

TOWARDS HYPERPARAMETER-FREE OPTIMIZATION WITH DIFFERENTIAL PRIVACY

Ruixuan Liu *
Emory University
ruixuan.liu2@emory.edu

Zhiqi Bu *
Amazon
woodyx218@gmail.com

ABSTRACT

Differential privacy (DP) is a privacy-preserving paradigm that protects the training data when training deep learning models. Critically, the performance of models is determined by the training hyperparameters, especially those of the learning rate schedule, thus requiring fine-grained hyperparameter tuning on the data. In practice, it is common to tune the learning rate hyperparameters through the grid search that (1) is computationally expensive as multiple runs are needed, and (2) increases the risk of data leakage as the selection of hyperparameters is data-dependent. In this work, we adapt the automatic learning rate schedule to DP optimization for any models and optimizers, so as to significantly mitigate or even eliminate the cost of hyperparameter tuning when applied together with automatic per-sample gradient clipping. Our hyperparameter-free DP optimization is almost as computationally efficient as the standard non-DP optimization, and achieves state-of-the-art DP performance on various language and vision tasks.

1 INTRODUCTION

The performance of deep learning models relies on a proper configuration of training hyperparameters. In particular, the learning rate schedule is critical to the optimization, as a large learning rate may lead to divergence, while a small learning rate may slowdown the converge too much to be useful. In practice, people have used heuristic learning rate schedules that are controlled by many hyperparameters. For example, many large language models including LLaMa2 (Touvron et al., 2023) uses linear warmup and cosine decay in its learning rate schedule, which are controlled by 3 hyperparameters. Generally speaking, hyperparameter tuning (especially for multiple hyperparameters) can be expensive for large datasets and large models.

To address this challenge, it is desirable or even necessary to determine the learning rate schedule in an adaptive and data-dependent way, without little if any manual effort. Recent advances such as D-adaptation (Defazio & Mishchenko, 2023), Prodigy (Mishchenko & Defazio, 2024), DoG (Ivgi et al., 2023), DoWG (Khaled et al., 2023), U-DoG (Kreisler et al., 2024), and GeN (Bu & Xu, 2024) have demonstrated the feasibility of automatic learning rate schedule, with some promising empirical results in deep learning (see a detailed discussion in Section 2.3).

While data-dependent learning rate are evolving in the standard non-DP regime, the adaptive data dependency can raise the privacy risks of memorizing and reproducing the training data. As an example, we consider the zeroth-order optimizer (ZO-SGD),

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_{\text{ZO-SGD}} \mathbf{z}_t$$

where \mathbf{w}_t is the model parameters, \mathbf{z}_t is a random vector that is independent of any data, and $\eta_{\text{ZO-SGD}} \approx \frac{\partial L}{\partial \mathbf{w}_t}^\top \mathbf{z}_t \in \mathbb{R}$ is the effective learning rate, with L being the loss value. In Malladi et al. (2023), ZO-SGD can effectively optimize models such as OPT 13B~60B in the few-shot and many-shot settings. Hence, by using an adaptive hyperparameter $\eta_{\text{ZO-SGD}}$, the models are able to learn and possibly memorize the training data even if the descent direction is data-independent. In fact, the private information can also leak through other hyperparameters and in other models,

*Equal contribution. This work does not relate to ZB’s position at Amazon.

where an outlier datapoint can be revealed via the membership inference attacks in support vector machines (Papernot & Steinke, 2021).

Specifically, in the regime of deep learning with differential privacy (DP), the privacy risks from the non-DP hyperparameter tuning could render the privacy guarantee non-rigorous. In practice, most work has leveraged the DP optimizers Abadi et al. (2016) to offer the privacy guarantee, where the privatization is only on the gradients, but not on the hyperparameters such as the clipping threshold R_g and the learning rate η . A common approach is to trial-and-error on multiple (R_g, η) pairs and select the best hyperparameters, as showcased by (Li et al., 2021; Kurakin et al., 2022) in Figure 8.

At high level, there are two approaches to accommodate the privacy risk in hyperparameter tuning: (1) The more explored approach is to assign a small amount of privacy budget to privatize the hyperparameter tuning. Examples include Renyi-DP tuning (Papernot & Steinke, 2021), DP-Hypo (Wang et al., 2023), DP-ZO-SGD (Tang et al., 2024; Liu et al., 2024; Zhang et al.) and many more (Liu & Talwar, 2019; Panda et al., 2024; Koskela & Kulkarni, 2023). However, these methods may suffer from worse performance due to larger DP noise addition from the reduced privacy budget, or high computation overhead. (2) The less explored approach is to adopt hyperparameter-free methods. For instance, (Bu et al., 2023b; Yang et al., 2022) replace the per-sample gradient clipping with the per-sample gradient normalization (i.e. setting $R_g \approx 0^+$), so as to remove the tuning of the hyperparameter R_g in DP optimizers.

In this work, we work towards the hyperparameter-free optimization with DP, by adapting learning-rate-free methods in DP optimization:

$$\text{DP-SGD} \implies \begin{cases} \text{Vanilla:} & \mathbf{w}_{t+1} = \mathbf{w}_t - \eta_{\text{manual}} \left(\sum_{i \in \mathcal{B}} \min\left\{\frac{R_g}{\|\mathbf{g}_i\|}, 1\right\} \mathbf{g}_i + \sigma R_g \mathbf{z} \right) \\ \text{Hyperparameter-free:} & \mathbf{w}_{t+1} = \mathbf{w}_t - \eta_{\text{GeN-DP}} \left(\sum_{i \in \mathcal{B}} \frac{\mathbf{g}_i}{\|\mathbf{g}_i\|} + \sigma \mathbf{z} \right) \end{cases} \quad (1)$$

Our contributions are summarized as follows:

1. We propose **HyFreeDP**, a hyperparameter-free DP framework that rigorously guarantees DP on the hyperparameter tuning of (R_g, η_t) , and works with any optimizer.
2. We apply the loss privatization to leverage GeN learning rate under DP, with a specific auto-regressive clipping threshold R_l that aims to minimize the clipping bias.
3. We give an end-to-end privacy accounting method that adds $< 1\%$ more gradient noise, while accurately capturing the loss curvature to determine the adaptive learning rate.
4. We show the strong performance and high efficiency of our method empirically.

2 PRELIMINARIES AND RELATED WORKS

2.1 DIFFERENTIALLY PRIVATE OPTIMIZATION

Overview of DP optimization. We aim to minimize the loss $\sum_i L(\mathbf{w}, \mathbf{x}_i)$ where $\mathbf{w} \in \mathbb{R}^d$ is the model parameters and \mathbf{x}_i is one of data points with $1 \leq i \leq N$. We denote the per-sample gradient as $\mathbf{g}_i(\mathbf{w}) := \frac{\partial L(\mathbf{w}, \mathbf{x}_i)}{\partial \mathbf{w}} \in \mathbb{R}^d$ and the mini-batch gradient at the t^{th} updating iteration as $\mathbf{m}_t \in \mathbb{R}^d$: for a batch size $B \leq N$,

$$\mathbf{m}_t(\{\mathbf{g}_i\}_{i=1}^B; R_g, \sigma_g) := \left[\sum_i^B c_i(R_g) \mathbf{g}_i + \sigma_g R_g \cdot \mathbf{z}_g \right] / B \quad (2)$$

where $\mathbf{z}_g \sim \mathcal{N}(0, \mathbf{I}_d)$ is the Gaussian noise, and $c_i(R_g) = \min(R_g / \|\mathbf{g}_i\|, 1)$ is the per-sample gradient clipping factor in (Abadi et al., 2016; De et al., 2022) with R_g being the clipping threshold.

In particular, \mathbf{m}_t reduces to the standard non-DP gradient $\sum_i \mathbf{g}_i / B$ when $\sigma_g = 0$ and R_g is sufficiently large (i.e., no clipping is applied). The clipped and perturbed batch gradient becomes the DP gradient $\mathbf{m} \equiv \mathbf{m}_{\text{DP}}$ whenever $\sigma_g > 0$, with stronger privacy guarantee for larger σ_g . On top of \mathbf{m} , models can be optimized using any optimizer such as SGD and AdamW through

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_t \cdot \mathbf{G}_t(\mathbf{m}_t(\{\mathbf{g}_i\}_{i=1}^B)) \quad (3)$$

in which \mathbf{G}_t is the post-processing such as momentum, adaptive pre-conditioning, and weight decay.

