpha \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right) \mathbf{g}_t^\top \right] + \eta \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\boldsymbol{\xi}_t \left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right)^\top \right] \\
& \quad - 2\eta\zeta\alpha \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{g}_t \left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right)^\top \right] + \eta^2 \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\boldsymbol{\xi}_t - 2\zeta\alpha\mathbf{g}_t) (\boldsymbol{\xi}_t - 2\zeta\alpha\mathbf{g}_t)^\top \right] \\
& = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right) \left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right)^\top \right] + \eta^2 \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\boldsymbol{\xi}_t - 2\zeta\alpha\mathbf{g}_t) (\boldsymbol{\xi}_t - 2\zeta\alpha\mathbf{g}_t)^\top \right] \\
& = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right) \left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right)^\top \right] + \eta^2 \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\boldsymbol{\xi}_t - 2\zeta\alpha\mathbf{g}_t) (\boldsymbol{\xi}_t - 2\zeta\alpha\mathbf{g}_t)^\top \right],
\end{aligned} \tag{37}$$

where in (37) we have used the fact that $\mathbb{E}_{(\mathbf{x}_t, y_t) \sim \mathcal{D}} [\boldsymbol{\xi}_t] = \mathbf{0}$ and $\mathbb{E}[\mathbf{g}_t] = \mathbf{0}$ has been sampled at each step independently of all other terms. Continuing to expand the above expression gives:

$$\begin{aligned}
\mathbf{C}_{t+1} & = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right) \left(\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) \right)^\top \right] \\
& \quad + \eta^2 \left(\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\boldsymbol{\xi}_t \boldsymbol{\xi}_t^\top] - 2\zeta\alpha \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\boldsymbol{\xi}_t \mathbf{g}_t^\top] - 2\zeta\alpha \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\mathbf{g}_t \boldsymbol{\xi}_t^\top] + 4\zeta^2\alpha^2 \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\mathbf{g}_t \mathbf{g}_t^\top] \right) \\
& = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{B}_t(\mathbf{w}_t - \mathbf{w}^*) (\mathbf{w}_t - \mathbf{w}^*)^\top \mathbf{B}_t^\top \right] + \eta^2 (\boldsymbol{\Sigma} - \mathbf{0} - \mathbf{0} + 4\zeta^2\alpha^2 \mathbf{I}) \\
& = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\mathbf{I} - \eta \mathbf{x}_t \mathbf{x}_t^\top) (\mathbf{w}_t - \mathbf{w}^*) (\mathbf{w}_t - \mathbf{w}^*)^\top (\mathbf{I} - \eta \mathbf{x}_t \mathbf{x}_t^\top) \right] + \eta^2 (\boldsymbol{\Sigma} + 4\zeta^2\alpha^2 \mathbf{I}) \\
& = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\mathbf{w}_t - \mathbf{w}^*) (\mathbf{w}_t - \mathbf{w}^*)^\top \right] - \eta \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\mathbf{x} \mathbf{x}^\top) (\mathbf{w}_t - \mathbf{w}^*) (\mathbf{w}_t - \mathbf{w}^*)^\top \right] \\
& \quad - \eta \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\mathbf{w}_t - \mathbf{w}^*) (\mathbf{w}_t - \mathbf{w}^*)^\top (\mathbf{x} \mathbf{x}^\top) \right] + \eta^2 \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\mathbf{x} \mathbf{x}^\top) (\mathbf{w}_t - \mathbf{w}^*) (\mathbf{x} \mathbf{x}^\top) \right] \\
& \quad + \eta^2 (\boldsymbol{\Sigma} + 4\zeta^2\alpha^2 \mathbf{I}) \\
& = \mathbf{C}_t - \eta \mathbf{H} \mathbf{C}_t - \eta \mathbf{C}_t \mathbf{H} + \eta^2 \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{x}^\top \mathbf{C}_t \mathbf{x} \mathbf{x} \mathbf{x}^\top \right] + \eta^2 (\boldsymbol{\Sigma} + 4\zeta^2\alpha^2 \mathbf{I}),
\end{aligned} \tag{38}$$

where in (38) we have again used the fact that $\mathbb{E}[\mathbf{g}_t] = \mathbf{0}$ has been sampled independently at each step. Comparing this with equation (2) of Jain et al. (2017), we find that the term $\gamma^2 \boldsymbol{\Sigma}$ has been

replaced with $\eta^2(\Sigma + 4\zeta^2\alpha^2\mathbf{I})$ where $\Sigma = \mathbb{E}_{(\mathbf{x}_t, y_t) \sim \mathcal{D}} [\xi_t \xi_t^\top]$. Therefore the bound for the Variance term follows almost similarly from the paper, except that Σ gets replaced with $\Sigma + 4\zeta^2\alpha^2\mathbf{I}$.

Using Lemma 18, we have the Variance Term contribution as

$$\begin{aligned} \frac{1}{2} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\bar{\mathbf{w}} - \mathbf{w}^*\|_{\mathbf{H}}^2 | \mathbf{w}_0 = \mathbf{w}^*] &\leq \frac{\text{Tr}(\mathbf{C}_\infty)}{\eta(N - (N/2 + 1) + 1)} \\ &\leq \frac{1}{N} \frac{\eta R_x^2}{1 - \eta R_x^2} d \|\Sigma + 4\zeta^2\alpha^2\mathbf{I}\|_{\mathbf{H}} + \frac{1}{N} \text{Tr}(\mathbf{H}^{-1}(\Sigma + 4\zeta^2\alpha^2\mathbf{I})). \end{aligned} \quad (39)$$

Using (36) and (39) in (34) in:

$$\begin{aligned} \mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) &\leq \frac{1}{2} \left(\sqrt{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\bar{\mathbf{w}} - \mathbf{w}^*\|_{\mathbf{H}}^2 | \nu_0 = \dots = \nu_{N-1} = 0]} \right. \\ &\quad \left. + \sqrt{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\bar{\mathbf{w}} - \mathbf{w}^*\|_{\mathbf{H}}^2 | \mathbf{w}_0 = \mathbf{w}^*]} \right)^2 \end{aligned}$$

gives us the risk bound as:

$$\begin{aligned} &\leq \frac{1}{2} \left(\sqrt{\frac{R_x^2}{2} e^{-\eta\mu(\frac{N}{2}+1)} \|\mathbf{w}^*\|^2} + \sqrt{\frac{1}{N} \frac{\eta R_x^2}{1 - \eta R_x^2} d \|\Sigma + 4\zeta^2\alpha^2\mathbf{I}\|_{\mathbf{H}} + \frac{1}{N} \text{Tr}(\mathbf{H}^{-1}(\Sigma + 4\zeta^2\alpha^2\mathbf{I}))} \right)^2 \\ &\leq \frac{R_x^2}{2} e^{-\eta\mu(\frac{N}{2}+1)} \|\mathbf{w}^*\|^2 + \frac{1}{N} \frac{\eta R_x^2}{1 - \eta R_x^2} d \|\Sigma + 4\zeta^2\alpha^2\mathbf{I}\|_{\mathbf{H}} + \frac{1}{N} \text{Tr}(\mathbf{H}^{-1}(\Sigma + 4\zeta^2\alpha^2\mathbf{I})). \end{aligned}$$

This gives us the required result. ■

Appendix B. Missing Proofs from Section 3.2

B.1. Proof of Privacy Guarantees

B.1.1. PROOF OF LEMMA 12

Proof Since each of the computation in Step 3 of DP-STAT is of sensitivity one and the DP noise variance is $\lceil \log_2(B/\Delta) \rceil \alpha^2$, each step is $\left(\rho_i = \frac{1}{2^{\lceil \log_2(B/\Delta) \rceil} \alpha^2}\right)$ -zCDP. Since we perform at most $\lceil \log_2(B/\Delta) \rceil$ of such computations, by standard zCDP property of Gaussian mechanism, and its corresponding composition property (Bun and Steinke, 2016), the overall $\rho = \sum_{i=1}^{\lceil \log_2(B/\Delta) \rceil} \rho_i = \frac{1}{2\alpha^2}$ and thus the privacy guarantee follows immediately.

