A TRAINABLE OPTIMIZER

Anonymous authors
Paper under double-blind review

ABSTRACT

The concept of learning to optimize involves utilizing a trainable optimization strategy rather than relying on manually defined full gradient estimations such as ADAM. We present a framework that jointly trains the full gradient estimator and the trainable weights of the model. Specifically, we prove that pseudo-linear TO (Trainable Optimizer), a linear approximation of the full gradient, matches SGD's convergence rate while effectively reducing variance. Pseudo-linear TO incurs negligible computational overhead, requiring only minimal additional tensor multiplications. To further improve computational efficiency, we introduce two simplified variants of Pseudo-linear TO. Experiments demonstrate that TO methods converge faster than benchmark algorithms (e.g., ADAM) in both strongly convex and non-convex settings, and fine tuning of an LLM.

1 INTRODUCTION

Gradient-based stochastic optimization methods are fundamental to solving many machine learning problems. Given a set of loss functions $\{f_n(w)\}_{n=1}^N$ where $w \in \mathbb{R}^d$, the full loss is defined as $F(w) = \frac{1}{N}\sum_{i=1}^N f_n(w)$. For a subset $\mathcal{B} \subset \{1,\dots,N\}$, the mini-batch loss is given by $F^{\mathcal{B}}(w) := \frac{1}{|\mathcal{B}|}\sum_{n\in\mathcal{B}} f_n(w)$. Let w^* be the global minimum of $F(w)$. Popular approaches, such as Stochastic Gradient Descent (SGD), RMSProp (Hinton et al., 2012), Adagrad (Duchi et al., 2011), and ADAM (Kingma & Ba, 2015), can be unified under a general framework. In such a framework, the update direction depends on past iterations $w_1,\dots,w_t$, past stochastic gradients $g_1,\dots,g_t$ and the optimizer's variables θ_t, following the update rule

$$w_{t+1} = w_t - \gamma_t \widehat{G}_t,$$

where γ_t is the learning rate and $\widehat{G}_t$ denotes the update direction.

Learning to Optimize (L2O) automates the development of optimization strategies by replacing traditional, manually designed $\widehat{G}_t$ functions with learned counterparts (Chen et al., 2022). The core premise of L2O is to employ machine learning models, typically neural networks, to discover task-specific optimization policies that outperform conventional handcrafted rules. Unlike classical optimizers with fixed update rules and hyperparameters, L2O models are meta-trained across diverse problems to generalize to unseen tasks, making them particularly effective for complex, high-dimensional spaces such as neural architecture search and scientific computing. L2O models contain an offline procedure, where the learnable optimizer is trained on a set of similar tasks, and an online procedure, where the trained optimizer optimizes a new unseen task from the same task distribution.

Generally, an L2O approach parameterizes the update direction $\widehat{G}_t$ as a function of past values $w_1,\dots,w_t$ of trainable model weights and past gradients $g_1,\dots,g_t$ with optimizer variable θ being fixed. For a given task, θ is updated at the end of the optimization procedure, based on the performance of the optimizer on that task. With a learned θ, an unseen task is further handled with the following update rule for the model weights

$$w_{t+1} = w_t - \gamma_t \widehat{G}_t(w_1,\dots,w_t; g_1,\dots,g_t, \theta).$$

Practically, for a given new task, similar tasks are not necessarily available, which makes the offline procedure of L2O not possible. We propose a departure from standard L2O methods—which train $\widehat{G}_t$ across broad problem distributions—by instead co-training the optimizer and weights of the model

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Algorithm 1 A General Trainable Optimizer

Require: Learning rates β_t, γ_t.
1: Initialize w_1, A_0 and b_0.
2: **for** $t = 1, \ldots, T$ **do**
3: Sample $\mathcal{B}_t \subset \{1, \ldots, N\}$ such that $|\mathcal{B}_t| = b$
4: $g_t \leftarrow \nabla F^{\mathcal{B}_t}(w_t)$
5: $\theta_{t+1} = \theta_t - \beta_t \nabla_{\theta_t} l(\theta_t; w_1, \ldots, w_t; g_1, \ldots, g_t)$
6: $w_{t+1} = w_t - \gamma_t \widehat{G}_t(w_1, \ldots, w_t; g_1, \ldots, g_t; \theta_t)$
7: **end for**

for a single task. Our new approach replaces the fixed multidimensional optimizer variable $\theta \in \mathbb{R}^{\bar{d}}$ with θ_t, which is updated along with w_t simultaneously at each iteration.

Algorithm 1 displays the general framework of our approach, Trainable Optimizer (TO). The trainable weights of the model are updated as

$$w_{t+1} = w_t - \gamma_t \widehat{G}_t(w_1, \ldots, w_t; g_1, \ldots, g_t; \theta_t). \tag{1}$$

Optimizer variable θ_t is updated based on the gradient of an approximation loss function $l(\theta_t; w_1, \ldots, w_t; g_1, \ldots, g_t)$, following the update rule:

$$\theta_{t+1} = \theta_t - \beta_t \nabla_{\theta_t} l(\theta_t; w_1, \ldots, w_t; g_1, \ldots, g_t). \tag{2}$$

Despite L2O's empirical success in accelerating convergence, its theoretical guarantees remain unstudied. We establish that for specific parameterizations of $\widehat{G}_t$ and setting of l to be an L_2 error with the ground truth being the gradient of F,

$$l(\theta_t; w_1, \ldots, w_t; g_1 \ldots, g_t) = \frac{1}{2} \left\| g_t - \widehat{G}_t(w_1, \ldots, w_t; g_1, \ldots, g_t, \theta_t) \right\|_2^2,$$

both the gradient approximation error $\left\| \widehat{G}_t - \nabla F(w_t) \right\|$ and the optimality gap $\|w_t - w^*\|$ converge to zero simultaneously.

The relationship between convergence and variance reduction in gradient-based methods has been well-studied. Johnson & Zhang (2013b) proposed SVRG, achieving variance reduction through snapshot gradients and bias correction via full-batch gradients, guaranteeing exponential convergence for strongly convex losses. Zaheer et al. (2018) proved ADAM's convergence under growing mini-batch sizes, while Qian & Klabjan (2020) and Wang & Klabjan (2025) further established a crucial link between variance reduction and convergence. A fundamental criterion for analyzing stochastic optimization methods involves verifying whether the estimation variance vanishes as $t \to \infty$. Our framework inherently satisfies this property, constituting a trainable variance reduction algorithm.