Hyper-parameters matter for DP optimization. Previous works (De et al., 2022; Li et al., 2021) reveal that the performance of DP optimization is sensitive to the hyper-parameter choices. On the one hand, DP by itself brings extra hyper-parameters, such as the gradient clipping threshold R_g , making the tuning more complex. An adaptive clipping method (Andrew et al., 2021) proposes to automatically learn R_g at each iteration, with an extra privacy budget that translates to worse accuracy. While there are methods that does not incur additional privacy budget. Automatic clipping (or per-sample normalization) (Bu et al., 2023b; Yang et al., 2022) is a technique that uses $c_i = 1/||g_i||$ to replace the R_g -dependent per-sample clipping, and thus removes the hyperparameter R_g from DP algorithms. On the other hand, hyper-parameter tuning for DP training has different patterns than non-DP training and cannot borrow previous experience on non-DP. For example, previous works (Li et al., 2021; De et al., 2022) observe that there is no benefit from decaying the learning rate during training and the optimal learning rate can be much higher than the optimal one in non-DP training.

2.2 END-TO-END DP GUARANTEE FOR OPTIMIZATION AND TUNING

However, hyper-parameter tuning enlarges privacy risks (Papernot & Steinke, 2021), and it is necessary to provide end-to-end privacy guarantee for DP. There are two existing technical paths to solve this problem:

- Multiple runs with multiple choices of hyper-parameters, and choose from these choices in a DP manner (Mohapatra et al., 2022; Papernot & Steinke, 2021; Wang et al., 2023; Liu & Talwar, 2019). This path requires some knowledge to determine the choices a priori and consumes a significant privacy budget as well as computation due to the multiple runs. For instance, (Liu & Talwar, 2019) showed that repeated searching with random iterations satisfies $(3\epsilon, 0)$ -DP if each run was $(\epsilon, 0)$ -DP.
- Making DP optimization hyper-parameter free, and only a single run is needed. For example, the automatic clipping (Bu et al., 2023b; Yang et al., 2022) eliminates the hyper-parameter R_g . Our work follows the second path, aiming to resolve tuning on η —the last critical hyper-parameter—by finding a universal configuration across datasets and tasks.

2.3 AUTOMATIC LEARNING RATE SCHEDULE

Automatic learning rate schedules (also known as learning-rate-free or parameter-free methods) have demonstrated promising performance in deep learning, with little to none manual efforts to select the learning rate. Specifically, D-adaptation (Defazio & Mishchenko, 2023), Prodigy (Mishchenko & Defazio, 2024), and DoG (Ivgi et al., 2023) (and its variants) have proposed to estimate $\eta \approx \frac{D}{G\sqrt{T}}$ where $D = ||w_0 - w_*||$ is the initialization-to-minimizer distance, G is the Lipschitz continuity constant, and T is the total number of iterations. These methods have their roots in the convergence theory under the convex and Lipschitz conditions, and may not be accurate when applied in the DP regime when the gradient is noisy. In fact, the estimation of D and G can deviate significantly from the truth when ϵ is small and the number of trainable parameters is large, i.e. the DP noise in gradient is large, as shown in Table 2.

Along an orthogonal direction, GeN (Bu & Xu, 2024) leverages the Taylor approximation of loss to set the learning rate η_{GeN} , without assuming the Lipschitz continuity or the knowledge of D . Given any descent vector \mathbf{m} , omitting the iteration index t for a brief notation, we can use the GeN learning rate in (4) to approximately minimize L :

$$\eta_{\text{GeN}}(\mathbf{m}) := \frac{\mathbf{G}^\top \mathbf{m}}{\mathbf{m}^\top \mathbf{H} \mathbf{m}} = \operatorname{argmin}_\eta L(\mathbf{w}) - \mathbf{G}^\top \mathbf{m} \eta + \frac{\mathbf{m}^\top \mathbf{H} \mathbf{m} \eta^2}{2} \approx \operatorname{argmin}_\eta L(\mathbf{w} - \eta \mathbf{m}) \quad (4)$$

in which $\mathbf{G} = \frac{\partial L}{\partial \mathbf{w}}$ is the gradient and $\mathbf{H} = \frac{\partial^2 L}{\partial \mathbf{w}^2}$ is the Hessian matrix. Notice that because L is approximated by a quadratic function, the minimizer η_{GeN} is unique and in closed form.

Numerically, η_{GeN} can be computed up to any precision by curve fitting or finite difference. Under the non-DP regime, given a series of η_i and loss values $L(\mathbf{w} - \eta_i \mathbf{m})$, Bu & Xu (2024) obtains the

numerator and denominator of η_{GeN} by solving the problem in (5):

$$\mathbf{m}^\top \mathbf{H} \mathbf{m}, \mathbf{G}^\top \mathbf{m} \approx \operatorname{argmin}_{a,b} \sum_i \left| L(\mathbf{w} - \eta_i \mathbf{m}) - \left(L(\mathbf{w}) - b\eta_i + a\frac{\eta_i^2}{2} \right) \right|^2 \quad (5)$$

Nevertheless, directly applying these non-DP automatic learning rate scheduler with DP gradient in (2) and using $\eta_{\text{GeN}}(\mathbf{m})$ will violate DP, because the learning rate estimation is obtained from forward passes on batches of private data. We defer the explanation and solution to Section 3.

3 LOSS VALUE PRIVATIZATION WITH MINIMAL CLIPPING BIAS

3.1 PRIVATIZED QUADRATIC FUNCTION

We emphasize that the learning rate $\eta_{\text{GeN}}(\mathbf{m})$ is not DP, because even though \mathbf{m} is privatized, the data is accessed without protection through \mathbf{G} and \mathbf{H} . To solve this issue, we introduce a privatized variant of (5),

$$(\mathbf{m}^\top \mathbf{H} \mathbf{m})_{\text{DP}}, (\mathbf{G}^\top \mathbf{m})_{\text{DP}} := \operatorname{argmin}_{a,b} \sum_i \left| \tilde{L}(\mathbf{w} - \eta_i \mathbf{m}_{\text{DP}}) - \left(\tilde{L}(\mathbf{w}) - b\eta_i + a\frac{\eta_i^2}{2} \right) \right|^2 \quad (6)$$

which not only replaces \mathbf{m} with the DP gradient \mathbf{m}_{DP} but also privatizes the loss by $\tilde{L}(\mathbf{w} - \eta_i \mathbf{m}_{\text{DP}})$ as in (7). The resulting learning rate is $\eta_{\text{GeN-DP}} = \frac{(\mathbf{G}^\top \mathbf{m}_{\text{DP}})_{\text{DP}}}{(\mathbf{m}^\top \mathbf{H} \mathbf{m})_{\text{DP}}}$, which is DP since every quantity in (6) is DP and because of the post-processing property. We now discuss the specifics of the loss value privatization $\tilde{L}(\mathbf{w} - \eta_i \mathbf{m}) \in \mathbb{R}$. In Table 1, we emphasize that the loss privatization is distinctively different from the gradient privatization because the loss is scalar, whereas the gradient is high-dimensional.

Table 1: Difference between the privatization of loss and gradient.

Aspects	Loss Privatization	Gradient Privatization
Dimension	1	d
Clipping Norm	L2 or L1	L2
Noise Magnitude	$\sqrt{\frac{2}{\pi}} \frac{\sigma_l R_l}{B}$	$\sqrt{d} \frac{\sigma_g R_g}{B}$
Key to Convergence	Clipping Bias	Noise Magnitude
Per-sample Operation	Clipping ($R_l \approx L$)	Normalization ($R_g \approx 0^+$)

From the perspective of per-sample clipping, the gradient is ubiquitously clipped on L2 norm, because $\|\mathbf{g}_i\|_2 \ll \|\mathbf{g}_i\|_1$ in large neural networks, and consequently a Gaussian noise is added to privatize the gradient. In contrast, we can apply L2 or L1 norm for the loss clipping, and add Gaussian (by default) or Laplacian noise to the loss, respectively.