To prove the utility guarantee, first note that by standard concentration inequality for Gaussian distribution and by union bound, it follows that w.p. at least $1 - \beta$, none of the noise added in $\lceil \log_2(B/\Delta) \rceil$ iterations of Step 3 exceeds $\Gamma = \alpha \sqrt{2 \log(B/\Delta) \log(\log(B/\Delta)/\beta)}$. We will condition the rest of the proof on this event.

For any value of γ_i in iteration i of Step 3 of DP-STAT, $|c_{\text{priv}} - c| \leq \Gamma \implies c - \Gamma \leq c_{\text{priv}} \leq c + \Gamma$. Now note that, 1) if $c_{\text{priv}} \geq s$ then $c + \Gamma \geq s$, and 2) if $c_{\text{priv}} \leq s$ then $c - \Gamma \leq s \implies c \leq s$ since the true count can never exceed s . Combining the two cases gives $s - \Gamma \leq c \leq s$.

Furthermore, in the doubling search of Algorithm DP-STAT, each choice for γ_i is of the form $2^{i+1}\Delta$. Therefore, if the loop breaks out at iteration i^* , there exists a $\gamma \in [2^{i^*-1}\Delta, 2^{i^*}\Delta]$ s.t.

$s - \Gamma \leq |\{i : |\langle \mathbf{x}_i, \mathbf{w} \rangle - y_i| \leq \gamma : i \in \{0, \dots, s\}\}| \leq s$. Hence, error in γ_{priv} estimation can only be off by a factor of two, and an additional discretization error of Δ . This implies the utility guarantee. ■

B.1.2. PROOF OF THEOREM 9

Proof Our analysis will broadly involve computing the Zero Mean Concentrated Differential Privacy (zCDP) parameters and then using them to compute the Approximate Differential Privacy parameters. The update Step 10 of Algorithm DP-AMBSSGD without the additive Gaussian noise is:

$$\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t - \frac{\eta}{b} \sum_{i=0}^{b-1} \text{clip}_{\zeta_t}(\mathbf{x}_{\tau(t)+i}(\langle \mathbf{x}_{\tau(t)+i}, \mathbf{w}_t \rangle - y_{\tau(t)+i})),$$

where $\text{clip}_{\zeta}(\boldsymbol{\nu}) = \boldsymbol{\nu} \cdot \max\left\{1, \frac{\zeta}{\|\boldsymbol{\nu}\|_2}\right\}$. Therefore, the local L_2 sensitivity of the \mathbf{w}_{t+1} due to a sample difference in the τ^{th} batch is $\Delta_2 = \frac{2\eta\zeta_t}{b}$.

Since $\mathbf{g}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{d \times d})$, the above step is $\left(\rho_{t,1} = \frac{\Delta_2^2}{2 \cdot \frac{4\eta^2\zeta_t^2\alpha^2}{b^2}} = \frac{1}{2\alpha^2}\right)$ -zCDP since DP noise standard deviation is $\eta \frac{2\zeta_t\alpha}{b}$. (Proposition 1.6 of [Bun and Steinke \(2016\)](#)) and from Lemma 12, Step 10 is $\left(\rho_{t,2} = \frac{1}{2\alpha^2}\right)$ -zCDP. Therefore by composition, each iteration step is $\left(\rho_t = \rho_{t,1} + \rho_{t,2} = \frac{1}{\alpha^2}\right)$ -zCDP. Observe that ρ_t is a constant given a fixed value of α .

Since each data sample $(\mathbf{x}_i, y_i) \forall i \in [N]$ appears in exactly one mini-batch and the algorithm DP-AMBSSGD takes only a single pass over the entire data, by parallel composition of zCDPs, the overall ρ for DP-AMBSSGD is given by $\rho = \max_{t \in [T]} \rho_t = \frac{1}{\alpha^2}$.

Recall ρ -zCDP for an algorithm is equivalent to obtaining a $(\mu, \mu\rho)$ -Renyi differential privacy (RDP) ([Mironov, 2017](#)) guarantee. In the following, we will optimize for $\mu \in [1, \infty)$ and demonstrate that for the choice of the noise multiplier α mentioned in the theorem statement satisfies (ε, δ) -DP, which would conclude the proof. Our analysis is similar to that of Theorem 1 in [Chien et al. \(2021\)](#).

Note that $(\mu, \mu\rho)$ -RDP $\implies (\varepsilon, \delta)$ Approximate Privacy where $\varepsilon = \mu\rho + \frac{\log(1/\delta)}{\mu-1} \forall \mu > 1$. Also note that $\varepsilon_{\min} = \rho + 2\sqrt{\rho \log(1/\delta)}$ is attained at $\frac{d\varepsilon}{d\mu} = 0 \implies \mu = 1 + \sqrt{\frac{\log(1/\delta)}{\rho}}$.

Consider a fixed ε . Since we want to minimize α (which scales as $1/\sqrt{\rho}$), we need to compute the maximum permissible ρ s.t. $\varepsilon_{\min}(\rho) \leq \varepsilon$. Since $\varepsilon_{\min}(\rho)$ is an increasing function of ρ (thus an increasing function of α) and a second order polynomial in $\sqrt{\rho}$ with root at $\sqrt{\rho} = \sqrt{\log(1/\delta) + \varepsilon_{\min}} - \sqrt{\log(1/\delta)}$, the maximum is achieved at $\varepsilon_{\min}(\rho) = \varepsilon$. Therefore,

$$\frac{1}{\alpha^2} = (\sqrt{\log(1/\delta) + \varepsilon} - \sqrt{\log(1/\delta)})^2 = \frac{\varepsilon^2}{(\sqrt{\log(1/\delta) + \varepsilon} + \sqrt{\log(1/\delta)})^2}.$$

Since the above value of α satisfies (ε, δ) -DP and

$$\frac{\varepsilon^2}{(\sqrt{\log(1/\delta) + \varepsilon} + \sqrt{\log(1/\delta)})^2} \geq \frac{\varepsilon^2}{4(\log(1/\delta) + \varepsilon)},$$

choosing $\alpha \geq \frac{2\sqrt{\log(1/\delta) + \varepsilon}}{\varepsilon}$ ensures (ε, δ) -DP. ■

B.2. Proof of Risk Bounds

B.2.1. AUXILIARY TECHNICAL LEMMAS

Here we first present a few results which will be used in proving our DP-AMBSSGD risk bounds.

Lemma 20 (From Jain et al. (2018)) *For any learning rate $\eta \leq \frac{b}{R_x^2 + (b-1)\|\mathbf{H}\|}$ after N iterations of non-private fixed mini-batch SGD Algorithm with batch size b , the bias error for the final iterate and tail-averaged iterate respectively are given by:*

$$\begin{aligned}\mathcal{L}(\mathbf{w}_N^{bias}) - \mathcal{L}(\mathbf{w}^*) &\leq \frac{\kappa}{2}(1 - \eta\mu)^N (\mathcal{L}(\mathbf{w}_0) - \mathcal{L}(\mathbf{w}^*)), \\ \mathcal{L}(\bar{\mathbf{w}}_{t:N}^{bias}) - \mathcal{L}(\mathbf{w}^*) &\leq \frac{2}{\eta^2 N^2 \mu^2} (1 - \eta\mu)^t (\mathcal{L}(\mathbf{w}_0) - \mathcal{L}(\mathbf{w}^*)),\end{aligned}$$

where \mathbf{w}_0 is the initialization of \mathbf{w} and $\bar{\mathbf{w}}_{t:N} := \frac{1}{N} \sum_{t'=t}^{t+N-1} \mathbf{w}_{t'}$.