In this work, we focus on pseudo-linear approximations, where $\widehat{G}_t$ is formulated as a linear function with regards to w_t, i.e., $\widehat{G}_t = A_t w_t + b_t$ with optimizer variables $A_t \in \mathbb{R}^{d \times d}$, $b_t \in \mathbb{R}^d$, and $\theta_t = (A_t, b_t)$. We call it pseudo since A_t and b_t are functions of $w_1, \ldots, w_{t-1}$ and $g_1, \ldots, g_{t-1}$. This pseudo-linear parameterization is motivated by two key considerations. First, it draws inspiration from Taylor expansion principles where lower-order terms often provide effective approximations. Second, it enables efficient computation, as the gradient of the approximation loss can be calculated explicitly by simple tensor multiplications, which requires negligible additional computational time.

Our theoretical analysis establishes that for strongly convex loss functions, the parameters w_t under the pseudo-linearly approximated gradients converge to the optimal point w^* at the rate $\mathbb{E}\left[\|w_t - w^*\|_2^2\right] \leq \mathcal{O}(1/t)$. Furthermore, we show that the approximation variance $\mathbb{E}\left[\left\| \widehat{G}_t - \nabla F(w_t) \right\|_2^2\right]$ also converges to zero at the same $\mathcal{O}(1/t)$ rate. Comparing with the variance of SGD and ADAM, which does not converge to zero as $t \to \infty$, these results confirm that our trained optimizer functions as an effective variance reduction algorithm.

It is worth highlighting the connection and distinction between the pseudo-linear TO and quasi-Newton methods. While both leverage linear structures to approximate gradient information, classical

quasi-Newton methods rely on deterministic full-gradient settings and construct curvature approximations (e.g., inverse Hessians), with theoretical guarantees that do not directly extend to stochastic regimes. In contrast, our framework is designed for stochastic optimization, where full gradients are usually practically unavailable and only mini-batch estimates can be used, making it more suitable for large-scale machine learning tasks. Furthermore, we emphasize that pseudo-linear TO is a specific instance of a broader framework: our main contribution is a general trainable optimizer formulation, where different choices of the approximation loss l and estimator $\widehat{G}_t$ can be incorporated beyond the linear case.

To address computational concerns, we propose two simplified variants of pseudo-linear TO: Diagonal TO restricts A_t to diagonal matrices and RankOne TO uses rank-one matrices for A_t. The two variants significantly reduce the number of variables in the optimizer, which are more suitable with the practical cases with limited memory. Notably, we prove that momentum-based SGD emerges as a special case of pseudo-linear TO when appropriate learning rates and initializations A_0 and b_0 are selected.

Our comprehensive experimental evaluation compares three TO variants (Pseudo-linear (Full-TO), Diagonal, RankOne) across diverse machine learning tasks spanning three categories: 1. Strongly Convex losses (e.g., L_2-regularized logistic regression), 2. Convex but not strongly convex losses (e.g., logistic regression without regularization), and 3. Non-convex losses (e.g., ResNet classification, Llama fine-tuning). The results demonstrate that TO significantly outperforms ADAM in both strongly convex and complex non-convex settings, while matching ADAM's performance on standard convex tasks. Moreover, our method offers practical advantages over quasi-Newton implementations such as L-BFGS (Liu & Nocedal, 1989) in terms of training loss under mini-batch settings.

Our main contributions are as follows. 1. We propose a new framework that updates the optimizer variables simultaneously with trainable weights of the model, based on (1) and (2). 2. We show the $\mathcal{O}(1/t)$ convergence of Pseudo-linear TO for strongly convex losses. 3. We propose two variants of pseudo-linear TO: Diagonal TO and RankOne TO. They reduce the memory requirement of pseudo-linear TO and are better suitable for large models. 4. We show by means of experiments that TO outperforms ADAM on strongly convex problems and complex non-convex problems.

In Section 2, we review related work on trainable optimizers and variance-reduction optimization methods. Section 3 presents the TO framework with linear approximation and demonstrates the convergence of model parameters to the optimal point while reducing the approximation variance to zero. This section also introduces two simplified versions—Diagonal TO and RankOne TO. In Section 4, we present numerical experiments illustrating the improved convergence rates of the proposed methods.

2 Related Work

2.1 Learning to Optimize

The L2O paradigm represents a significant shift from traditional optimization methods, replacing theoretically-derived update rules with data-driven approaches that leverage machine learning. Pioneering work by Andrychowicz et al. (2016) demonstrated that recurrent neural networks (RNNs) can learn effective optimization strategies that outperform conventional methods like SGD. Subsequent advances incorporated reinforcement learning (RL) frameworks (Li & Malik, 2017; Bello et al., 2017), enabling the discovery of specialized optimization algorithms for diverse applications ranging from neural architecture search to black-box optimization. Our work extends these foundations while introducing several key innovations. First, unlike standard L2O approaches that solely optimize the target model's parameters, our framework simultaneously optimizes both the model parameters and the optimization strategy itself. This co-optimization enables dynamic adaptation to problem-specific characteristics. Second, while most L2O methods lack theoretical guarantees, we provide rigorous convergence proofs showing simultaneous convergence of both trainable weights of the model and the gradient approximation error. By employing a carefully designed linear approximation, our approach achieves both computational efficiency and theoretical soundness.

2.2 VARIANCE REDUCTION

Variance reduction has emerged as a fundamental technique for improving stability and the convergence rate of stochastic optimization methods. One particularly influential method is Stochastic Variance Reduced Gradient (SVRG), originally introduced by Johnson & Zhang (2013a). SVRG operates by maintaining a snapshot model to effectively correct the noisy estimates obtained from stochastic gradients, thereby making the entire optimization process both faster and more numerically stable. Another important approach is SAGA, proposed by Defazio et al. (2014), which builds upon SVRG's foundation while streamlining the process through an efficient history of gradients, making it more practical for large-scale implementations. Recent work includes the introduction of Variance Reduced Adam (VRADAM) by Wang & Klabjan (2025), which combines variance reduction techniques with ADAM. This hybrid approach successfully resolves ADAM's well-known divergence issues while carefully preserving its advantages in training speed and the convergence rate for challenging deep learning tasks. While other recent works, such as Huang et al. (2022) and Cutkosky & Orabona (2019), have demonstrated various approaches combining variance reduction techniques with modern optimization algorithms, our work makes distinct contributions by specifically focusing on utilizing the idea of trainable optimizers for variance reduction. We provide theoretical and empirical evidence showing that the gradient variance can be systematically reduced through the optimization process itself as the trainable optimizer variables are progressively refined during training.