From the perspective of noising, the expected noise magnitude for loss privatization is $\mathbb{E}|z_l| \frac{\sigma_l R_l}{B} = \sqrt{\frac{2}{\pi}} \frac{\sigma_l R_l}{B}$ where $z_l \sim N(0,1)$ is the noise on loss, and that for gradient privatization is $\mathbb{E}\|\mathbf{z}_g\| \frac{\sigma_g R_g}{B} \approx \sqrt{d} \frac{\sigma_g R_g}{B}$ where the gradient noise vector $\mathbf{z}_g \sim N(0, \mathbf{I}_d)$ by the law of large numbers. On the one hand, the gradient noise \mathbf{z}_g is large and requires small R_g to suppress the noise magnitude, as many works have use very small R (Li et al., 2021; De et al., 2022). This leads to the automatic clipping in Bu et al. (2023b) when the gradient clipping effectively becomes the gradient normalization as $R_g \rightarrow 0^+$. In fact, each per-sample gradient (as a vector) has a magnitude and a direction, and the normalization neglects some if not all magnitude information about the per-sample gradients. On the other hand, we must not use a small R_l for the loss clipping because the per-sample loss (as a scalar) only has the magnitude information. We will show by Theorem 1 that the choice of threshold R_l creates a bias-variance tradeoff between the clipping and the noising for the loss privatization:

$$\tilde{L} = \frac{1}{B} \left[\sum_i \min\left(\frac{R_l}{L_i}, 1\right) L_i + \sigma_l R_l \cdot N(0,1) \right] \quad (7)$$

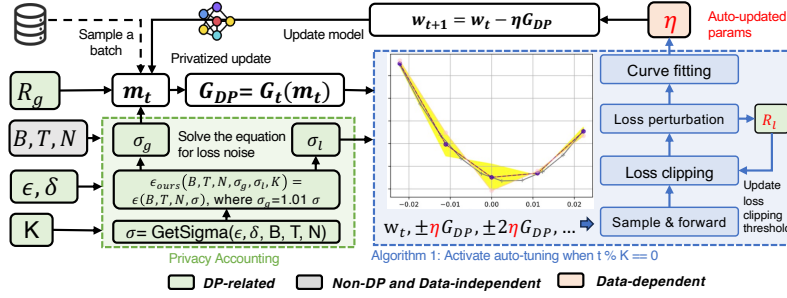


Figure 1: HyFreeDP overview with three types of hyper-parameters in the DP training. HyFreeDP saves tuning efforts via automatically tuning hyper-parameters in red text, and sets other parameters as default constants. We showcase with 5 points in curve fitting.

We note (7) is a private mean estimation that has been heavily studied in previous works (Biswas et al., 2020; Kamath et al., 2020), though many use an asymptotic threshold like $R_l + O(\log B)$, whereas Theorem 1 is more suited for our application in practice.

3.2 BIAS-VARIANCE TRADE-OFF IN LOSS PRIVATIZATION

Theorem 1. *The per-sample clipping bias of (7) is*

$$\left| \mathbb{E}(\tilde{L}) - \frac{\sum_i L_i}{B} \right| = \left| \frac{1}{B} \sum_i \min\left(\frac{R_l}{L_i}, 1\right) L_i - \frac{1}{B} \sum_i L_i \right| = \left| \frac{1}{B} \sum_i (L_i - R_l) \mathbb{I}(L_i > R_l) \right|$$

which is monotonically decreasing in R_l , and converges to $[\mathbb{E}(L_i | L_i > R_l) - R_l] \cdot \mathbb{P}(L_i > R_l)$ as $B \rightarrow \infty$. In contrast, the noise variance is $\text{Var}(\tilde{L}) = (\sigma_l R_l / B)^2$ which is increasing in R_l .

In words, a large R_l reduces the clipping bias but magnifies the noise, and vice versa for a small R_l . We propose to use $R_l \approx L$ so that the clipping bias is close to zero (i.e. \tilde{L} is approximately unbiased), and the loss noise shown in Table 1 is reasonably small for large batch size.

To put this into perspective, we give the explicit form of clipping bias when L_i follows a Gaussian distribution in Corollary 1.

Corollary 1. *Suppose $L_i \sim N(\mu, \xi^2)$, then the asymptotic clipping bias in Theorem 1 is*

$$[\mathbb{E}(L_i | L_i > R_l) - R_l] \cdot \mathbb{P}(L_i > R_l) = \xi[\phi(\alpha) - \alpha(1 - \Phi(\alpha))], \quad (8)$$

where $\alpha = \frac{R_l - \mu}{\xi}$, ϕ is the probability density function and Φ is the cumulative distribution function of standard normal distribution. The term (8) is strictly decreasing in α as well as R_l (see Figure 7).

4 ALGORITHM

4.1 HYPERPARAMETER-FREE DP OPTIMIZATION

We present our algorithm in Algorithm 1, which is DP as guaranteed in Theorem 2, almost as efficient as the standard non-DP optimization by Section 4.4, and highly accurate and fast in convergence as demonstrated in Section 5. Importantly, we have split the hyperparameters into three classes, as shown in Figure 1:

- **DP-related** hyperparameters that do not depend on the tasks, such as the gradient noise σ_g , the loss noise σ_l , and the update interval K , can be set as default constants. For example, we fix $R_g \rightarrow 0^+$ and re-scale the learning rate by $1/R_g$ according to automatic clipping and set $K = 5$.
- Training hyperparameters that are robust to different models and datasets, which we view as **data-independent**, need-not-to-search, and not violating DP, such as the batch size B and the number of iterations T . We also fix other hyperparameters not explicitly displayed

in Algorithm 1, e.g. throughout this paper, we fix the momentum coefficients and weight decay $(\beta_1, \beta_2, \text{weight_decay}) = (0.9, 0.999, 0.01)$, which is the default in Pytorch AdamW.

- Training hyperparameters that are **data-dependent**, which requires dynamical searching under DP, such as R_l and η . For these hyperparameters, Algorithm 1 adopts multiple *auto-regressive* designs, i.e. the variables to use in the t -th iteration is based on the $(t-1)$ -th iteration, which has already been privatized. These auto-regressive designs allow new variables to preserve DP by the post-processing property of DP¹.

To be specific, we have used $\tilde{L}_{t-1}^{(0)}$ as the loss clipping threshold R_l for the next iteration in Line 7, because loss values remain similar values within a few iterations; In practice, we set a more conservative loss clipping threshold $R_l = \sum \tilde{L}_{t-1}^{(k)}$ to avoid the clipping bias. We have used the previous η to construct next-loss in Line 6, which in turn will determine the new η by Line 9.

Algorithm 1 Hyperparameter-free Optimization with Differential Privacy

- 1: **INPUT:** initial $\eta=1e-4$, initial $R_l = 1$
 - 2: Forward pass to compute per-sample losses $L_{t,i}^{(0)} = L(\mathbf{w}_t, \mathbf{x}_i)$
 - 3: Compute the mini-batch loss $L_t^{(0)} = \frac{1}{B} \sum_i L_{t,i}^{(0)}$
 - 4: Back-propagate from $L_t^{(0)}$ to compute \mathbf{m}_{DP} in (2) with automatic clipping
 - 5: Post-process \mathbf{m}_{DP} by any optimizer $\mathbf{G}_{\text{DP}} := \mathbf{G}(\mathbf{m})$ in (3)
 - 6: **if** $t\%K == 0$ (e.g. $K = 10$) **then**
 - 7: Forward pass to get per-sample losses $L_{t,i}^{(\pm 1)} = L(\mathbf{w}_t \pm \eta \mathbf{G}_{\text{DP}}, \mathbf{x}_i)$
 - 8: Privatize losses $\tilde{L}_t^{(k)}$ by (7) with $R_l = \tilde{L}_{t-1}^{(0)}$ for $k \in \{-1, 0, +1\}$
 - 9: Fit the quadratic function in (6) from $\{-\eta, 0, \eta\}$ to $\{\tilde{L}_t^{(-1)}, \tilde{L}_t^{(0)}, \tilde{L}_t^{(+1)}\}$
 - 10: Extract coefficients of the fitted quadratic function $(\mathbf{m}^\top \mathbf{H} \mathbf{m})_{\text{DP}}, (\mathbf{G}^\top \mathbf{m})_{\text{DP}}$
 - 11: Update η with $\eta_{\text{GeN-DP}} = \frac{(\mathbf{G}^\top \mathbf{m})_{\text{DP}}}{(\mathbf{m}^\top \mathbf{H} \mathbf{m})_{\text{DP}}} \approx \text{argmin}_\eta \tilde{L}(\mathbf{w}_t - \eta \mathbf{G}_{\text{DP}})$
 - 12: **end if**
 - 13: Update $\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{G}_{\text{DP}}$
-




Figure 2: Impact of loss value clipping and perturbation on curve fitting along different training iterations on CIFAR100 with Vit-Small fully fine-tuning, with zero in x-axis denotes the current \mathbf{w}_t . We use 5 points for the ease of illustration and use 3 points in Algorithm 1 and experiments.