Lemma 21 (From Jain et al. (2018)) *For any learning rate $\eta \leq \frac{b}{R_x^2 + (b-1)\|\mathbf{H}\|}$ after N iterations of non-private fixed mini-batch SGD Algorithm with batch size b the variance error for the final iterate and tail-averaged iterate respectively are given by:*

$$\mathcal{L}(\mathbf{w}_N^{variance}) - \mathcal{L}(\mathbf{w}^*) \leq \frac{\eta}{2b} \text{Tr}(\mathbf{H}\mathcal{T}_b^{-1}\boldsymbol{\Sigma}), \quad L(\bar{\mathbf{w}}_{t:N}^{variance}) - \mathcal{L}(\mathbf{w}^*) \leq \frac{2}{Nb} \text{Tr}(\mathcal{T}_b^{-1}\boldsymbol{\Sigma}),$$

where $\bar{\mathbf{w}}_{t:N} := \frac{1}{N} \sum_{t'=t}^{t+N-1} \mathbf{w}_{t'}$. Further, if we define,

$$\mathbf{A} := \left(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \frac{\eta}{b} \cdot (b-1)\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}} \right)^{-1} \boldsymbol{\Sigma}, \quad \bar{\mathbf{w}}_{t:N} := \frac{1}{N} \sum_{t'=t}^{t+N-1} \mathbf{w}_{t'}, \quad \eta \leq \frac{2b}{R_x^2 \cdot \frac{d\|(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}})^{-1}\boldsymbol{\Sigma}\|_2}{\text{Tr}((\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}})^{-1}\boldsymbol{\Sigma})}},$$

then,

$$\text{Tr}(\mathcal{T}_b^{-1}\boldsymbol{\Sigma}) \leq \text{Tr}(\mathbf{A}) + \frac{\frac{\eta R_x^2}{2b} d\|(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}})^{-1}\boldsymbol{\Sigma}\|_2}{\left(1 - \frac{\eta}{2b} \cdot (R_x^2 + (b-1)\|\mathbf{H}\|_2)\right) \left(1 - \eta \frac{b-1}{2b} \|\mathbf{H}\|_2\right)} \leq 2\text{Tr}(\mathbf{H}^{-1}\boldsymbol{\Sigma}).$$

Lemma 22 *If $\eta \leq \frac{b}{R_x^2 + (b-1)\|\mathbf{H}\|}$, $\tau(t) = t \cdot (b + s)$ and $\mathbf{B}_t := \left(\mathbf{I} - \frac{\eta}{b} \sum_{i=0}^{b-1} \mathbf{x}_{\tau(t)+s+i} \mathbf{x}_{\tau(t)+s+i}^\top\right)$, then $\forall j$, $\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{B}_j^\top \mathbf{B}_j \right] \preceq \mathbf{I} - \eta \mathbf{H}$.*

Proof

$$\begin{aligned}\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{B}_j^\top \mathbf{B}_j \right] &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{I} - \frac{\eta}{b} \sum_{i=0}^{b-1} \mathbf{x}_{\tau(j)+s+i} \mathbf{x}_{\tau(j)+s+i}^\top \right) \left(\mathbf{I} - \frac{\eta}{b} \sum_{i=0}^{b-1} \mathbf{x}_{\tau(j)+s+i} \mathbf{x}_{\tau(j)+s+i}^\top \right) \right] \\ &= \mathbf{I} - \frac{2\eta}{b} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\sum_{i=0}^{b-1} \mathbf{x}_{\tau(j)+s+i} \mathbf{x}_{\tau(j)+s+i}^\top \right] \\ &\quad + \frac{\eta^2}{b^2} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\sum_{i=0}^{b-1} \sum_{m=0}^{b-1} \mathbf{x}_{\tau(j)+s+i} \mathbf{x}_{\tau(j)+s+i}^\top \mathbf{x}_{\tau(j)+s+m} \mathbf{x}_{\tau(j)+s+m}^\top \right] \\ &= \mathbf{I} - \frac{2\eta}{b} b \mathbf{H} + \frac{\eta^2}{b^2} \left(b \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\|\mathbf{x}\|^2 \mathbf{xx}^\top \right] + b(b-1) \left(\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{xx}^\top \right] \right)^2 \right) \\ &\preceq \mathbf{I} - 2\eta \mathbf{H} + \frac{\eta^2}{b} \left(R_x^2 \mathbf{H} + (b-1)\|\mathbf{H}\| \mathbf{H} \right) \preceq \mathbf{I} - \eta \mathbf{H} \left(2 - \frac{\eta}{b} (R_x^2 + (b-1)\|\mathbf{H}\|) \right).\end{aligned}\tag{40}$$

Now, $\eta \leq \frac{b}{R_x^2 + (b-1)\|\mathbf{H}\|} \implies \frac{\eta}{b}(R_x^2 + (b-1)\|\mathbf{H}\|) \leq 1$. Using this in (40) gives

$$\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{B}_j^\top \mathbf{B}_j \right] \preceq \mathbf{I} - \eta \mathbf{H}.$$

■

Lemma 23 *If $\boldsymbol{\nu}_j = \frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+i} - \frac{2\zeta_j \alpha}{b} \mathbf{g}_j$, $\boldsymbol{\xi}_{\tau(t)+s+i} := z_{\tau(t)+s+i} \mathbf{x}_{\tau(t)+s+i}$ and $\tau(t) = t \cdot (b + s)$, then $\forall j$*

$$\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\boldsymbol{\nu}_j \boldsymbol{\nu}_j^\top \right] = \frac{1}{b} \boldsymbol{\Sigma} + \frac{4\alpha^2}{b^2} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\zeta_j^2 \right] \mathbf{I}.$$

Proof The covariance of $\boldsymbol{\nu}_j$ is given by $\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\boldsymbol{\nu}_j \boldsymbol{\nu}_j^\top \right]$

$$\begin{aligned} &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+i} - \frac{2\zeta_j \alpha}{b} \mathbf{g}_j \right) \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+i} - \frac{2\zeta_j \alpha}{b} \mathbf{g}_j \right)^\top \right] \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+i} \right) \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+i} \right)^\top \right] + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\frac{4\zeta_j^2 \alpha^2}{b^2} \mathbf{g}_j \mathbf{g}_j^\top \right] \\ &\quad - \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\frac{2\zeta_j \alpha}{b^2} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+i} \right) \mathbf{g}_j^\top \right] - \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{g}_j \left(\frac{2\zeta_j \alpha}{b^2} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+i} \right)^\top \right] \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+i} \right) \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+i} \right)^\top \right] - \mathbf{0} \\ &\quad - \mathbf{0} + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\frac{4\zeta_j^2 \alpha^2}{b^2} \mathbf{g}_j \mathbf{g}_j^\top \right] \end{aligned} \tag{41}$$

$$\begin{aligned} &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+i} \right) \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+k} \right)^\top \right] + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\frac{4\zeta_j^2 \alpha^2}{b^2} \mathbf{g}_j \mathbf{g}_j^\top \right] \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\frac{1}{b^2} \sum_{i=0}^{b-1} \sum_{k=0}^{b-1} \boldsymbol{\xi}_{\tau(j)+s+i} \boldsymbol{\xi}_{\tau(j)+s+k}^\top \right] + \frac{4\alpha^2}{b^2} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\zeta_j^2 \right] \mathbf{I} \\ &= \frac{1}{b} \boldsymbol{\Sigma} + \frac{4\alpha^2}{b^2} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\zeta_j^2 \right] \mathbf{I}, \end{aligned} \tag{42}$$

where in (41), we have used the fact that $\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\boldsymbol{\xi}_{\tau(j)+s+i} \boldsymbol{\xi}_{\tau(j)+s+k}^\top \right] = \mathbf{0}$ if $k \neq i$ due to the independence of the samples in (42) and the fact that $\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{g}_j \right] = \mathbf{0}$ has been sampled independently at each step.