3 TRAINABLE OPTIMIZER WITH PSEUDO-LINEAR APPROXIMATION

3.1 THE PSEUDO-LINEAR TO ALGORITHM

In this section, we present our approach for approximating the full gradient with a trainable pseudo-linear function. Consider a stationary point w^* of F, where we assume that F is second-order differentiable at w^*. Drawing inspiration from the Taylor expansion of $\nabla F(w)$ centered at w^*, we have $\nabla F(w) \approx \nabla F(w^*) + \nabla^2 F(w^*)(w - w^*) = \nabla^2 F(w^*)(w - w^*)$, which yields a linear function of w. We formulate the approximated gradient as $\widehat{G}_t = A_t w_t + b_t$, where A_t and b_t represent the variables of the optimizer. The corresponding approximation loss is defined by $l((A_t, b_t); w_t; g_t) := \frac{1}{2} \left\| g_t - A_t w_t - b_t \right\|_2^2$, with gradients computed as

$$\begin{aligned} \nabla_{A_t} l((A_t, b_t); w_t; g_t) &= -(g_t - A_t w_t - b_t) w_t^\top \\ \nabla_{b_t} l((A_t, b_t); w_t; g_t) &= -(g_t - A_t w_t - b_t). \end{aligned}$$

We introduced the Trainable Optimizer (TO) with pseudo-linear approximation in Algorithm 2. Notably, lines 5 and 6 of the algorithm implement updates to the parameters A and b through the gradient descent on the loss function $l((A_t, b_t); w_t; g_t)$.

The pseudo-linear TO framework admits an important interpretation as a generalized momentum-based SGD algorithm. This connection becomes exact under specific initialization conditions: when we set $A_0 = 0, b_0 = 0, \alpha_t = 0$ and maintain a constant $\beta_t = \beta$ for some $0 < \beta < 1$, the update rules in Algorithm 2 are simplified to $A_t = 0, b_t = \beta g_t + (1 - \beta) b_{t-1}, \widehat{G}_t = b_t$, which recovers precisely the standard momentum-based SGD formulation. The initialization of parameters A_0 and b_0 presents practical challenges. However, based on our theoretical analysis and empirical observations, we strongly recommend initializing both to zero, as this initialization strategy provides effective practical performance during the initial training phase.

3.2 ANALYSES

In this section we theoretically prove the convergence of Pseudo-linear TO under strongly convex loss. Theoretical analyses of gradient-based algorithms often rely on assumptions of bounded iterates and gradients. To this end, we modify Algorithm 2 by changing line 8 to $w_{t+1} \leftarrow \Pi_{\mathcal{F}}(w_t - \gamma_t \widehat{G}_t)$. The projection operator $\Pi_{\mathcal{F}}(w)$, which maps vectors to the bounded feasible set $\mathcal{F}$, is mathematically defined as $\Pi_{\mathcal{F}}(w) := \arg\min_{w' \in \mathcal{F}} \|w - w'\|_2$. Let $D_w = \sup_{w \in \mathcal{F}} \|w\|_2 < \infty$. In our context, we include additional optimizer variables A_t and b_t. In this section, we first show that under our arrangement of the algorithm, the trajectory of gradients are uniformly bounded. Then, we show that A_t and b_t are also bounded. All proofs are in the appendix.

Algorithm 2 Pseudo-linear TO (Full-TO)

Require: Learning rates α_t, β_t, γ_t.
1: Initialize w_1, A_0 and b_0.
2: **for** $t = 1, \ldots, T$ **do**
3: Sample $\mathcal{B}_t \subset \{1, \ldots, N\}$ such that $|\mathcal{B}_t| = b$
4: $g_t \leftarrow \nabla F^{\mathcal{B}_t}(w_t)$
5: $A_t \leftarrow A_{t-1} + \alpha_t(g_t - A_{t-1}w_t - b_{t-1})w_t^\top$
6: $b_t \leftarrow b_{t-1} + \beta_t(g_t - A_{t-1}w_t - b_{t-1})$
7: $\widehat{G}_t \leftarrow A_t w_t + b_t$
8: $w_{t+1} \leftarrow w_t - \gamma_t \widehat{G}_t$
9: **end for**

For the theoretical analysis, we make the following assumptions. We assume that $\mathcal{F}$ is closed convex. This implies that w_{t+1} is a unique minimizer. Let $\bar{\mathcal{F}}$ be any open set that contains $\mathcal{F}$. Since $\bar{\mathcal{F}}$ is open, the derivatives on $\mathcal{F}$ can be appropriately defined.

Assumption 3.1. The full loss F and minibatch losses $F^{\mathcal{B}}$ for all $\mathcal{B}$ satisfy the following conditions.

1. Feasibility set $\mathcal{F}$ must be such that $w^* \in \mathcal{F}$, where w^* is an optimal solution to the unconstrained problem.
2. For all $\mathcal{B} \subset \{1, \ldots, N\}$ with $|\mathcal{B}| = b$, $\nabla F^{\mathcal{B}}(w)$ is Lipschitz continuous on $\bar{\mathcal{F}}$, i.e., there exists $L > 0$ such that for all $w, w' \in \bar{\mathcal{F}}$, we have $\left\|\nabla F^{\mathcal{B}}(w) - \nabla F^{\mathcal{B}}(w')\right\| \leq L \left\|w - w'\right\|_2$.
3. There exists $0 < D_V < +\infty$ such that for all $w \in \bar{\mathcal{F}}$, $\mathbb{E}_{\mathcal{B}}\left[\left\|\nabla F^{\mathcal{B}}(w) - \nabla F(w)\right\|_2^2\right] \leq D_V$, where $\mathcal{B}$ is a random subset of $\{1, \ldots, N\}$ with $|\mathcal{B}| = b$.

The conditions are commonly made when analyzing convergence of gradient based optimization methods. Remarkably, we avoid making the uniform bounded gradient assumption, which contradicts strong convexity. However, given the Assumption and the bound for the trajectory w_t, i.e., $\|w_t\|_2 \leq D_w$ for all t, the stochastic gradients g_t in each iterate are also uniformly bounded for each t, as shown in the following proposition.

Proposition 3.2. *Given $\|w_t\| \leq D_w$, there exists $0 < D_G < +\infty$ such that for all t, $\|g_t\|_2 \leq D_G$.*

Furthermore, under specific requirements on learning rates, the optimizer variables A_t and b_t are also bounded. In what follows, all matrix norms are spectral.

Proposition 3.3. *Let $S_\alpha := \sum_{s=1}^{\infty} \alpha_s < 1/D_w^2 < \infty$ and let $\beta_t \leq 1$ for all t. There exist D_A and D_b independent on t, such that $\|A_t\|_2 \leq D_A$ and $\|b_t\|_2 \leq D_b$ for all t.*

The requirement $\sum_{t=1}^{\infty} \alpha_s < \infty$ can be simply satisfied by $\alpha_s = 6/\pi^2 D_w^2 s^2$. We prove this proposition by properly applying triangle inequalities and the principle of induction.