4.2 PRIVACY GUARANTEE

Theorem 2. Algorithm 1 is $(\epsilon_{\text{ours}}, \delta)$ -DP, where ϵ_{ours} depends on the batch size B , the number of iterations T , the noises (σ_g, σ_l) , and the update interval K . In contrast, vanilla DP-SGD is $(\epsilon_{\text{vanilla}}, \delta)$ -DP, where $\epsilon_{\text{vanilla}}$ depends on B, T, σ . Furthermore, we have $\epsilon_{\text{ours}}(B, T, N, \sigma_g, \sigma_l, K) > \epsilon_{\text{vanilla}}(B, T, N, \sigma)$ if $\sigma_g = \sigma$.

We omit the concrete formulae of ϵ because it depends on the choice of privacy accountants. For example, if we use μ -GDP as an asymptotic estimation, then we show in Appendix B that

$$\mu_{\text{vanilla}} = \sqrt{\left(\frac{B}{N}\right)^2 T(e^{1/\sigma^2} - 1)}, \mu_{\text{ours}} = \sqrt{\mu_{\text{vanilla}}^2 + \left(\frac{B}{N}\right)^2 \frac{3T}{K}(e^{1/\sigma_l^2} - 1)} \quad (9)$$

¹The post-processing of DP ensures that if X is (ϵ, δ) -DP then $g(X)$ is also (ϵ, δ) -DP for any function g .

which can translate into (ϵ, δ) -DP by Equation (6) in Bu et al. (2020). Note in our experiments, we use the improved RDP as the privacy accountant.

Theorem 2 and (9) show that given the same (B, T, σ) , our Algorithm 1 uses more privacy budget because we additionally privatize the loss, whereas the vanilla DP-SGD does not protect the hyperparameter tuning. To maintain the same privacy budget as vanilla DP-SGD, we use $\approx 99\%$ budget for the gradient privatization and $\approx 1\%$ budget for the loss privatization. Therefore, we need $\approx 1\%$ larger gradient noise $\sigma_g = \gamma\sigma$ and then select σ_l based on

$$\epsilon_{\text{ours}}(B, T, N, \gamma\sigma, \sigma_l, K) = \epsilon_{\text{vanilla}}(B, T, N, \sigma), \text{ where we can use } \gamma \leq 1.01. \quad (10)$$

4.3 END-TO-END NOISE DETERMINATION

We use `autoDP` library² to compute σ_l based on σ_g in (10). We give more details of our implementation in Appendix C. We visualize both noises in Figure 3 with RDP accounting and Gaussian and mechanisms for loss privatization, dashed line indicates the case when σ_g and σ_l are set equally. More examples for different mechanisms or different accounting are shown in Appendix. Note that we only adds a little more gradient noise, hence σ_g introduces negligible accuracy drop to Algorithm 1, as empirically shown in Section 5. Additionally, we demonstrate in Figure 2 that σ_l has negligible interference with the precision of estimating η_{GeN} .

4.4 EFFICIENCY OF ALGORITHM

We illustrate that Algorithm 1 can be almost as efficient as the standard non-DP optimization, in terms of training time and memory cost. We identify three orthogonal components that are absent from non-DP optimization: **(1) Gradient privatization.** DP optimization (including vanilla DP-SGD) always requires per-sample gradient clipping. Due to the high dimension of gradients, this could incur high cost in memory and time if implemented inefficiently. We directly leverage the recent advances like ghost clipping and book-keeping (BK) which have allowed DP optimization to be almost as efficient as non-DP optimization, up to 256 GPUs and 100B parameters. **(2) Loss privatization.** The cost of loss privatization alone is $O(B)$ and thus negligible, compared to the forward passes and back-propagation which are $O(Bd)$. **(3) Learning rate computation.** The cost of computing $\eta_{\text{GeN-DP}}$ in (6) mainly comes from the additional forward passes³ for $L_{t,i}^{(\pm 1)}$. Given that the back-propagation approximately costs $2\times$ the training time of forward pass, the optimizer without GeN learning rate (non-DP or DP) roughly uses 3 units of time at each iteration. In contrast, Algorithm 1 uses $3 + 2/K \approx 3.2$ units if we set $K = 10$ with $< 7\%$ overhead. We emphasize that the actual overhead to training time is even lower, because the training also includes non-optimization operations such as data loading and inter-GPU communication. In short, the efficiency gap between Algorithm 1 and non-DP optimization is negligible in practice.




Figure 3: Gradient and loss noise.

5 EXPERIMENTS

To comprehensively evaluate the effectiveness of the proposed method, we conduct experiments on both computer vision tasks and natural language tasks, across different model architectures (ViT (Yuan et al., 2021), GPT2 (Radford et al., 2019) and LLaMa2-7B (Touvron et al., 2023)) and fine-tuning paradigms (Full, BitFit (Zaken et al., 2022) and LoRA (Hu et al., 2021)). BitFit only tunes the bias terms of a pre-trained model, while LoRA only tunes the injected low-rank matrices, both keeping all other parameters frozen. We use the privacy budget $\epsilon = \{1, 3, 8\}$ with $\delta \ll N^{-1.1}$ for training dataset with N samples⁴ and also perform non-DP baseline ($\epsilon = \infty$).

²<https://github.com/yuxiangw/autodp>

³The number of $\{\eta_i\}_i$ in (6) is at least two since there are two unknown variables. More η_i may stabilize the algorithm, at cost of more forward passes, longer training time, and using more privacy budget.

⁴<https://github.com/lxuechen/private-transformers>

Table 2: Performance comparison of HyFreeDP with other baselines. We use cosine learning rate decay for NonDP-GS w/ LS in the BitFit fine-tuning. Detailed results are provided in the Appendix.

Fully Fine-Tune		Vit-Small					Vit-Base				
Privacy budget	Method	CIFAR10	CIFAR100	SVHN	GTSRB	Food101	CIFAR10	CIFAR100	SVHN	GTSRB	Food101
$\epsilon = 1$	NonDP-GS	96.49	79.70	89.86	67.00	70.38	95.20	67.18	89.10	81.11	58.80
	D-adaptation	23.74	0.80	15.27	1.86	1.17	19.15	0.83	18.56	4.72	1.48
	Prodigy	27.54	0.80	15.27	1.86	1.19	19.15	0.83	18.56	4.72	1.43
	DP-hyper	92.98	74.63	34.58	28.23	19.06	94.86	6.59	79.94	77.85	57.02
	HyFreeDP	96.36	78.17	91.15	74.68	67.98	95.76	67.17	88.24	79.00	58.16
$\epsilon = 3$	NonDP-GS	96.86	83.69	91.60	83.08	74.58	95.84	79.03	91.74	91.03	66.87
	D-adaptation	40.99	0.94	17.22	2.65	1.37	30.29	1.07	19.77	5.68	1.68
	Prodigy	77.18	0.95	18.11	2.68	1.47	30.29	1.07	19.77	5.68	1.65
	DP-hyper	95.10	78.82	48.82	42.02	36.21	95.75	19.92	87.82	90.25	66.09
	HyFreeDP	96.89	81.62	92.21	89.92	75.24	97.29	80.34	91.71	91.76	73.46
NonDP	NonDP-GS	98.33	89.41	95.58	98.66	85.19	98.88	92.40	96.87	98.41	89.61
	D-adaptation	54.29	86.35	97.09	97.38	74.75	63.37	88.41	97.02	98.65	78.35
	Prodigy	97.85	89.21	96.87	98.31	87.49	98.59	91.01	97.14	98.77	89.47
	HyFreeDP	98.32	90.84	96.86	99.00	86.38	98.83	92.58	96.75	97.37	88.73
BitFit Fine-Tune		Vit-Small					Vit-Base				
Privacy budget	Method	CIFAR10	CIFAR100	SVHN	GTSRB	Food101	CIFAR10	CIFAR100	SVHN	GTSRB	Food101
$\epsilon = 1$	NonDP-GS	96.74	84.22	90.18	86.20	77.39	97.34	84.97	90.91	87.15	76.61
	NonDP-GS w/ LS	97.40	81.36	67.95	48.64	74.55	97.90	81.87	79.86	55.66	73.76
	Prodigy	96.36	82.23	90.30	86.85	76.36	97.13	84.87	91.08	81.54	78.72
	HyFreeDP	96.20	83.84	90.49	88.08	79.27	97.42	84.01	92.21	81.83	79.25
$\epsilon = 3$	NonDP-GS	97.13	85.95	91.42	91.50	80.37	97.73	87.00	92.20	91.46	80.06
	NonDP-GS w/ LS	97.61	84.87	78.91	60.13	78.08	98.05	85.48	86.22	69.03	78.71
	Prodigy	95.72	83.63	91.90	91.64	79.20	96.96	86.26	92.13	90.27	81.98
	HyFreeDP	97.09	86.07	91.64	92.38	84.55	97.65	87.04	92.37	90.67	84.53
NonDP	NonDP-GS	97.91	89.39	91.88	95.20	86.29	98.56	91.59	94.42	92.87	88.37
	NonDP-GS w/ LS	97.89	89.40	91.98	90.31	86.27	98.56	91.61	94.43	92.90	88.39
	Prodigy	97.88	87.85	95.19	95.34	86.56	98.41	90.70	96.02	95.02	88.60
	HyFreeDP	97.97	89.86	93.60	95.19	87.24	98.43	91.69	95.48	95.23	89.06