■

B.2.2. PROOF OF LEMMA 13

Proof Since there is no clipping, the t^{th} iteration update is given by:

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \frac{\eta}{b} \sum_{i=0}^{b-1} \mathbf{x}_{\tau(t)+s+i} (\langle \mathbf{x}_{\tau(t)+s+i}, \mathbf{w}_t \rangle - y_{\tau(t)+s+i}) - \frac{2\eta\zeta_t\alpha}{b} \mathbf{g}_t,$$

where $\tau(t) = t \cdot (b + s)$.

The update step can be written as:

$$\begin{aligned} \mathbf{w}_{t+1} - \mathbf{w}^* &= \left(\mathbf{I} - \frac{\eta}{b} \sum_{i=0}^{b-1} \mathbf{x}_{\tau(t)+s+i} \mathbf{x}_{\tau(t)+s+i}^\top \right) (\mathbf{w}_t - \mathbf{w}^*) + \frac{\eta}{b} \sum_{i=0}^{b-1} z_{\tau(t)+s+i} \mathbf{x}_{\tau(t)+s+i} - \frac{2\eta\zeta_t\alpha}{b} \mathbf{g}_t \\ \implies \boldsymbol{\omega}_{t+1} &= \left(\mathbf{I} - \frac{\eta}{b} \sum_{i=0}^{b-1} \mathbf{x}_{\tau(t)+s+i} \mathbf{x}_{\tau(t)+s+i}^\top \right) \boldsymbol{\omega}_t + \frac{\eta}{b} \sum_{i=0}^{b-1} z_{\tau(t)+s+i} \mathbf{x}_{\tau(t)+s+i} - \frac{2\eta\zeta_t\alpha}{b} \mathbf{g}_t \\ &= \left(\mathbf{I} - \frac{\eta}{b} \sum_{i=0}^{b-1} \mathbf{x}_{\tau(t)+s+i} \mathbf{x}_{\tau(t)+s+i}^\top \right) \boldsymbol{\omega}_t + \frac{\eta}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t)+s+i} - \frac{2\eta\zeta_t\alpha}{b} \mathbf{g}_t \\ &= \mathbf{B}_t \boldsymbol{\omega}_t + \eta \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t)+s+i} - \frac{2\zeta_t\alpha}{b} \mathbf{g}_t \right) \\ &= \mathbf{B}_t \boldsymbol{\omega}_t + \eta \boldsymbol{\nu}_t, \end{aligned}$$

where

$$\begin{aligned} \boldsymbol{\omega}_t &:= \mathbf{w}_t - \mathbf{w}^*, \quad \mathbf{B}_t := \left(\mathbf{I} - \frac{\eta}{b} \sum_{i=0}^{b-1} \mathbf{x}_{\tau(t)+s+i} \mathbf{x}_{\tau(t)+s+i}^\top \right), \\ \boldsymbol{\xi}_{\tau(t)+s+i} &:= z_{\tau(t)+s+i} \mathbf{x}_{\tau(t)+s+i}, \quad \boldsymbol{\nu}_t := \frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t)+s+i} - \frac{2\zeta_t\alpha}{b} \mathbf{g}_t. \end{aligned}$$

Therefore, $\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_t - \mathbf{w}^*\|_{\mathbf{H}}^2]$

$$\begin{aligned} &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left\| \mathbf{B}_{t-1} (\mathbf{w}_{t-1} - \mathbf{w}^*) + \eta \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t-1)+s+i} - \frac{2\zeta_{t-1}\alpha}{b} \mathbf{g}_{t-1} \right) \right\|_{\mathbf{H}}^2 \right] \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\|\mathbf{B}_{t-1} (\mathbf{w}_{t-1} - \mathbf{w}^*)\|_{\mathbf{H}}^2 \right] \\ &\quad + 2 \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\eta \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t-1)+s+i} - \frac{2\zeta_{t-1}\alpha}{b} \mathbf{g}_{t-1} \right)^\top \mathbf{H} (\mathbf{B}_{t-1} (\mathbf{w}_{t-1} - \mathbf{w}^*)) \right] \\ &\quad + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left\| \eta \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t-1)+s+i} - \frac{2\zeta_{t-1}\alpha}{b} \mathbf{g}_{t-1} \right) \right\|_{\mathbf{H}}^2 \right] \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\|\mathbf{B}_{t-1} (\mathbf{w}_{t-1} - \mathbf{w}^*)\|_{\mathbf{H}}^2 \right] + 0 + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left\| \eta \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t-1)+s+i} - \frac{2\zeta_{t-1}\alpha}{b} \mathbf{g}_{t-1} \right) \right\|_{\mathbf{H}}^2 \right]. \end{aligned} \tag{43}$$

That is,

$$\begin{aligned}
& \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_t - \mathbf{w}^*\|_{\mathbf{H}}^2] = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\mathbf{w}_{t-1} - \mathbf{w}^*)^\top \mathbf{B}_{t-1}^\top \mathbf{H} \mathbf{B}_{t-1} (\mathbf{w}_{t-1} - \mathbf{w}^*) \right] \\
& + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\eta^2 \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t-1)+s+i} \right)^\top \mathbf{H} \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t-1)+s+i} \right) \right] \\
& - 2 \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\eta^2 \frac{2\zeta_{t-1}\alpha}{b} \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t-1)+s+i} \right)^\top \mathbf{H} \mathbf{g}_{t-1} \right] + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\eta^2 \frac{4\zeta_{t-1}^2\alpha^2}{b^2} \mathbf{g}_{t-1}^\top \mathbf{H} \mathbf{g}_{t-1} \right] \\
& \leq (1 - \eta\mu) \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{t-1} - \mathbf{w}^*\|_{\mathbf{H}}^2] + \frac{\eta^2}{b} \text{Tr}(\mathbf{H}\boldsymbol{\Sigma}) - 0 + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\eta^2 \frac{4\zeta_{t-1}^2\alpha^2}{b^2} \right] \text{Tr}(\mathbf{H}) \quad (44) \\
& \leq (1 - \eta\mu) \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{t-1} - \mathbf{w}^*\|_{\mathbf{H}}^2] + \frac{\eta^2}{b} \text{Tr}(\mathbf{H}\boldsymbol{\Sigma}) \\
& + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\eta^2 \frac{4\alpha^2}{b^2} \cdot 3K_2^2 R_x^2 \log^{4a} N \left(\|\mathbf{H}\| \|\mathbf{w}_{t-1} - \mathbf{w}^*\|^2 + \sigma^2 + \Delta^2 \right) \right] \text{Tr}(\mathbf{H}) \quad (45) \\
& = \left(1 - \left(\eta\mu - 12\eta^2 \frac{\alpha^2}{b^2} K_2^2 R_x^2 \kappa \log^{4a} N \text{Tr}(\mathbf{H}) \right) \right) \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{t-1} - \mathbf{w}^*\|_{\mathbf{H}}^2] \\
& + \frac{\eta^2}{b} \text{Tr}(\mathbf{H}\boldsymbol{\Sigma}) + 12\eta^2 \frac{\alpha^2}{b^2} K_2^2 R_x^2 \log^{4a} N (\sigma^2 + \Delta^2) \text{Tr}(\mathbf{H}), \quad (46)
\end{aligned}$$

where in (43) and (44) we have used the fact that $\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\boldsymbol{\xi}_{\tau(t-1)+s+i}] = \mathbf{0}$, $\mathbb{E}[\mathbf{g}_{t-1}] = \mathbf{0}$ is sampled independently of all the other terms at each step and Lemmas 22 and 23. In (45) we have used the value:

$$\zeta_t \leq R_x K_2 \log^{2a} N \left(\sqrt{\|\mathbf{H}\|} \|\mathbf{w}_t - \mathbf{w}^*\| + \sigma + \Delta \right),$$

where $\Delta = \frac{\|\mathbf{w}^*\|_{\mathbf{H}} + \sigma}{N^{100}}$ as well as the fact $\mathbf{I}\mu \preceq \mathbf{H} \implies \|\mathbf{H}\| \|\mathbf{w}_{t-1} - \mathbf{w}^*\|^2 \leq \kappa \|\mathbf{w}_{t-1} - \mathbf{w}^*\|_{\mathbf{H}}^2$. Since $T \cdot (b + s) = N$, therefore if

$$\begin{aligned}
\frac{\eta\mu}{2} & \geq 12\eta^2 \frac{\alpha^2}{b^2} K_2^2 R_x^2 \kappa \log^{4a} N \text{Tr}(\mathbf{H}) \\
\implies \left(\frac{N}{T} - s \right)^2 & \geq \frac{24\eta\alpha^2 K_2^2 R_x^2 \kappa \log^{4a} N \text{Tr}(\mathbf{H})}{\mu},
\end{aligned}$$

then (46) gives $\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_t - \mathbf{w}^*\|_{\mathbf{H}}^2]$

$$\begin{aligned}
& \leq (1 - \eta\mu/2) \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{t-1} - \mathbf{w}^*\|_{\mathbf{H}}^2] + \frac{\eta^2}{b} \text{Tr}(\mathbf{H}\boldsymbol{\Sigma}) + 12\eta^2 \frac{\alpha^2}{b^2} K_2^2 R_x^2 \log^{4a} N (\sigma^2 + \Delta^2) \text{Tr}(\mathbf{H}) \\
& \leq (1 - \eta\mu/2)^t \|\mathbf{w}_0 - \mathbf{w}^*\|_{\mathbf{H}}^2 + \frac{2}{\eta\mu} \left(\frac{\eta^2}{b} \text{Tr}(\mathbf{H}\boldsymbol{\Sigma}) + 12\eta^2 \frac{\alpha^2}{b^2} K_2^2 R_x^2 \log^{4a} N (\sigma^2 + \Delta^2) \text{Tr}(\mathbf{H}) \right) \\
& \leq e^{-\eta\mu t/2} \|\mathbf{w}^*\|_{\mathbf{H}}^2 + \frac{2\eta}{\mu b} \text{Tr}(\mathbf{H}\boldsymbol{\Sigma}) + \frac{24\eta\alpha^2}{\mu b^2} K_2^2 R_x^2 \log^{4a} N (\sigma^2 + \Delta^2) \text{Tr}(\mathbf{H}).
\end{aligned}$$

Note that the Bias Term decay rate is $\eta\mu/2$. For $t = T/2 = \frac{N}{2(b+s)}$ and $T = c_1\kappa \log N$, we have

$$\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}} [\|\mathbf{w}_{T/2} - \mathbf{w}^*\|_{\mathbf{H}}^2] \leq \frac{\|\mathbf{w}^*\|_{\mathbf{H}}^2}{N^{\frac{\eta\mu}{4}c_1\kappa}} + \frac{2\eta}{\mu b} \text{Tr}(\mathbf{H}\Sigma) + \frac{24\eta\alpha^2}{\mu b^2} K_2^2 R_x^2 \log^{4a} N (\sigma^2 + \Delta^2) \text{Tr}(\mathbf{H}). \quad (47)$$

Since $\zeta_t \leq R_x K_2 \log^{2a} N \left(\sqrt{\|\mathbf{H}\|} \|\mathbf{w}_t - \mathbf{w}^*\| + \sigma + \Delta \right)$, the bound on $\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}} [\zeta_t^2]$

$$\begin{aligned} \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}} [\zeta_t^2] &\leq 3K_2^2 R_x^2 \log^{4a} N \left(\kappa \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}} [\|\mathbf{w}_t - \mathbf{w}^*\|_{\mathbf{H}}^2] + \sigma^2 + \Delta^2 \right) \\ &\leq 3K_2^2 R_x^2 \log^{4a} N \left(\kappa \left(e^{-\eta\mu/2t} \|\mathbf{w}^*\|_{\mathbf{H}}^2 + \frac{2\eta}{\mu b} \text{Tr}(\mathbf{H}\Sigma) \right) \right. \\ &\quad \left. + \frac{24\eta\alpha^2}{\mu b^2} K_2^2 R_x^2 \log^{4a} N (\sigma^2 + \Delta^2) \text{Tr}(\mathbf{H}) \right) + \sigma^2 + \Delta^2 \end{aligned} \quad (48)$$

is decreasing with t (w.p. $\geq 1 - \frac{1}{\text{Poly}(N)}$).

Furthermore, from Lemma 23,

$$\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}} [\boldsymbol{\nu}_t \boldsymbol{\nu}_t^\top] = \frac{1}{b} \Sigma + \frac{4\alpha^2}{b^2} \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}} [\zeta_t^2] \mathbf{I} \quad \forall t \quad (49)$$

Thus, if we define ζ s.t.

$$\begin{aligned} \zeta^2 &= \max\{\text{Upper-Bound}_{(\mathbf{x},y)\sim\mathcal{D}} [\zeta_{T/2}^2], \dots, \text{Upper-Bound}_{(\mathbf{x},y)\sim\mathcal{D}} [\zeta_T^2]\} \\ &= \text{Upper-Bound}_{(\mathbf{x},y)\sim\mathcal{D}} [\zeta_{T/2}^2] \\ &= 3K_2^2 R_x^2 \log^{4a} N \left(\kappa \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}} [\|\mathbf{w}_{T/2} - \mathbf{w}^*\|_{\mathbf{H}}^2] + \sigma^2 + \Delta^2 \right), \end{aligned} \quad (50)$$

where in the last step we have used (48), then using (49),

$$\begin{aligned} \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}} [\boldsymbol{\nu}_t \boldsymbol{\nu}_t^\top] &= \frac{1}{b} \Sigma + \frac{4\alpha^2}{b^2} \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}} [\zeta_t^2] \mathbf{I} \\ &\preceq \frac{1}{b} \Sigma + \frac{4\alpha^2}{b^2} \zeta^2 \mathbf{I} \quad \forall t \in \{T/2, \dots, T\}. \end{aligned} \quad (51)$$

This implies that we can perform the tail-averaged iterate analysis by restarting the algorithm with the initial value of $\mathbf{w}_0 = \mathbf{w}_{T/2}$ obtained in (47) and replacing $\zeta_{T/2}^2$ with ζ as defined in (50) for the remaining $T/2$ iterations.