Since $F(w) = \frac{1}{\binom{N}{b}} \sum_{\mathcal{B}:|\mathcal{B}|=b} F^{\mathcal{B}}(w)$, Item 3 in Assumption naturally implies that

$$\|\nabla F(w) - \nabla F(w')\|_2 \leq \frac{1}{\binom{N}{b}} \sum_{\mathcal{B}:|\mathcal{B}|=b} \|\nabla F^{\mathcal{B}}(w) - \nabla F^{\mathcal{B}}(w')\|_2 \leq L\|w - w'\|_2,$$

i.e., ∇F is also Lipschitz continuous. Therefore we have $\|\nabla F(w_t)\|_2 \leq L\|w_t - w^*\| \leq 2LD_w < D_G$ and we have $\max\left\{\|\nabla F(w_t)\|_2, \left\|\widehat{G}_t\right\|_2\right\} \leq D_G$. Using the previously defined constants L, D_A and D_G, we present the convergence result of Pseudo-linear TO for strongly convex losses.

Theorem 3.4. *Given Assumption and that F is a c strongly convex function, setting $\gamma_t = \gamma/(t + \mu)$ and $\beta_t = \beta/(t - 1 + \mu)$, where γ and β satisfy*

$$\gamma > \frac{1}{c} \quad \text{and} \quad \beta > 1 + \frac{4\gamma^2 L^2 \sqrt{D_A^2 + L^2}}{c(\gamma c - 1)} + 2\gamma\sqrt{D_A^2 + L^2},$$

and $\alpha_t = \alpha/(t-1+\mu)^2$ *with* $\sum_{s=1}^{\infty} \alpha_s < 1/D_w^2$*, and* μ *to be large enough such that* $\alpha_1 < 1, \beta_1 < 1$*,* $\gamma_1 < 1$*, and* $\alpha D_w^2/\mu + \beta < \mu$*, there exist* $M_1, M_2 > 0$ *such that for all* t*, we have*

$$\mathbb{E}\left[\|w_t - w^*\|_2^2\right] \leq \frac{M_1}{t+\mu}, \quad \mathbb{E}\left[\left\|\widehat{G}_t - \nabla F(w_t)\right\|_2^2\right] \leq \frac{M_2}{t+\mu}.$$

To prove the theorem, we jointly bound $\mathbb{E}\left[\|w_t - w^*\|_2^2\right]$ and $\mathbb{E}\left[\left\|\widehat{G}_t - \nabla F(w_t)\right\|_2^2\right]$ recurrently and apply induction. The expectations are over history up to t. Theorem 3.4 shows that for strongly convex losses, Pseudo-linear TO achieves an $\mathcal{O}(1/t)$ convergence rate, matching the rates of SGD and Momentum. Comparing to ADAM, which has strongly convex examples with non-convergent variance (Wang & Klabjan, 2025), namely $\mathbb{E}\left[\|\widehat{G}_t^{(\mathrm{ADAM})} - \nabla F(w_t)\|^2\right] \not\to 0$, the convergence result for Pseudo-linear TO automatically achieves the reduction of variance.

Relaxing the requirement of strong convexity, we also present the convergence of Pseudo-linear TO for convex but not strongly convex losses. Alternatively, we assume that F has a nonempty and bounded set of global minimizers. We specifically let the feasible set $\mathcal{F}$ equal to $\mathcal{F}_c := \{w \in \mathbb{R}^d : F(w) \leq c\}$, where c is a constant greater than $\min_{w\in\mathbb{R}^d} F(w)$ so that $\mathcal{F}_c$ is nonempty . We show that each $w \in \mathcal{F}_c$ has uniform upper bound of $\|w\|_2$ and $\|\nabla F(w)\|_2$. Therefore the propositions bounding $\|A\|_2$ and $\|b\|_2$ also hold. The proof is in the appendix.

Theorem 3.5. *Given Assumption and that* F *is convex with nonempty and bounded set of minimizers, setting* $\gamma_t = \gamma/(t+\mu)$ *and* $\beta_t = \beta/(t-1+\mu)$*, where* γ *and* β *satisfy* $\gamma < \frac{1}{L}$ *and* $\beta > 1 + 6\gamma\sqrt{D_A^2 + L^2}$*, and* $\alpha_t = \alpha/(t-1+\mu)^2$ *with* $\sum_{s=1}^{\infty} \alpha_s < 1/D_w^2$*, and* μ *to be large enough such that* $\alpha_1 < 1, \beta_1 < 1$*,* $\gamma_1 < 1$*, and* $\alpha D_w^2/\mu + \beta < \mu$*, we have*

$$\min_{1\leq t\leq T} \mathbb{E}\left[\|\widehat{G}_t - G_t\|_2^2\right] \leq \mathcal{O}\left(\frac{1}{\log T}\right), \quad \min_{1\leq t\leq T} \mathbb{E}\left[\|G_t\|_2^2\right] \leq \mathcal{O}\left(\frac{1}{\log T}\right).$$

3.3 Simplified Variants of Pseudo-linear TO

Algorithm 3 Diagonal TO (Diag-TO)

Require: Learning rates $\alpha_t, \beta_t, \gamma_t$.
Initialize w_1, a_0 and b_0.
for $t = 1, \ldots, T$ **do**
  Sample $\mathcal{B}_t \subset \{1, \ldots, N\}$ such that $|\mathcal{B}_t| = b$
  $g_t \leftarrow \nabla F^{\mathcal{B}_t}(w_t)$
  $a_t \leftarrow a_{t-1} + \alpha_t(g_t - a_{t-1} \odot w_t - b_{t-1}) \odot w_t$
  $b_t \leftarrow b_{t-1} + \beta_t(g_t - a_{t-1} \odot w_t - b_{t-1})$
  $\widehat{G}_t \leftarrow a_t \odot w_t + b_t$
  $w_{t+1} \leftarrow w_t - \gamma_t \widehat{G}_t$
end for

Algorithm 4 RankOne TO (RO-TO)

Require: Learning rates $\alpha_t, \beta_t, \gamma_t$.
Initialize w_1, a_0, c_0 and b_0.
for $t = 1, \ldots, T$ **do**
  Sample $\mathcal{B}_t \subset \{1, \ldots, N\}$ such that $|\mathcal{B}_t| = b$
  $g_t \leftarrow \nabla F^{\mathcal{B}_t}(w_t)$
  $a_t \leftarrow a_{t-1} + \alpha_t(g_t - a_{t-1}c_{t-1}^\top w_t - b_{t-1})c_{t-1}^\top w_t$
  $c_t \leftarrow c_{t-1} + \alpha_t(g_t - a_{t-1}c_{t-1}^\top w_t - b_{t-1})^\top a_{t-1} w_t$
  $b_t \leftarrow b_{t-1} + \beta_t(g_t - a_{t-1}c_{t-1}^\top w_t - b_{t-1})$
  $\widehat{G}_t \leftarrow a_t c_t^\top w_t + b_t$
  $w_{t+1} \leftarrow w_t - \gamma_t \widehat{G}_t$
end for

A notable limitation of Pseudo-linear TO lies in its requirement of memory. For model weights $w \in \mathbb{R}^d$, the optimizer variables A and b introduce $d^2 + d$ new parameters, which is impractical for large models, when d is large. To address this challenge, we propose two memory-efficient variants of the fully parameterized Pseudo-linear TO (Full-TO) introduced in Algorithm 2.