As the first hyperparameter-free method for differentially private optimization, we compare HyFreeDP with the following baselines: 1) **NonDP-GS**: We manually perform grid search over a predefined range of learning rates, selecting the best without accumulating privacy budget across runs. This serves as the performance upper bound since tuning is non-DP. We also experiment with a manually tuned learning rate scheduler, noted as NonDP-GS w/ LS. We search the learning rate over the range $[5e-5, 1e-4, 5e-4, 1e-3, 5e-3]$ based on previous works (Bu et al., 2023b;a) to cover suitable η for various DP levels. 2) **DP-hyper** (Liu & Talwar, 2019; Wang et al., 2023; Papernot & Steinke, 2021): We simulate with a narrow range around the optimal η , spending 85% of the privacy budget for DP training with the searched η . 3) **D-Adaptation** (Defazio & Mishchenko, 2023) and **Prodigy** (Mishchenko & Defazio, 2024): Both are state-of-the-art learning rate tuning algorithms in non-DP optimization. We adopt their optimizers with recommended hyperparameters, alongside DP-specific clipping and perturbation. 4) **HyFreeDP**: We initialize $R_l = 1$ and $\eta = 1e - 4$ by default, allowing automatic updates in training.

5.1 IMAGE CLASSIFICATION TASKS

Experimental setups. As the main result shown in Table 2, we compare HyFreeDP to other baselines by experiments on CIFAR10, CIFAR100, SVHN, GTSRB and Food101 for models of Vit-Small and Vit-Base.

Evaluation results. HyFreeDP almost outperforms all end-to-end DP baselines. While non-DP automatic learning rate schedulers (D-adaptation and Prodigy) sometimes match grid-searched constants, they perform poorly in DP training, especially with tighter privacy budgets and when the number of trainable parameters is large, confirming our analysis in Section 2.3. Although DP-hyper surpasses these schedulers—likely due to our intentionally narrow search range—finding suitable ranges remains challenging as training dynamics vary by dataset and privacy budget. Even in this optimistic setting, DP-hyper underperforms due to increased gradient noise. In contrast, HyFreeDP achieves consistent performance across various ϵ values and datasets without manual learning rate tuning, thanks to our gradient noise control and efficient learning rate estimation with minimal loss value perturbation.

HyFreeDP achieves comparable or superior performance to NonDP-GS baseline without manual learning rate tuning in BitFit experiments. NonDP-GS w/ LS and Prodigy show improved stability compared to full fine-tuning, suggesting the sensitivity to training paradigms and trainable model size. In Figure 4, we illustrate training dynamics (R_l , η , loss, test accuracy) across meth-

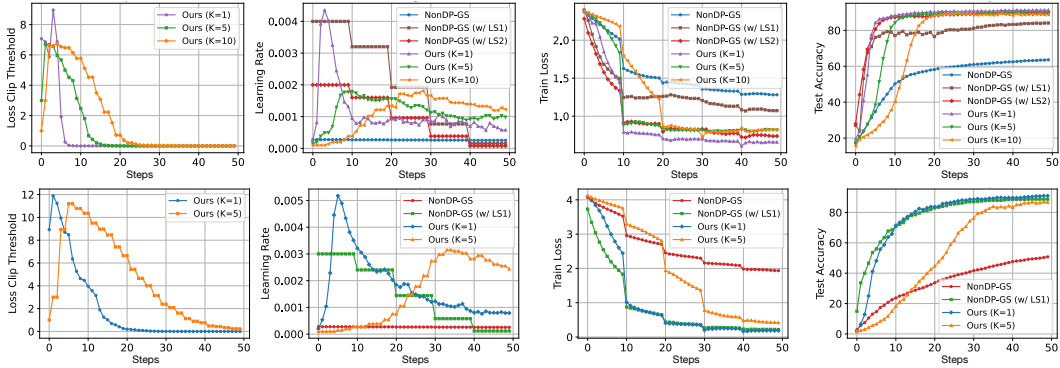


Figure 4: Automatic learning of clipping threshold, learning rate, training loss, and testing accuracy for SVHN (top) and GTSRB (bottom). HyFreeDP schedules R_l and η during training, approaching the manually tuned baseline with end-to-end DP guarantees, and is robust to varying intervals K .

ods. HyFreeDP automatically finds optimal learning rate schedules, with clipping thresholds peaking early and decreasing gradually, enabling more accurate rate estimation as training progresses. While updating with $K = 1$ yields optimal convergence, HyFreeDP remains robust to less frequent updates (e.g., $K = 5$), balancing tuning cost and convergence speed.

5.2 NATURAL LANGUAGE GENERATION TASKS

Experimental setups. We conduct experiments on E2E dataset with a table to text task on GPT-2 model, and also evaluate the language generation task with PubMed dataset by fine-tuning LLaMa2-7B model with LoRA (Hu et al., 2021) for demonstrating the scalability and generality of HyFreeDP. We follow the experimental setups based on previous works (Bu et al., 2024) and use the dataset provided by Yu et al. (2022; 2023) which contains over 75,000 abstracts of medical papers that were published after the cut-off date of LLaMa2. Based on the non-DP experience, LoRA typically requires a magnitude greater learning rate than full fine-tuning⁵, thus we scale up our default initial learning by $\times 10$. We tune the best learning rate for LLaMa2-7B with LoRA fine-tuning on 4,000 samples of PubMed for NonDP-GS when training on the full dataset.

Table 3: Performance comparison on GPT-2 for E2E dataset with different privacy budgets. Best end-to-end DP results are bolded, and results surpassing the manually tuned baseline are underlined.

Full Fine-Tune		$\epsilon = 3$					$\epsilon = 8$				
Model	Method	BLEU	CIDEr	METEOR	NIST	ROUGE.L	BLEU	CIDEr	METEOR	NIST	ROUGE.L
GPT-2	NonDP-GS	0.583	1.566	0.367	5.656	0.653	0.612	1.764	0.385	6.772	0.664
	D-Adaptation	0.000	0.000	0.003	0.082	0.016	0.000	0.000	0.000	0.000	0.000
	Prodigy	0.082	0.000	0.157	1.307	0.239	0.012	0.000	0.003	0.000	0.003
	HyFreeDP	0.585	1.564	0.365	5.736	0.636	0.612	1.768	0.378	6.702	0.655

Evaluation results. As shown in Table 3, we observe that even when the privacy budget is not small (e.g., $\epsilon = 8$), non-DP automatic learning rate scheduler does not perform well. We find that HyFreeDP consistently obtains a comparable performance as the NonDP-GS baseline without extra tuning. In Figure 5, we observe that HyFreeDP automatically discovers a learning rate schedule that achieves better generalization performance compared to the early-stopped NonDP-GS. The automatically determined learning rate (η) reveals a consistent pattern across model scales: an




Figure 5: Training dynamics of Llama2-7B on PubMed

⁵See this reference as an example. Note that LoRA may use a similar learning rate if it is set proportional to rank (Biderman et al., 2024).

“increase-then-decrease” trajectory that holds true from smaller models (Vit-Small) to larger models (LLaMa2-7B).

Table 4: Comparison of model performance in minutes with and without HyFreeDP across various datasets. The coefficients represent the ratio relative to the w/o auto configuration.

Models	Dataset	K=1	K=5	K=10	w/o HyFreeDP
Llama2-7B (LoRA-FT)	PubMed (4k)	409.750 ($\times 2.040$)	244.333 ($\times 1.217$)	222.583 ($\times 1.108$)	200.833 ($\times 1.000$)
GPT2	E2E	163.333 ($\times 1.888$)	97.167 ($\times 1.123$)	94.983 ($\times 1.098$)	86.500 ($\times 1.000$)
Vit-base	CIFAR100	152.617 ($\times 1.370$)	118.483 ($\times 1.063$)	113.317 ($\times 1.017$)	111.433 ($\times 1.000$)
Vit-base (BitFit-FT)	CIFAR100	113.450 ($\times 1.654$)	74.733 ($\times 1.089$)	73.817 ($\times 1.076$)	68.600 ($\times 1.000$)
Vit-small	SVHN	102.000 ($\times 1.255$)	84.500 ($\times 1.040$)	82.800 ($\times 1.019$)	81.250 ($\times 1.000$)

5.3 EFFICIENCY COMPARISON

Based on Section 4.4, we compare the training efficiency of HyFreeDP with a single run of DP training using the same non-DP and data-independent hyperparameters, as shown in Table 4. For LLaMa2-7B, we sample 4,000 records from PubMed and train for 3 epochs on a single A100 (80GB). For smaller datasets and models, we use previous setups on Titan RTX (24GB). HyFreeDP introduces less than $\times 2$ overhead compared to a single run of DP training, even with frequent updates at $K = 1$. Smaller models or those using LoRA or BitFit have lower additional costs, especially with $K = 1$, and the gap narrows as K increases, approaching a cost factor of $1 \times$.