Thus, now consider re-running the algorithm for $T/2$ iterations with the initialization $\mathbf{w}_0 = \mathbf{w}_{T/2}$ and constant $\zeta_t = \zeta$ defined above. The t^{th} iteration update is given by:

$$\implies \boldsymbol{\omega}_{t+1} = \mathbf{B}_t \boldsymbol{\omega}_t + \eta \left(\frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t)+s+i} - \frac{2\zeta\alpha}{b} \mathbf{g}_t \right) = \mathbf{B}_t \boldsymbol{\omega}_t + \eta \boldsymbol{\nu}_t, \quad (52)$$

where, $\boldsymbol{\nu}_t := \frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t)+s+i} - \frac{2\zeta\alpha}{b} \mathbf{g}_t$.
Expanding the recursion thus gives

$$\begin{aligned} \boldsymbol{\omega}_{t+1} &= \mathbf{B}_t \boldsymbol{\omega}_{t-1} + \eta \boldsymbol{\nu}_t = \mathbf{B}_t \mathbf{B}_{t-1} \dots \mathbf{B}_0 \boldsymbol{\omega}_0 + \eta (\boldsymbol{\nu}_t + \mathbf{B}_t \boldsymbol{\nu}_{t-1} + \dots + \mathbf{B}_t \dots \mathbf{B}_1 \boldsymbol{\nu}_0) \\ &= \mathbf{B}_t \mathbf{B}_{t-1} \dots \mathbf{B}_0 \boldsymbol{\omega}_0 + \eta (\boldsymbol{\nu}_t + \sum_{k=0}^{t-1} (\mathbf{B}_t \dots \mathbf{B}_{t-k} \boldsymbol{\nu}_{t-1-k})) \\ &:= \mathbf{Q}_{(0,t)} \boldsymbol{\omega}_0 + \eta \sum_{j=0}^t \mathbf{Q}_{(j+1,t)} \boldsymbol{\nu}_j := \boldsymbol{\omega}_{t+1}^{bias} + \boldsymbol{\omega}_{t+1}^{variance}, \end{aligned} \quad (53)$$

where, $\mathbf{Q}_{(j,t)} := \mathbf{B}_t \mathbf{B}_{t-1} \dots \mathbf{B}_j$ s.t. $\mathbf{Q}_{(j,t)} := \mathbf{I}$ if $j > t$, $\boldsymbol{\omega}_t^{bias} := \mathbf{Q}_{(0,t-1)} \boldsymbol{\omega}_0$, $\boldsymbol{\omega}_t^{variance} := \eta \sum_{j=0}^t \mathbf{Q}_{(j+1,t)} \boldsymbol{\nu}_{\ell,j}$.

Now, note that both the recursive form (52) and the expanded form (53) are the same as that obtained in equation (12) in Jain et al. (2018) which analyzed a non-private version of fixed mini-batch SGD Algorithm, except that we now have the term $\boldsymbol{\nu}_t = \frac{1}{b} \sum_{i=0}^{b-1} \boldsymbol{\xi}_{\tau(t)+s+i} - \frac{2\zeta\alpha}{b} \mathbf{g}_t$ in place of $\boldsymbol{\zeta}_{j,b}$.

Following along similar lines, we denote the tail averaged iterate as (Note the change in indexing since we are rerunning the algorithm with $\mathbf{w}_0 = \mathbf{w}_{T/2}$ for only $T/2$ iterations):

$$\bar{\mathbf{w}} := \frac{2}{T} \sum_{t'=1}^{T/2} \mathbf{w}_{t'}.$$

Then we have a similar tail averaged version of $\boldsymbol{\omega}$:

$$\begin{aligned} \bar{\boldsymbol{\omega}} &= \frac{2}{T} \sum_{t'=1}^{T/2} \boldsymbol{\omega}_{t'} = \frac{2}{T} \sum_{t'=1}^{T/2} (\boldsymbol{\omega}_{t'}^{bias} + \boldsymbol{\omega}_{t'}^{variance}) = \underbrace{\frac{2}{T} \sum_{t'=1}^{T/2} \boldsymbol{\omega}_{t'}^{bias}}_{\text{Tail Averaged Bias}} + \underbrace{\frac{2}{T} \sum_{t'=1}^{T/2} \boldsymbol{\omega}_{t'}^{variance}}_{\text{Tail Averaged Bias}} \\ &:= \bar{\boldsymbol{\omega}}^{bias} + \bar{\boldsymbol{\omega}}^{variance}. \end{aligned}$$

The overall error can then be bounded as:

$$\begin{aligned} \mathcal{L}(\mathbf{w}_t) - \mathcal{L}(\mathbf{w}^*) &= \frac{1}{2} \langle \mathbf{H}, \mathbb{E}_{(\mathbf{x},y) \sim \mathcal{D}} [\boldsymbol{\omega}_t \boldsymbol{\omega}_t^\top] \rangle = \frac{1}{2} \langle \mathbf{H}, \mathbb{E}_{(\mathbf{x},y) \sim \mathcal{D}} [(\boldsymbol{\omega}_t^{bias} + \boldsymbol{\omega}_t^{variance})(\boldsymbol{\omega}_t^{bias} + \boldsymbol{\omega}_t^{variance})^\top] \rangle \\ &\leq \langle \mathbf{H}, (\mathbb{E}_{(\mathbf{x},y) \sim \mathcal{D}} [\boldsymbol{\omega}_t^{bias} (\boldsymbol{\omega}_t^{bias})^\top] + \mathbb{E}_{(\mathbf{x},y) \sim \mathcal{D}} [\boldsymbol{\omega}_t^{variance} (\boldsymbol{\omega}_t^{variance})^\top]) \rangle \\ &= 2 \left(\frac{1}{2} \langle \mathbf{H}, \mathbb{E}_{(\mathbf{x},y) \sim \mathcal{D}} [\boldsymbol{\omega}_t^{bias} (\boldsymbol{\omega}_t^{bias})^\top] \rangle + \frac{1}{2} \langle \mathbf{H}, \mathbb{E}_{(\mathbf{x},y) \sim \mathcal{D}} [\boldsymbol{\omega}_t^{variance} (\boldsymbol{\omega}_t^{variance})^\top] \rangle \right) \\ &= 2 \left(\underbrace{(\mathcal{L}(\mathbf{w}_t^{bias}) - \mathcal{L}(\mathbf{w}^*))}_{\text{Bias Term Risk}} + \underbrace{(\mathcal{L}(\mathbf{w}_t^{variance}) - \mathcal{L}(\mathbf{w}^*))}_{\text{Variance Term Risk}} \right), \end{aligned} \quad (54)$$

where, $\mathcal{L}(\mathbf{w}_t^{bias}) - \mathcal{L}(\mathbf{w}^*) := \frac{1}{2} \langle \mathbf{H}, \mathbb{E}_{(\mathbf{x},y) \sim \mathcal{D}} [\boldsymbol{\omega}_t^{bias} (\boldsymbol{\omega}_t^{bias})^\top] \rangle$, and

$\mathcal{L}(\mathbf{w}_t^{variance}) - \mathcal{L}(\mathbf{w}^*) := \frac{1}{2} \langle \mathbf{H}, \mathbb{E}_{(\mathbf{x},y) \sim \mathcal{D}} [\boldsymbol{\omega}_t^{variance} (\boldsymbol{\omega}_t^{variance})^\top] \rangle$.

Similarly, the overall error in the tail averaged iterate will be bounded as:

$$\begin{aligned}
 \mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) &= \frac{1}{2} \langle \mathbf{H}, \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\bar{\boldsymbol{\omega}} \bar{\boldsymbol{\omega}}^\top] \rangle \\
 &= \frac{1}{2} \langle \mathbf{H}, \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\frac{2}{T} \sum_{t'=1}^{T/2} \boldsymbol{\omega}_{t'} \right) \left(\frac{2}{T} \sum_{t'=1}^{T/2} \boldsymbol{\omega}_{t'} \right)^\top \right] \rangle \leq \frac{4}{\eta T^2} \sum_{t'=1}^{T/2} \text{Tr}(\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\boldsymbol{\omega}_{t'} \boldsymbol{\omega}_{t'}^\top]) \\
 &= \frac{4}{\eta T^2} \sum_{t'=1}^{T/2} \text{Tr}(\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\mathbf{B}_{t'-1} \boldsymbol{\omega}_{t'-1} \boldsymbol{\omega}_{t'-1}^\top \mathbf{B}_{t'-1}^\top + \eta^2 \boldsymbol{\nu}_{t'-1} \boldsymbol{\nu}_{t'-1}^\top]) \\
 &= \frac{4}{\eta T^2} \sum_{t'=1}^{T/2} \text{Tr}(\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\mathbf{B}_{t-1} \boldsymbol{\omega}_{t-1} \boldsymbol{\omega}_{t-1}^\top \mathbf{B}_{t-1}^\top]) + \frac{\eta^2}{b} \boldsymbol{\Sigma} + \frac{\eta^2 \alpha^2}{b^2} \zeta^2 \mathbf{I} \tag{55} \\
 &\leq \frac{8}{\eta T^2} \sum_{t'=1}^{T/2} \text{Tr}(\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\boldsymbol{\omega}_{t'}^{bias} (\boldsymbol{\omega}_{t'}^{bias})^\top]) + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\boldsymbol{\omega}_{t'}^{variance} (\boldsymbol{\omega}_{t'}^{variance})^\top]) \\
 &:= \underbrace{\mathcal{L}(\bar{\mathbf{w}}^{bias}) - \mathcal{L}(\mathbf{w}^*)}_{\text{Tail Averaged Bias Risk}} + \underbrace{\mathcal{L}(\bar{\mathbf{w}}^{variance}) - \mathcal{L}(\mathbf{w}^*)}_{\text{Tail Averaged Variance Risk}}, \tag{56}
 \end{aligned}$$