- **Diagonal TO** (Algorithm 3) enforces $A_t = \mathrm{diag}\{a_t\}$, where $a_t \in \mathbb{R}^d$, reducing the parameter count to $\mathcal{O}(d)$. Here, $\odot$ denotes element-wise vector multiplication.
- **RankOne TO** (Algorithm 4) employs factorization $A_t = a_t c_t^\top$ with $a_t, c_t \in \mathbb{R}^d$, similarly achieving memory complexity $\mathcal{O}(d)$ while maintaining a richer parameter interaction.

The gradient update rules in both algorithms (Lines 5-6 of Algorithm 3 and Lines 5-7 of Algorithm 4) are derived from the corresponding structured gradient computations for each parameterization. These modifications enlarge the scale of problems suitable under the Pseudo-linear TO framework. It is easy to extend the proof of Theorem 3.4 and show the convergence of Diagonal and RankOne TO.

4 Experiments

4.1 Datasets and Implementations

We conduct comprehensive empirical evaluations comparing various TO implementations against ADAM, Momentum and SAGA across multiple real-world datasets.

- **MNIST** (Deng, 2012): A handwritten digit dataset comprising of 60,000 grayscale images (28×28 pixels)
- **News20** (Lang, 1995): Approximately 18,000 newsgroup posts across 20 topics
- **CovType** (Blackard et al., 1998): Forest cover type prediction dataset with 581,012 samples and 54 features across 7 classes
- **SmallNorb** (LeCun et al., 2004): 3D object recognition dataset containing images of 50 toys from 5 categories
- **KDD** (Tavallaee et al., 2009): Network intrusion detection subset from KDD CUP 99
- **CIFAR-10** (Krizhevsky & Hinton, 2009): 60,000 color images across 10 object classes
- **Alpaca** (Taori et al., 2023): 52,000 instruction-following demonstrations generated by text-davinci-003

The models employed are logistic regressions or deep neural networks. The architectures are described later in the appendix. The last dataset corresponds to fine-tuning of an LLM. The experiments span strongly convex, convex, non-convex and LLM tasks. Given the negligible computational overhead beyond gradient computation, we focus on convergence speed measured in epochs. For each experiment, we perform extensive hyperparameter tuning:

- ADAM learning rates: $\gamma \in [10^{-5},\ 2\times10^{-5},\ 5\times10^{-5},\ 10^{-4},\ 2\times10^{-4},\ 5\times10^{-4},\ 10^{-3},\ 2\times10^{-3},\ 5\times10^{-3}]$,
- TO methods: $\gamma \in [10^{-3},\ 2\times10^{-3},\ 5\times10^{-3},\ 0.01,\ 0.02,\ 0.05,\ 0.1,\ 0.2,\ 0.5]$,
- TO-specific parameters: $\alpha, \beta \in$ [0.0, 0.01, 0.1, 0.5, 1.0].

All the hyperparameters are tested under constant and exponentially decaying scheme. The exponential decay rates are selected among $[0.6,\ 0.8,\ 0.95]$. The trainable weights of logistic regression and FFN are initialized from standard normal sampling. The mini-batch size is set to be 64.

We use the full training loss as our primary evaluation metric. To ensure statistical reliability, we repeat each experiment five times with different random seeds. Using ADAM as our baseline, we calculate the relative performance difference ρ for each method. Let loss_k be the training loss of a considered algorithm in epoch k. The relative difference is defined as

$$\rho = \mathrm{avg}\left(\frac{\min_{1\le k\le T}\{\text{ADAM loss}_k\} - \min_{1\le k\le T}\{\mathrm{loss}_k\}}{\min_{1\le k\le T}\{\text{ADAM loss}_k\}}\right),$$

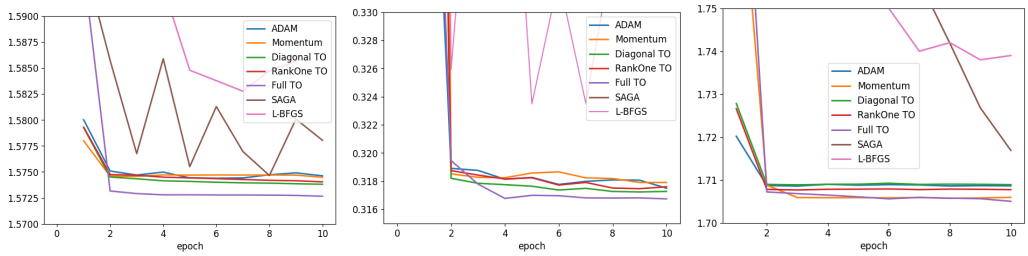

(a) CovType (b) KDD (c) MNIST

Figure 1: Training loss curves for logistic regression ($\lambda = 0$) on different datasets

Table 1: Memory requirements of Full-TO and its memory-efficient variants

	Full-TO	Diag-TO	RO-TO
Logistic-CovType	150 KB	1.5 KB	2.3 KB
Logistic-MNIST	400.8 MB	80.1 KB	120.1 KB
FFN-CovType	381.6 MB	78.1 KB	117.2 KB
FFN-MNIST	39.9 GB	808.7 KB	1.2 MB

where T is the number of epochs and the average is taken across multiple runs. Metric $\rho > 0$ indicates superior performance to ADAM in terms of achieved minimum loss. We complement this with two-sided Wald tests on the relative differences, reporting significance levels s. Our statistical analysis follows these conventions: $\mathbf{s < 0.05}$: Significantly better than ADAM, $\mathbf{s > 0.95}$: Significantly worse than ADAM, $\mathbf{0.05 \leq s \leq 0.95}$: No significant difference from ADAM.

Our theoretical analyses show that pseudo-linear and other variants of TO achieve an $\mathcal{O}(1/t)$ convergence rate for strongly convex losses, and a convergence of the approximate variance with the same convergence rate. This suggests the expectation of TO's advantage on strongly convex task and noisy and complex tasks.

4.2 Convex Experiments

For convex losses, we train logistic regression models on five of the previously mentioned datasets. Since our theoretical analyses primarily address strongly convex losses, we emphasize results using logistic regression with cross-entropy loss and L2 regularization. Our experimental protocol involves three key validation steps to determine the optimal regularization parameter λ. First, we randomly partition each dataset into training and validation sets. Second, we train models using ADAM with λ values from the discrete set $[0,\ 5 \times 10^{-4},\ 10^{-3},\ 5 \times 10^{-3},\ 0.01,\ 0.05]$, selecting the λ yielding best validation performance with ADAM. When the optimal λ equals zero, we conduct additional evaluations by randomly sampling five normally distributed λ values with mean being 0 and standard deviation of 1, and filter the negative values, in order to thoroughly compare TO variants against ADAM on strongly convex losses. The final selected λ values for each dataset appear in the appendix.