6 DISCUSSION AND CONCLUSION

In conclusion, we tackle the challenge of hyperparameter tuning in differential privacy (DP) by introducing a hyperparameter-free DP training method that privately and automatically updates the learning rate. Combined with automatic clipping, our approach reduces tuning efforts and ensures end-to-end DP during training. This bridges the gap between hyperparameter-free methods in non-DP settings and DP optimization, opening promising avenues for future research.

ACKNOWLEDGMENTS

R.L. is partially supported by National Institutes of Health grant (R01LM013712), National Science Foundation grants (CNS-2124104 and CNS-2125530). We sincerely thank Prof. Li Xiong for her invaluable support and guidance throughout this work. We also appreciate the insightful comments and constructive feedback from the reviewers and the area chair.

REFERENCES

- Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308–318, 2016.
- Galen Andrew, Om Thakkar, Brendan McMahan, and Swaroop Ramaswamy. Differentially private learning with adaptive clipping. *Advances in Neural Information Processing Systems*, 34:17455–17466, 2021.
- Dan Biderman, Jacob Portes, Jose Javier Gonzalez Ortiz, Mansheej Paul, Philip Greengard, Connor Jennings, Daniel King, Sam Havens, Vitaliy Chiley, Jonathan Frankle, et al. Lora learns less and forgets less. *Transactions on Machine Learning Research*, 2024.
- Sourav Biswas, Yihe Dong, Gautam Kamath, and JU COINPRESS. Practical private mean and covariance estimation. *Preprint. Available at*, 2020.
- Zhiqi Bu and Shiyun Xu. Gradient descent with generalized newton’s method. In *The Thirteenth International Conference on Learning Representations*, 2024.
- Zhiqi Bu, Jinshuo Dong, Qi Long, and Weijie J Su. Deep learning with gaussian differential privacy. *Harvard data science review*, 2020(23), 2020.

- Zhiqi Bu, Ruixuan Liu, Yu-Xiang Wang, Sheng Zha, and George Karypis. On the accuracy and efficiency of group-wise clipping in differentially private optimization. *arXiv preprint arXiv:2310.19215*, 2023a.
- Zhiqi Bu, Yu-Xiang Wang, Sheng Zha, and George Karypis. Automatic clipping: Differentially private deep learning made easier and stronger. *Advances in Neural Information Processing Systems*, 36, 2023b.
- Zhiqi Bu, Yu-Xiang Wang, Sheng Zha, and George Karypis. Differentially private bias-term fine-tuning of foundation models. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=fqeANcjBMT>.
- Soham De, Leonard Berrada, Jamie Hayes, Samuel L Smith, and Borja Balle. Unlocking high-accuracy differentially private image classification through scale. *arXiv preprint arXiv:2204.13650*, 2022.
- Aaron Defazio and Konstantin Mishchenko. Learning-rate-free learning by d-adaptation. In *International Conference on Machine Learning*, pp. 7449–7479. PMLR, 2023.
- Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 84(1):3–37, 2022.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2021.
- Maor Ivgi, Oliver Hinder, and Yair Carmon. Dog is sgd’s best friend: A parameter-free dynamic step size schedule. In *International Conference on Machine Learning*, pp. 14465–14499. PMLR, 2023.
- Gautam Kamath, Vikrant Singhal, and Jonathan Ullman. Private mean estimation of heavy-tailed distributions. In *Conference on Learning Theory*, pp. 2204–2235. PMLR, 2020.
- Ahmed Khaled, Konstantin Mishchenko, and Chi Jin. Dowg unleashed: An efficient universal parameter-free gradient descent method. *Advances in Neural Information Processing Systems*, 36:6748–6769, 2023.
- Antti Koskela and Tejas D Kulkarni. Practical differentially private hyperparameter tuning with subsampling. *Advances in Neural Information Processing Systems*, 36:28201–28225, 2023.
- Itai Kreisler, Maor Ivgi, Oliver Hinder, and Yair Carmon. Accelerated parameter-free stochastic optimization. In *The Thirty Seventh Annual Conference on Learning Theory*, pp. 3257–3324. PMLR, 2024.
- Alexey Kurakin, Shuang Song, Steve Chien, Roxana Geambasu, Andreas Terzis, and Abhradeep Thakurta. Toward training at imagenet scale with differential privacy. *arXiv preprint arXiv:2201.12328*, 2022.
- Xuechen Li, Florian Tramèr, Percy Liang, and Tatsunori Hashimoto. Large language models can be strong differentially private learners. In *International Conference on Learning Representations*, 2021.
- Jingcheng Liu and Kunal Talwar. Private selection from private candidates. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 298–309, 2019.
- Zhihao Liu, Jian Lou, Wenjie Bao, Zhan Qin, and Kui Ren. Differentially private zeroth-order methods for scalable large language model finetuning. *arXiv preprint arXiv:2402.07818*, 2024.
- Sadhika Malladi, Tianyu Gao, Eshaan Nichani, Alex Damian, Jason D Lee, Danqi Chen, and Sanjeev Arora. Fine-tuning language models with just forward passes. *Advances in Neural Information Processing Systems*, 36:53038–53075, 2023.
- Konstantin Mishchenko and Aaron Defazio. Prodigy: an expeditiously adaptive parameter-free learner. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 35779–35804, 2024.

- Shubhankar Mohapatra, Sajin Sasy, Xi He, Gautam Kamath, and Om Thakkar. The role of adaptive optimizers for honest private hyperparameter selection. In *Proceedings of the aaai conference on artificial intelligence*, volume 36, pp. 7806–7813, 2022.
- Ashwinee Panda, Xinyu Tang, Saeed Mahloujifar, Vikash Sehwal, and Prateek Mittal. A new linear scaling rule for private adaptive hyperparameter optimization. In *International Conference on Machine Learning*, pp. 39364–39399. PMLR, 2024.
- Nicolas Papernot and Thomas Steinke. Hyperparameter tuning with renyi differential privacy. In *International Conference on Learning Representations*, 2021.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Xinyu Tang, Ashwinee Panda, Milad Nasr, Saeed Mahloujifar, and Prateek Mittal. Private fine-tuning of large language models with zeroth-order optimization. In *ICML 2024 Workshop on Foundation Models in the Wild*, 2024.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Hua Wang, Sheng Gao, Huanyu Zhang, Weijie J Su, and Milan Shen. Dp-hypo: an adaptive private hyperparameter optimization framework. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pp. 41868–41891, 2023.
- Xiaodong Yang, Huishuai Zhang, Wei Chen, and Tie-Yan Liu. Normalized/clipped sgd with perturbation for differentially private non-convex optimization. *arXiv preprint arXiv:2206.13033*, 2022.
- Da Yu, Saurabh Naik, Arturs Backurs, Sivakanth Gopi, Huseyin A Inan, Gautam Kamath, Janardhan Kulkarni, Yin Tat Lee, Andre Manoel, Lukas Wutschitz, et al. Differentially private fine-tuning of language models. In *International Conference on Learning Representations (ICLR)*, 2022.
- Da Yu, Arturs Backurs, Sivakanth Gopi, Huseyin Inan, Janardhan Kulkarni, Zinan Lin, Chulin Xie, Huishuai Zhang, and Wanrong Zhang. Training private and efficient language models with synthetic data from llms. In *Socially Responsible Language Modelling Research*, 2023.
- Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 558–567, 2021.
- Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 1–9, 2022.
- Liang Zhang, Kiran Koshy Thekumparampil, Sewoong Oh, and Niao He. Dpzero: Dimension-independent and differentially private zeroth-order optimization. In *International Workshop on Federated Learning in the Age of Foundation Models in Conjunction with NeurIPS 2023*.