where in (55) we have used Lemma 23 and the discussion around (51), and

$$\begin{aligned}
 \mathcal{L}(\bar{\mathbf{w}}^{bias}) - \mathcal{L}(\mathbf{w}^*) &:= \frac{8}{\eta T^2} \sum_{t'=1}^{T/2} \text{Tr}(\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\boldsymbol{\omega}_{t'}^{bias} (\boldsymbol{\omega}_{t'}^{bias})^\top]), \\
 \mathcal{L}(\bar{\mathbf{w}}^{variance}) - \mathcal{L}(\mathbf{w}^*) &:= \frac{8}{\eta T^2} \sum_{t'=1}^{T/2} \text{Tr}(\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\boldsymbol{\omega}_{t'}^{variance} (\boldsymbol{\omega}_{t'}^{variance})^\top]).
 \end{aligned}$$

Bias Term Analysis:

By virtue of the form of Bias Variance Decomposition in (54) and (56), the Bias Term error for the tail averaged iterate is the same as that of the non-private fixed mini-batch SGD case.

Therefore, using Lemmas 10 and 11 of Jain et al. (2018) in the form of Lemma 20, we get

$$\mathcal{L}(\bar{\mathbf{w}}^{bias}) - \mathcal{L}(\mathbf{w}^*) \leq \frac{8}{\eta^2 \mu^2 T^2} (1 - \eta \mu) (\mathcal{L}(\mathbf{w}_0) - \mathcal{L}(\mathbf{w}^*)) \leq \frac{8}{\eta^2 \mu^2 T^2} e^{-\eta \mu} (\mathcal{L}(\mathbf{w}_0) - \mathcal{L}(\mathbf{w}^*)). \tag{57}$$

Variance Term Analysis:

Note that the Variance Term denotes starting the initialising \mathbf{w}_0 with the ground truth \mathbf{w}^* and letting the inherent noise $\boldsymbol{\xi}$ and additive DP Gaussian noise $\frac{2\zeta\alpha}{b} \mathbf{g}$ to drive the updates. We can therefore write down the Variance Term $\boldsymbol{\omega}^{variance}$ updates as:

$$\boldsymbol{\omega}_t^{variance} = \mathbf{B}_{t-1} \boldsymbol{\omega}_{t-1}^{variance} + \eta \boldsymbol{\nu}_{t-1}$$

with $\boldsymbol{\omega}_0^{variance} = \mathbf{0}$.

Defining $\Phi_t^{\text{variance}} := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\boldsymbol{\omega}_t^{\text{variance}} (\boldsymbol{\omega}_t^{\text{variance}})^\top]$, we have

$$\begin{aligned}
\Phi_t^{\text{variance}} &:= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\boldsymbol{\omega}_t^{\text{variance}} (\boldsymbol{\omega}_t^{\text{variance}})^\top \right] \\
&= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\mathbf{B}_{t-1} \boldsymbol{\omega}_{t-1}^{\text{variance}} + \eta \boldsymbol{\nu}_{t-1}) (\mathbf{B}_{t-1} \boldsymbol{\omega}_{t-1}^{\text{variance}} + \eta \boldsymbol{\nu}_{t-1})^\top \right] \\
&= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\mathbf{B}_{t-1} \boldsymbol{\omega}_{t-1}^{\text{variance}}) (\mathbf{B}_{t-1} \boldsymbol{\omega}_{t-1}^{\text{variance}})^\top \right] + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\eta \boldsymbol{\nu}_{t-1} (\mathbf{B}_{t-1} \boldsymbol{\omega}_{t-1}^{\text{variance}})^\top \right] \\
&\quad + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\mathbf{B}_{t-1} \boldsymbol{\omega}_{t-1}^{\text{variance}}) \eta \boldsymbol{\nu}_{t-1}^\top \right] + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\eta \boldsymbol{\nu}_{t-1} \eta \boldsymbol{\nu}_{t-1}^\top \right] \\
&= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[(\mathbf{B}_{t-1} \boldsymbol{\omega}_{t-1}^{\text{variance}}) (\mathbf{B}_{t-1} \boldsymbol{\omega}_{t-1}^{\text{variance}})^\top \right] + \mathbf{0} + \mathbf{0} + \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\eta \boldsymbol{\nu}_{t-1} \eta \boldsymbol{\nu}_{t-1}^\top \right] \\
&= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{B}_{t-1} \Phi_{t-1}^{\text{variance}} \mathbf{B}_{t-1}^\top \right] + \eta^2 \left(\frac{\boldsymbol{\Sigma}}{b} + \frac{4\zeta^2 \alpha^2}{b^2} \mathbf{I} \right),
\end{aligned} \tag{58}$$

where in (58) we have again used the fact that $\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\mathbf{g}_{t-1}] = \mathbf{0}$ has been sampled independently at each step and in the last step we have used Lemma 23.

$$\begin{aligned}
&\text{Using the definition of } \mathbf{B}_{t-1} \text{ gives } \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{B}_{t-1} \Phi_{t-1}^{\text{variance}} \mathbf{B}_{t-1}^\top \right] \\
&= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\left(\mathbf{I} - \frac{\eta}{b} \sum_{i=0}^{b-1} \mathbf{x}_{\tau(t-1)+s+i} \mathbf{x}_{\tau(t-1)+s+i}^\top \right) \Phi_{t-1}^{\text{variance}} \left(\mathbf{I} - \frac{\eta}{b} \sum_{i=0}^{b-1} \mathbf{x}_{\tau(t-1)+s+i} \mathbf{x}_{\tau(t-1)+s+i}^\top \right)^\top \right] \\
&:= (\mathcal{I} - \eta \mathcal{T}_b) \Phi_{t-1}^{\text{variance}},
\end{aligned}$$

where \mathcal{T}_b represents the linear operator:

$$\mathcal{T}_b := \mathcal{H}_L + \mathcal{H}_R - \frac{\eta}{b} \mathcal{M} - \eta \frac{b-1}{b} \mathcal{H}_L \mathcal{H}_R$$

and $\mathcal{H}_L = \mathbf{H} \otimes \mathbf{I}$ and $\mathcal{H}_R = \mathbf{I} \otimes \mathbf{H}$ represent the left and right multiplication linear operators respectively.