Due to memory constraints, we exclude the full version of pseudo-linear TO results for SmallNorb and News20 datasets. For experiments in which Full-TO is feasible, we report the memory requirements of Full-TO and its two memory-efficient variants in Table 1. In the appendix, we present the detailed values of ρ and significance for all strongly convex experiments. The results show that Diagonal TO and RankOne TO substantially reduce the memory requirements compared to Full-TO, making them more suitable for high-dimensional optimization problems. Our analysis of 17 strongly convex experiments reveals that TO variants significantly outperform ADAM in 16 cases. TO variants also outperform Momentum in 15 cases. The complete values of ρ including Momentum are in the appendix.

To investigate performance on convex but non-strongly-convex losses, we conduct parallel experiments without regularization $\lambda = 0$ on the same datasets. In Figure 1, we plot the curve of training loss of the $\lambda = 0$ experiments for the three datasets where Full TO could be evaluated (CovType, MNIST and KKD) and find that pseudo-linear TO exceeds the performances of baselines such as

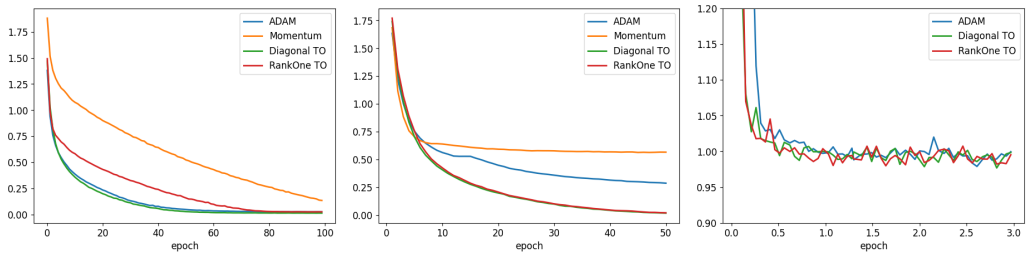

(a) ResNet18-CIFAR10 (b) ResNet50-CIFAR10 (c) Llama7b-Alpaca

Figure 2: Training loss curves for non-convex tasks

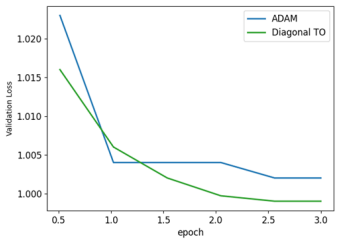

Figure 3: Validation loss curve for Llama7b-alpaca task

ADAM, Momentum, SAGA and L-BFGS. Additional results, such as training losses for different additional datasets, and the standard deviation of the train loss among different random seeds can be found in the appendix. Our empirical findings suggest initializing with $\alpha \approx 0.01$ while $\beta \approx 1$. Exponential decaying of α and β also plays crucial rules practically. These settings provide robust performance across different datasets.

4.3 Nonconvex Experiments

For non-convex optimization tasks, we conduct three experimental evaluations. First, we train FFNs on both the CovType and KDD datasets. Second, we train ResNets of varying depths (18 and 50 layers) on the CIFAR-10 classification task. Third, we perform fine-tuning of a selected subset of pretrained weights from the Llama-7b model using the Alpaca dataset.

Table 2: Values of ρ for non-convex tasks

Model	Dataset	Full-TO	Diag-TO	RO-TO
FFN	CovType	-1.3×10^{-4}	-1.8×10^{-4}	-2.0×10^{-4}
	KDD	-1.1×10^{-3}	-3.2×10^{-3}	-2.1×10^{-3}
ResNet18	CIFAR10	N/A	$\mathbf{0.5}$	8.2×10^{-3}
ResNet50		N/A	$\mathbf{0.9}$	$\mathbf{0.9}$
Llama-7b	Alpaca	N/A	$\mathbf{2.1 \times 10^{-3}}$	-9.2×10^{-4}

The relative performance differences of TO variants are reported in Table 2. Figure 2 shows the curves of training loss for the tasks of training CIFAR10 with ResNet and finetuning Llama. These experiments demonstrate TO's superior performance on challenging non-convex problems. For the Llama fine-tuning task, TO achieves faster initial convergence compared to ADAM. In the ResNet classification tasks, TO consistently discovers solutions with better final optimality. To investigate the effect of TO on the generalization performance of LLMs, we evaluate the validation loss in the Llama-7b Alpaca experiment, with the results shown in Figure 3. We observe that Diagonal TO consistently outperforms ADAM in terms of validation-loss convergence, particularly during the early stages of training. Moreover, Diagonal TO achieves a lower final validation loss than ADAM. Given its rapid loss reduction in the initial iterations, these results suggest that TO is especially well suited for LLM training scenarios where computational budgets or training time are limited.

AI USE STATEMENT

In this work, we used generative AI tools for polishing writing. We have not used generative AI tools for retrieval or discovery, research ideation or execution, drafting sections of the paper, generating synthetic datasets or proving mathematical claims. We have reviewed all AI-assisted work. All LLM polished paragraphs were verified by two authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

Marcin Andrychowicz, Misha Denil, Sergio Gómez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL https://proceedings.neurips.cc/paper_files/paper/2016/file/fb87582825f9d28a8d42c5e5e5e8b23d-Paper.pdf.

Irwan Bello, Barret Zoph, Vijay Vasudevan, and Quoc V. Le. Neural optimizer search with reinforcement learning. In *International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 459–468, International Convention Centre, Sydney, Australia, 2017. PMLR. URL http://proceedings.mlr.press/v70/bello17a.html.

Jock A. Blackard, Denis J. Dean, and Charles W. Anderson. UCI machine learning repository: Covertype data set. https://archive.ics.uci.edu/ml/datasets/covertype, 1998.

Tianlong Chen, Xiaohan Chen, Wuyang Chen, Howard Heaton, Jialin Liu, Zhangyang Wang, and Wotao Yin. Learning to optimize: A primer and a benchmark. *Journal of Machine Learning Research*, 23(189):1–59, 2022. URL http://jmlr.org/papers/v23/21-0308.html.

A. Cutkosky and F. Orabona. Momentum-based variance reduction in non-convex sgd. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 15210–15219, 2019.

Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL https://proceedings.neurips.cc/paper_files/paper/2014/file/ede7e2b6d13a41ddf9f4bdef84fdc737-Paper.pdf.

Li Deng. The MNIST database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.

John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7), 2011.

Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky. Neural networks for machine learning: lecture 6a overview of mini-batch gradient descent. 2012. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf.

Feihu Huang, Junyi Li, and Heng Huang. Super-adam: Faster and universal framework of adaptive gradients, 2022. URL https://arxiv.org/abs/2106.08208.

Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013a. URL https://proceedings.neurips.cc/paper_files/paper/2013/file/ac1dd209cbcc5e5d1c6e28598e8cbbe8-Paper.pdf.

Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in Neural Information Processing Systems*, 2013b.

Diederik P Kingma and Jimmy Ba. ADAM: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.

Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009. https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.222.9220&rep=rep1&type=pdf.

Ken Lang. Newsweeder: Learning to filter netnews. In *Proceedings of the Twelfth International Conference on Machine Learning*, pp. 331–339, 1995.

Yann LeCun, Fu Jie Huang, and Léon Bottou. Learning methods for generic object recognition with invariance to pose and lighting. *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2:II–104 Vol.2, 2004.

Ke Li and Jitendra Malik. Learning to optimize. In *International Conference on Learning Representations (ICLR)*, 2017.

Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical Programming*, 45(1-3):503–528, 1989.

Xin Qian and Diego Klabjan. The impact of the mini-batch size on the variance of gradients in stochastic gradient descent. *arXiv preprint arXiv:2004.13146*, 2020.

Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of ADAM and beyond. *International Conference on Learning Representations*, 2018.

Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.

Mahbod Tavallaee, Ebrahim Bagheri, Wei Lu, and Ali A Ghorbani. A detailed analysis of the KDD cup 99 data set. In *2009 IEEE symposium on Computational Intelligence for Security and Defense Applications*, 2009.

Ruiqi Wang and Diego Klabjan. Divergence results and convergence of a variance reduced version of adam. *arXiv:2210.05607*, 2025. URL https://arxiv.org/abs/2210.05607.

Manzil Zaheer, Sashank Reddi, Devendra Sachan, Satyen Kale, and Sanjiv Kumar. Adaptive methods for nonconvex optimization. *Advances in Neural Information Processing Systems*, 2018.

A Proofs

A.1 Technical Lemmas

Lemma A.1. *For all t, $\|w_{t+1} - w_t\|_2 \leq \gamma_t \left\|\widehat{G}_t\right\|_2$.*

Proof. (Reddi et al., 2018) shows that for all u and u', $\|\Pi_{\mathcal{F}}(u) - \Pi_{\mathcal{F}}(u')\|_2 \leq \|u - u'\|_2$. In our case, we have $w_t \in \mathcal{F}$, therefore $\|w_{t+1} - w_t\|_2 = \|\Pi_{\mathcal{F}}(w_t - \gamma_t \widehat{G}_t) - \Pi_{\mathcal{F}}(w_t)\|_2 \leq \|\gamma_t \widehat{G}_t\|_2$. □

Lemma A.2. *Let $0 < k_2 < k_1$ and $b_1 < 0 < b_2$, for any x_0 and y_0. There exists $x \geq x_0$ and $y \geq y_0$ such that $k_2 x + b_2 \leq y \leq k_1 x + b_1$.*

Proof. Let $x^* = \frac{b_2 - b_1}{k_1 - k_2}$, $y^* = \frac{k_1 b_2 - k_2 b_1}{k_1 - k_2}$, and let $x_n = x^* + n$, $y_n = y^* + \frac{k_1 + k_2}{2} n$. It is easy to check that $k_2 x_n + b_2 \leq y_n \leq k_1 x_n + b_1$ for all $n > 0$. We can pick n large enough such that $x_n \geq x_0$ and $y_n \geq y_0$. □

Lemma A.3. *Let $F : \mathbb{R}^d \to \mathbb{R}$ be differentiable and convex, with a globally L-Lipschitz continuous gradient for some $L > 0$. Suppose that the set of global minimizers*

$$S := \arg\min_{w \in \mathbb{R}^d} F(w)$$

is nonempty and bounded, and let $F_* := \min_{w \in \mathbb{R}^d} F(w) > -\infty$. For any finite $c \geq F_*$, let us define the level set

$$\mathcal{F}_c := \{w \in \mathbb{R}^d : F(w) \leq c\}.$$

Then $\mathcal{F}_c$ is nonempty, compact, and convex, and

$$\sup_{w\in\mathcal{F}_c} \|\nabla F(w)\|_2 \le \sqrt{2L(c-F_*)}.$$

Moreover, the Euclidean projection $\Pi_{\mathcal{F}_c}(y)$ exists and is unique for every $y \in \mathbb{R}^d$. For every $x \in \mathcal{F}_c$ and $y \in \mathbb{R}^d$, it satisfies

$$\|\Pi_{\mathcal{F}_c}(y) - x\|_2 \le \|y - x\|_2, \qquad F(\Pi_{\mathcal{F}_c}(y)) \le F(y).$$

Proof. Since F is continuous and convex, $\mathcal{F}_c$ is closed and convex. It is nonempty because $S \subseteq \mathcal{F}_c$.

We first prove that $\mathcal{F}_c$ is bounded. Suppose otherwise. Fix $x_* \in S$ and choose a sequence $\{z_k\} \subseteq \mathcal{F}_c$ such that

$$r_k := \|z_k - x_*\|_2 \longrightarrow \infty.$$

After discarding finitely many terms, define $u_k := (z_k - x_*)/r_k$. By compactness of the unit sphere, we may consider a subsequence such that $u_k \to u$ with $\|u\|_2 = 1$. For any fixed $t > 0$ and all sufficiently large k, convexity gives

$$\begin{aligned} F_* &\le F(x_* + tu_k) \\ &= F\left(\left(1 - \frac{t}{r_k}\right) x_* + \frac{t}{r_k} z_k\right) \\ &\le \left(1 - \frac{t}{r_k}\right) F(x_*) + \frac{t}{r_k} F(z_k) \\ &\le F_* + \frac{t}{r_k}(c - F_*). \end{aligned}$$

Letting $k \to \infty$ and using continuity of F, we obtain

$$F(x_* + tu) = F_*, \qquad t \ge 0.$$

Thus the entire ray $\{x_* + tu : t \ge 0\}$ lies in S, contradicting the boundedness of S. Hence $\mathcal{F}_c$ is bounded and therefore compact.

Next, the descent lemma yields, for every $w \in \mathbb{R}^d$,

$$F_* \le F\left(w - \frac{1}{L}\nabla F(w)\right) \le F(w) - \frac{1}{2L}\|\nabla F(w)\|_2^2.$$

Consequently, for every $w \in \mathcal{F}_c$,

$$\|\nabla F(w)\|_2^2 \le 2L\big(F(w) - F_*\big) \le 2L(c - F_*),$$

which proves the uniform gradient bound.

Since $\mathcal{F}_c$ is nonempty, closed, and convex, its Euclidean projection exists and is unique. We fix $w \in \mathcal{F}_c$ and $y \in \mathbb{R}^d$, and set $p = \Pi_{\mathcal{F}_c}(y)$. The variational characterization of the projection gives

$$(y - p)^\top (w - p) \le 0.$$

Expanding $\|y - x\|_2^2$ therefore yields

$$\|p - x\|_2^2 \le \|y - x\|_2^2 - \|y - p\|_2^2 \le \|y - x\|_2^2,$$

which proves the distance inequality.