A EXPERIMENTAL DETAILS

Image classification. In full fine-tuning, we tried to integrate a constant linear scheduler as a common practice, but the result does not steadily outperform the tuned NonDP-GS, so we omit the results in Appendix Table 5. We do not apply the linear scheduler to Prodigy and D-adaptation for keeping the originally recommended configuration. The results with $\epsilon = 8$ are shown in Table 7 with the consistent conclusions across different datasets.

Privacy Budget	Vit-Small					Vit-Base				
	CIFAR10	CIFAR100	SVHN	GTSRB	Food101	CIFAR10	CIFAR100	SVHN	GTSRB	Food101
$\epsilon = 1$	88.02	3.46	29.00	10.16	9.80	96.62	58.74	89.13	64.39	60.74
$\epsilon = 3$	92.48	8.00	35.54	17.68	20.40	97.11	75.73	91.79	82.31	69.09
$\epsilon = 8$	93.79	15.04	43.73	24.15	31.05	97.41	81.31	93.05	88.79	73.65
NonDP	98.00	89.11	96.55	96.94	84.55	98.88	92.51	97.27	98.29	89.10

Table 5: Full fine-tuning results by adding a linear scheduler to NonDP-GS across different datasets with privacy budgets for Vit-Small and Vit-Base models. Results show that directly integrating a constant learning rate scheduler in NonDP does not hurt performance but the DP training performance is sensitive to the learning rate scheduler.

In Table 2, we report the test accuracy averaged over 3 checkpoints. For demonstrating the stability of different methods, we present the mean accuracy with standard deviation in Table 6 for low privacy budget regime and NonDP.




Figure 6: Privacy composition for gradient and loss values privatization, with privacy accountants of RDP and GDP, and loss perturbation of Gaussian and Laplacian mechanisms. Gray dashed line indicates the naive and even budget splitting for every access to private gradient and loss, which results in larger noise magnitude on gradient especially when the adjustment is frequent. The privacy accounting strategy proposed in HyFreeDP effectively restrains the privacy budget consumption on hyper-parameter tuning and spending it wisely by only perturbing a single-dimensional loss value.

Clipping bias. Additionally, we demonstrate the loss clipping bias in Figure 7.




Figure 7: $\phi(\alpha) - \alpha(1 - \Phi(\alpha))$ which is minimized at $\alpha = \infty$.

Table 6: Averaged test accuracy with standard deviation of HyFreeDP with other baselines

Fully Fine-Tune			Vit-Small				Vit-Base				
Privacy budget	Method	CIFAR10	CIFAR100	SVHN	GTSRB	Food101	CIFAR10	CIFAR100	SVHN	GTSRB	Food101
$\epsilon = 1$	NonDP-GS	96.49 \pm 0.02	79.7 \pm 0.13	89.86 \pm 0.21	67.00 \pm 0.44	70.38 \pm 0.23	95.2 \pm 0.02	67.18 \pm 0.52	89.1 \pm 0.08	81.11 \pm 0.27	58.8 \pm 0.32
	D-adaptation	23.74 \pm 0.19	0.8 \pm 0.01	15.27 \pm 0.08	1.86 \pm 0	1.17 \pm 0.01	19.15 \pm 0.41	0.83 \pm 0.31	18.56 \pm 0.74	4.72 \pm 0.02	1.48 \pm 0.10
	Prodigy	27.54 \pm 0.61	0.80 \pm 0.01	15.27 \pm 0.08	1.86 \pm 0.00	1.19 \pm 0.00	19.15 \pm 0.41	0.83 \pm 0.00	18.56 \pm 0.04	4.72 \pm 0.02	1.43 \pm 0.00
	DP-hyper	92.98 \pm 0.14	74.63 \pm 0.09	34.58 \pm 0.58	28.23 \pm 0.91	19.06 \pm 0.59	94.86 \pm 0.10	6.59 \pm 0.00	79.94 \pm 0.04	77.85 \pm 0.16	57.02 \pm 0.00
	HyFreeDP	96.36 \pm 0.03	78.17 \pm 0.12	91.15 \pm 0.06	74.68 \pm 0.54	67.98 \pm 0.09	95.76 \pm 0.03	67.17 \pm 0.11	88.24 \pm 0.14	79.00 \pm 0.51	58.16 \pm 0.07
$\epsilon = 3$	NonDP-GS	96.86 \pm 0.03	83.69 \pm 0.12	91.6 \pm 0.11	83.08 \pm 0.69	74.58 \pm 0.23	95.84 \pm 0.07	79.03 \pm 0.18	91.74 \pm 0.04	91.03 \pm 0.11	66.87 \pm 0.06
	D-adaptation	40.99 \pm 0.92	0.94 \pm 0.01	17.22 \pm 0.99	2.65 \pm 0.03	1.37 \pm 0.01	30.29 \pm 0.51	1.07 \pm 0.01	19.77 \pm 0.04	5.68 \pm 0.00	1.68 \pm 0.01
	Prodigy	77.18 \pm 1.85	0.95 \pm 0.01	18.11 \pm 0.16	2.68 \pm 0.05	1.47 \pm 0.02	30.29 \pm 0.51	1.07 \pm 0.01	19.77 \pm 0.04	5.68 \pm 0.00	1.65 \pm 0.02
	DP-hyper	95.10 \pm 0.04	78.82 \pm 0.14	48.82 \pm 1.15	42.02 \pm 0.47	36.21 \pm 0.8	95.75 \pm 0.09	19.92 \pm 0.73	87.82 \pm 0.19	90.25 \pm 0.42	66.09 \pm 0.05
	HyFreeDP	96.89 \pm 0.06	81.62 \pm 0.06	92.21 \pm 0.12	89.92 \pm 0.57	75.24 \pm 0.06	97.29 \pm 0.02	80.34 \pm 0.05	91.71 \pm 0.18	91.76 \pm 0.48	73.46 \pm 0.06
NonDP	NonDP-GS	98.33 \pm 0.01	89.41 \pm 0.18	95.58 \pm 0.14	98.66 \pm 0.02	85.19 \pm 0.12	98.88 \pm 0.01	92.4 \pm 0.04	96.87 \pm 0.11	98.41 \pm 0.02	89.61 \pm 0.08
	D-adaptation	54.29 \pm 0.53	86.35 \pm 0.06	97.09 \pm 0.08	97.38 \pm 0.31	74.75 \pm 1.00	63.37 \pm 1.10	88.41 \pm 0.05	97.02 \pm 0.09	98.65 \pm 0.01	78.35 \pm 1.17
	Prodigy	97.85 \pm 0.03	89.21 \pm 0.07	96.87 \pm 0.07	98.31 \pm 0.14	87.49 \pm 0.08	98.59 \pm 0.06	91.01 \pm 0.10	97.14 \pm 0.06	98.77 \pm 0.14	89.47 \pm 0.06
	HyFreeDP	98.32 \pm 0.03	90.84 \pm 0.02	96.86 \pm 0.02	99.00 \pm 0.02	86.38 \pm 0.03	98.83 \pm 0.01	92.58 \pm 0.02	96.75 \pm 0.05	97.37 \pm 0.05	88.73 \pm 0.02

Table 7: Performance comparison of HyFreeDP to other baselines. We use cosine learning rate decay for NonDP-GS w/ LS baseline in the BitFit fine-tuning setting.

Full Fine-Tune			Vit-Small				Vit-Base				
Privacy budget	Method	CIFAR10	CIFAR100	SVHN	GTSRB	Food101	CIFAR10	CIFAR100	SVHN	GTSRB	Food101
$\epsilon = 8$	NonDP-GS	96.28	84.99	92.53	89.97	77.08	96.14	82.52	92.38	94.91	71.05
	D-adaptation	78.31	1.07	19.10	3.67	1.76	40.90	1.25	20.86	6.84	1.94
	Prodigy	95.74	1.29	20.75	4.86	2.92	45.89	1.25	20.86	6.84	1.90
	DP-hyper	95.70	80.58	64.72	50.10	49.18	96.28	36.46	90.27	94.62	70.63
	HyFreeDP	97.04	86.00	95.15	88.06	80.61	97.79	87.57	95.00	94.07	81.91
BitFit Fine-Tune			Vit-Small				Vit-Base				
Privacy budget	Method	CIFAR10	CIFAR100	SVHN	GTSRB	Food101	CIFAR10	CIFAR100	SVHN	GTSRB	Food101
$\epsilon = 8$	NonDP-GS	97.23	86.56	92.23	93.34	82.12	97.81	88.05	93.38	93.04	81.88
	NonDP-GS w/ LS	97.63	86.21	83.17	67.35	79.85	98.08	87.17	88.25	75.61	81.18
	HyFreeDP	97.31	88.43	89.17	93.54	74.61	97.90	90.34	92.36	93.14	87.14

B PROOFS

Proof of Theorem 1. It’s not hard to see

$$\begin{aligned} \mathbb{E}(\tilde{L}) - \frac{\sum_i L_i}{B} &= \frac{1}{B} \sum_i \min\left(\frac{R_i}{L_i}, 1\right) L_i - \frac{1}{B} \sum_i L_i \\ &= \frac{1}{B} \sum_i (R_i - L_i) \mathbb{I}(L_i > R_i) = -\frac{1}{B} \sum_i \text{ReLU}(L_i - R_i) \end{aligned}$$

in which the non-negative $\text{ReLU}(x) = x \cdot \mathbb{I}(x > 0)$. Taking the absolute value, we have

$$\left| \mathbb{E}(\tilde{L}) - \frac{\sum_i L_i}{B} \right| = \frac{1}{B} \sum_i \text{ReLU}(L_i - R_i) = \left| \frac{1}{B} \sum_i (L_i - R_i) \mathbb{I}(L_i > R_i) \right|$$

which is decreasing in R_i because ReLU is increasing in its input $(L_i - R_i)$, and this input is decreasing in R_i .