Continuing to expand the recursion, we get

$$\begin{aligned}
\Phi_t^{\text{variance}} &= (\mathcal{I} - \eta \mathcal{T}_b) \Phi_{t-1}^{\text{variance}} + \eta^2 \left(\frac{\boldsymbol{\Sigma}}{b} + \frac{4\zeta^2 \alpha^2}{b^2} \mathbf{I} \right) \\
&= \frac{\eta^2}{b} \left(\sum_{k=0}^{t-1} (\mathcal{I} - \eta \mathcal{T}_b)^k \right) \left(\boldsymbol{\Sigma} + \frac{4\zeta^2 \alpha^2}{b} \mathbf{I} \right).
\end{aligned}$$

Furthermore, this sequence of covariances is non-decreasing w.r.t t :

$$\Phi_{t+1}^{\text{variance}} - \Phi_t^{\text{variance}} = \frac{\eta^2}{b} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\mathbf{Q}_{1,t+1} \left(\boldsymbol{\Sigma} + \frac{4\zeta^2 \alpha^2}{b} \mathbf{I} \right) \mathbf{Q}_{1,t+1}^\top \right] \geq 0.$$

By virtue of the form of the Variance Potential Decomposition above, the Variance error at iteration ℓ for the final as well as tail averaged iterate is the same as of the non-private fixed mini-batch SGD case, except that we now have $\boldsymbol{\Sigma} + \frac{4\zeta^2 \alpha^2}{b} \mathbf{I}$ in place of $\boldsymbol{\Sigma}$ of Jain et al. (2018).

Therefore the bound for the Variance Term follows almost similarly from the paper, except that Σ gets replaced with $\Sigma + \frac{4\zeta^2\alpha^2}{b}\mathbf{I}$. Using Lemmas 13, 14 and 15 of Jain et al. (2018) in the form of Lemma 21, we obtain the tail averaged iterate for our case of additive DP noise as

$$\mathcal{L}(\bar{\mathbf{w}}^{\text{variance}}) - \mathcal{L}(\mathbf{w}^*) \leq \frac{4}{Tb} \text{Tr} \left(\mathcal{T}_b^{-1} \left(\Sigma + \frac{4\zeta^2\alpha^2}{b} \mathbf{I} \right) \right),$$

such that if

$$\mathbf{A} := \left(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \frac{\eta}{b} \cdot (b-1) \mathcal{H}_{\mathcal{L}} \mathcal{H}_{\mathcal{R}} \right)^{-1} \left(\Sigma + \frac{4\zeta^2\alpha^2}{b} \mathbf{I} \right)$$

and $\eta \leq \frac{2b}{R_x^2 \frac{d \left\| (\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}})^{-1} \left(\Sigma + \frac{4\zeta^2\alpha^2}{b} \mathbf{I} \right) \right\|_2}{\text{Tr} \left((\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}})^{-1} \left(\Sigma + \frac{4\zeta^2\alpha^2}{b} \mathbf{I} \right) \right)}}$, then

$$\begin{aligned} \text{Tr} \left(\mathcal{T}_b^{-1} \left(\Sigma + \frac{4\zeta^2\alpha^2}{b} \mathbf{I} \right) \right) &\leq \text{Tr}(\mathbf{A}) + \frac{\frac{\eta R_x^2}{2b} d \left\| (\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}})^{-1} \left(\Sigma + \frac{4\zeta^2\alpha^2}{b} \mathbf{I} \right) \right\|_2}{\left(1 - \frac{\eta}{2b} \cdot (R_x^2 + (b-1) \|\mathbf{H}\|_2) \right) \left(1 - \eta \frac{b-1}{2b} \|\mathbf{H}\|_2 \right)} \\ &\leq 2 \text{Tr}(\mathbf{H}^{-1} (\Sigma + \frac{4\zeta^2\alpha^2}{b} \mathbf{I})). \end{aligned}$$

Thus we get:

$$\mathcal{L}(\bar{\mathbf{w}}^{\text{variance}}) - \mathcal{L}(\mathbf{w}^*) \leq \frac{8}{Tb} \left(\text{Tr}(\mathbf{H}^{-1} \Sigma) + \frac{4\zeta^2\alpha^2}{b} \text{Tr}(\mathbf{H}^{-1}) \right). \quad (59)$$

Combining the tail averaged iterate's Bias Error (57) and Variance Error (59) in (56), we get that at for each iteration ℓ , the overall risk for the tail averaged iterate as $\mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*)$

$$\begin{aligned} &= \underbrace{\mathcal{L}(\bar{\mathbf{w}}^{\text{bias}}) - \mathcal{L}(\mathbf{w}^*)}_{\text{Tail Averaged Bias Term Risk}} + \underbrace{\mathcal{L}(\bar{\mathbf{w}}^{\text{variance}}) - \mathcal{L}(\mathbf{w}^*)}_{\text{Tail Averaged Variance Term Risk}} \\ &\leq \frac{8}{\eta^2 \mu^2 T^2} e^{-\eta \mu} (\mathcal{L}(\mathbf{w}_0) - \mathcal{L}(\mathbf{w}^*)) + \frac{8}{Tb} \left(\text{Tr}(\mathbf{H}^{-1} \Sigma) + \frac{4\zeta^2\alpha^2}{b} \text{Tr}(\mathbf{H}^{-1}) \right) \end{aligned} \quad (60)$$

Using the value of ζ^2 from (50) and initialization $\mathbf{w}_0 = \mathbf{w}_{T/2}$ for the second run in (60) gives with probability $1 - \frac{1}{\text{Poly}(N)}$, $\mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*)$

$$\begin{aligned} &\leq \frac{4}{\eta^2 \mu^2 T^2} e^{-\eta \mu} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{T/2} - \mathbf{w}^*\|_{\mathbf{H}}^2] + \frac{8}{Tb} \left(\text{Tr}(\mathbf{H}^{-1} \Sigma) \right. \\ &\quad \left. + 3K_2^2 R_x^2 \log^{4a} N \left(\kappa \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{T/2} - \mathbf{w}^*\|_{\mathbf{H}}^2] + \sigma^2 + \Delta^2 \right) \cdot \frac{4\alpha^2}{b} \text{Tr}(\mathbf{H}^{-1}) \right) \\ &\leq \left(\frac{4}{\eta^2 \mu^2 T^2} e^{-\eta \mu} + \frac{96\alpha^2}{Tb^2} K_2^2 R_x^2 \kappa \log^{4a} N \text{Tr}(\mathbf{H}^{-1}) \right) \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{T/2} - \mathbf{w}^*\|_{\mathbf{H}}^2] \\ &\quad + \frac{8}{Tb} \left(\text{Tr}(\mathbf{H}^{-1} \Sigma) + \frac{12\alpha^2}{b} K_2^2 R_x^2 \log^{4a} N (\sigma^2 + \Delta^2) \text{Tr}(\mathbf{H}^{-1}) \right). \end{aligned} \quad (61)$$

Putting the value of $\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\|\mathbf{w}_{T/2} - \mathbf{w}^*\|_{\mathbf{H}}^2]$ from (47) in above thus gives

$$\begin{aligned} \mathcal{L}(\bar{\mathbf{w}}) - \mathcal{L}(\mathbf{w}^*) &\leq \left(\frac{4}{\eta^2 \mu^2 T^2} e^{-\eta \mu} + \frac{96\alpha^2}{Tb^2} K_2^2 R_x^2 \kappa \log^{4a} N \text{Tr}(\mathbf{H}^{-1}) \right) \\ &\quad \left(\frac{\|\mathbf{w}^*\|_{\mathbf{H}}^2}{N^{\frac{\eta \mu}{4} c_1 \kappa}} + \frac{2\eta}{\mu b} \text{Tr}(\mathbf{H}\mathbf{\Sigma}) + \frac{24\eta\alpha^2}{\mu b^2} K_2^2 R_x^2 \log^{4a} N (\sigma^2 + \Delta^2) \text{Tr}(\mathbf{H}) \right) \\ &\quad + \frac{8}{Tb} \left(\text{Tr}(\mathbf{H}^{-1}\mathbf{\Sigma}) + \frac{12\alpha^2}{b} K_2^2 R_x^2 \log^{4a} N (\sigma^2 + \Delta^2) \text{Tr}(\mathbf{H}^{-1}) \right). \end{aligned}$$

which is the required result. ■