As $B \rightarrow \infty$, the clipping bias tends to

$$\begin{aligned} \mathbb{E}(\text{ReLU}(L_i - R_i)) &= \mathbb{E}(\text{ReLU}(L_i - R_i) | L_i > R_i) \cdot \mathbb{P}(L_i > R_i) \\ &= \mathbb{E}((L_i - R_i) | L_i > R_i) \cdot \mathbb{P}(L_i > R_i) = [\mathbb{E}(L_i | L_i > R_i) - R_i] \cdot \mathbb{P}(L_i > R_i) \end{aligned}$$

where we have used $\text{ReLU}(x) = x$ when $x > 0$. \square

Proof of Corollary 1. The key part in (8) is $\mathbb{E}(L_i | L_i > R_i)$, which is the expectation of the truncated normal distribution by one-side truncation. It is known that for $\alpha = \frac{R_i - \mu}{\xi}$,

$$\mathbb{E}(L_i | L_i > R_i) = \mu + \xi \frac{\phi(\alpha)}{1 - \Phi(\alpha)}, \quad \mathbb{P}(L_i > R_i) = 1 - \Phi(\alpha)$$

The proof is complete by inserting these quantities. \square

Proof of Theorem 2. All privacy budget of Algorithm 1 goes into two components: privatizing the gradient (with noise level σ_g) and privatizing the loss (with noise level σ_l).

Under the same (B, T, N, σ_g) , we have T mechanisms of gradient privatization, each of (ϵ_g, δ_g) -DP and $3T/K$ mechanisms of loss privatization, each of (ϵ_l, δ_l) -DP. Hence it is clear that $\epsilon_{\text{ours}} > \epsilon_{\text{vanilla}}$.

To be more specific, we demonstrate with μ -GDP. The vanilla DP-SGD is μ -GDP with

$$\mu_{\text{vanilla}} = \frac{B}{N} \sqrt{T(e^{1/\sigma_g^2} - 1)}$$

which is the same as the gradient privatization component of our DP-SGD. We additionally spend

$$\mu_l = \frac{B}{N} \sqrt{\frac{3T}{K}(e^{1/\sigma_l^2} - 1)}$$

leading to a total budget of

$$\mu_{\text{ours}} = \sqrt{\mu_{\text{vanilla}}^2 + \mu_l^2}$$

by Corollary 3.3 in Dong et al. (2022). It is clear $\mu_{\text{ours}} > \mu_{\text{vanilla}}$. \square

C END-TO-END PRIVACY ACCOUNTING AND INVERSE

We demonstrate how to determine (σ_g, σ_l) given (ϵ, δ) -DP budget. In vanilla DP optimization, we can leverage privacy accountants such as RDP, GDP, PRV, etc. Each accountant is a function whose input is hyperparameters $(B, T, N, \delta, \sigma)$ and the output is ϵ (see an example in (9)). In this section, we denote any accountant as f , so that

$$f(\sigma; B, N, \delta) = \epsilon' \quad (11)$$

$$f^T(\sigma; B, N, \delta) = \epsilon \quad (12)$$

$$f^{-T}(\epsilon; B, N, \delta) = \sigma \quad (13)$$

in which ϵ' is the single-iteration budget, f^T means a composition of T iterations, and f^{-T} is the inverse function known as `GetSigma` in Figure 1.

In this work, we develop an end-to-end privacy accountant to compose both the gradient privatization and the loss privatization. Our accountant takes the input $(B, T, N, \delta, \sigma_l, \gamma, K)$, where $\gamma = 1.01$ by default and can take smaller value for larger models or smaller (ϵ, δ) budget. Therefore, $\sigma_g = \gamma\sigma$.

Firstly, we call $f^T(\gamma\sigma; B, N, \delta) = \hat{\epsilon}$ to get the reference budget $\hat{\epsilon}$ which is strictly smaller than ϵ because f is monotonically decreasing in its input. Then we guess the loss noise and call $f^{3T/K}(\sigma_l; B, N, \delta) = \epsilon_l$ since there are $3T/K$ rounds of loss privatization. We continue our guess until

$$f^{3T/K}(\sigma_l; B, N, \delta) + f^T(\gamma\sigma; B, N, \delta) = \epsilon_l + \hat{\epsilon} = \epsilon$$

Notice the left hand side is monotonically decreasing in σ_l . Hence we use bisection method to find the unique solution σ_l , at an exponentially fast speed.

Algorithm 2 End-to-End Privacy Accounting

- 1: **INPUT:** The end-to-end DP budget (ϵ, δ) , `GetSigma`(\cdot), `Compose`(\cdot), `Solve`(\cdot)
 - 2: **OUTPUT:** Noise magnitude for gradient and loss privatization σ_g and σ_l
 - 3: \triangleright Compute the gradient noise scale σ by assuming there is only a single training run
 - 4: $\sigma = \text{GetSigma}(\epsilon, \delta, B, T, N)$
 - 5: \triangleright Compute the σ_g in Algorithm 1 with a controlled noise increase
 - 6: $\sigma_g = \gamma \cdot \sigma$ with the constant γ slightly greater than 1 (e.g, $\gamma = 1.01$)
 - 7: \triangleright Define the composition function with input variable as the loss noise scale c
 - 8: $\epsilon_{\text{ours}}(c|\sigma_g, \delta, B, N, T, K) = \text{Compose}(f_g^T(\sigma_g; B, N, \delta), f_l^{3T/K}(c; B, N, \delta))$
 - 9: \triangleright Solve the minimization of the scalar function respect to c
 - 10: $\sigma_l = \text{Solve}(c, \epsilon) = \arg \min_c |\epsilon_{\text{ours}}(c) - \epsilon|$
-

In the above, we have used the functions `GetSigma`(\cdot)⁶, `Compose`(\cdot)⁷, and any root-finding method `Solve`(\cdot) such as the bi-section in `scipy` library. We highlight that Algorithm 1 and

⁶https://github.com/yuxiangw/autodp/blob/master/example/example_calibrator.py

⁷https://github.com/yuxiangw/autodp/blob/master/example/example_composition.py

Algorithm 2 are sufficiently flexible to work with the general DP notions, including GDP, Renyi DP, tCDP, per-instance DP, per-user DP, etc.

D MISC

Hyper-parameter matters for DP optimization. Previous works De et al. (2022); Li et al. (2021) reveal that the performance of DP optimization is sensitive to the hyper-parameter choices, as we cited in Figure 8.




Figure 8: Hyperparameter tuning of (R_g, η) cited from Figure 1 of Bu et al. (2023b). Here R_g is the clipping norm and η is the learning rate. Left: BLEU score of GPT2 on E2E dataset Li et al. (2021). Right: Test accuracy of ResNet18 on ImageNet Kurakin et al. (2022).

End-to-end DP guarantee for optimization and tuning. As shown in Table 8, we compare our work with representative works that try to ensure end-to-end DP for both optimization and tuning.

Table 8: Comparison of hyperparameter search strategies. Budget splitting percentages for Adaptive Clipping and DP-hyper are estimated values, as the actual percentages depend on specific datasets.

Method	Searching η	Searching R_g	% Budget on Hyperparam
Vanilla Abadi et al. (2016)	✓	✓	0%
Automatic Clipping Bu et al. (2023b)	✓	✗	0%
Adaptive Clipping Andrew et al. (2021)	✓	✗	≈20%
DP-hyper Papernot & Steinke (2021)	✓	✗	≈20%
Ours (this work)	✗	✗	<1%