Finally, if $y \in \mathcal{F}_c$, then $p = y$, so $F(p) = F(y)$. If $y \notin \mathcal{F}_c$, then $F(y) > c$, whereas $p \in \mathcal{F}_c$ implies

$$F(p) \le c < F(y).$$

Thus $F(F_{\mathcal{F}_c}(y)) \le F(y)$ in both cases. □

A.2 Proof to Proposition 3.2

Proof. By definition g_t reads

$$\begin{aligned}\|g_t\|_2 &= \|\nabla F^{\mathcal{B}_t}(w_t)\|_2 \\ &\leq \|\nabla F^{\mathcal{B}_t}(w_t) - \nabla F^{\mathcal{B}_t}(w^*)\|_2 + \|\nabla F^{\mathcal{B}_t}(w^*)\|_2 \\ &\leq L\|w_t - w^*\| + \max_{\mathcal{B}} \|\nabla F^{\mathcal{B}}(w^*)\|_2 \\ &\leq 2LD_w + \max_{\mathcal{B}} \|\nabla F^{\mathcal{B}}(w^*)\|_2,\end{aligned}$$

where the inequalities hold according to Assumption. The proposition is proved by letting $D_G = 2LD_w + \max_{\mathcal{B}} \|\nabla F^{\mathcal{B}}(w^*)\|_2$. □

A.3 Proof to Proposition 3.3

Proof. Consider

$$\begin{aligned}D_b &= \max\left\{\|b_0\|_2, \frac{1}{1 - S_\alpha D_w^2}\left(D_G(1 + S_\alpha D_w^2) + \|A_0\|_2 D_w\right)\right\} \\ D_A &= \|A_0\|_2 + S_\alpha D_w(D_b + D_G).\end{aligned}$$

Apparently $\|b_0\|_2 \leq D_b$. We assume that for all $0 \leq t \leq T-1$, $\|b_t\| \leq D_b$. Reformulate the updating rule for A_t yields

$$\begin{aligned}A_t &= A_{t-1} + \alpha_t(g_t - A_{t-1}w_t - b_{t-1})w_t^\top \\ &= A_{t-1}(I_d - \alpha_t w_t w_t^\top) + \alpha_t(g_t - b_t)w_t^\top.\end{aligned}$$

Then, for all $1 \leq t \leq T-1$,

$$\begin{aligned}\|A_t\|_2 &= \left\|A_{t-1}(I_d - \alpha_t w_t w_t^\top) + \alpha_t(g_t - b_t)w_t^\top\right\|_2 \\ &\leq \left\|I_d - \alpha_t w_t w_t^\top\right\|_2 \|A_{t-1}\|_2 + \alpha_t \|g_t - b_t\|_2 \|w_t\|_2 \\ &\leq \|A_{t-1}\|_2 + \alpha_t(D_b + D_G)D_w,\end{aligned}$$

where the last inequality holds because $\|g_t\| \leq D_G$. We recursively obtain

$$\begin{aligned}\|A_t\|_2 &\leq \|A_0\|_2 + \sum_{s=1}^{t} \alpha_s(D_b + D_G)D_w \\ &< \|A_0\|_2 + S_\alpha(D_b + D_G)D_w = D_A.\end{aligned}$$

Furthermore, we have

$$\begin{aligned}\|b_T\|_2 &= \|(1 - \beta_T)b_{T-1} + \beta_T(g_T - A_{T-1}w_T)\|_2 \\ &\leq (1 - \beta_T)\|b_{T-1}\|_2 + \beta_T(\|g_T\|_2 + \|A_{T-1}\|_2\|w_T\|_2) \\ &\leq (1 - \beta_T)D_b + \beta_T(D_G + D_A D_w) \\ &\leq D_b.\end{aligned}$$

The last inequality holds because

$$\begin{aligned}D_b &= \max\left\{\|b_0\|_2, \frac{1}{1 - S_\alpha D_w^2}\left(D_G(1 + S_\alpha D_w^2) + \|A_0\|_2 D_w\right)\right\} \\ &\geq \frac{1}{1 - S_\alpha D_w^2}\left(D_G(1 + S_\alpha D_w^2) + \|A_0\|_2 D_w\right),\end{aligned}$$

which implies that

$$(1 - S_\alpha D_w^2)D_b \geq D_G(1 + S_\alpha D_w^2) + \|A_0\|_2 D_w,$$

and therefore

$$\begin{aligned}D_b &\geq D_b S_\alpha D_w^2 + D_G(1 + S_\alpha D_w^2) + \|A_0\|_2 D_w \\ &= D_G + D_w(\|A_0\|_2 + S_\alpha D_w(D_G + D_b)) = D_G + D_w D_A.\end{aligned}$$

According to the principle of induction, this finishes the proof. □

A.4 PROOF TO THEOREM 3.4

Proof. We start with bounding the approximation loss of the full gradient. Letting $G_t = \nabla F(w_t)$, the approximation error reads

$$\begin{aligned}
\widehat{G}_t - G_t &= A_t w_t + b_t - G_t \\
&= \left(A_{t-1} + \alpha_t \left(g_t - A_{t-1}w_t - b_{t-1}\right) w_t^\top\right) w_t + b_{t-1} + \beta_t \left(g_t - A_{t-1}w_t - b_{t-1}\right) - G_t \\
&= A_{t-1}w_t + b_{t-1} + \left(\alpha_t \|w_t\|_2^2 + \beta_t\right)\left(g_t - A_{t-1}w_t - b_{t-1}\right) - G_t \\
&= (1-\delta_t)\left(A_{t-1}w_t + b_{t-1} - G_t\right) + \delta_t \left(g_t - G_t\right),
\end{aligned}$$

where $\delta_t = \alpha_t \|w_t\|_2^2 + \beta_t$. Letting $\delta = \alpha D_w^2/\mu + \beta$, apparently we have $\delta_t = (\alpha\|w_t\|_2^2/(t-1+\mu) + \beta)/(t-1+\mu) < \delta/(t-1+\mu)$. Taking the square norm of the approximation error yields

$$\begin{aligned}
\left\|\widehat{G}_t - G_t\right\|_2^2 &= (1-\delta_t)^2 \left\|A_{t-1}w_t + b_{t-1} - G_t\right\|_2^2 + \delta_t^2 \left\|g_t - G_t\right\|_2^2 \\
&\quad +2(1-\delta_t)\delta_t \left(A_{t-1}w_t + b_{t-1} - G_t\right)^\top \left(g_t - G_t\right).
\end{aligned}$$

Letting $\mathbb{E}_t[\cdot] = \mathbb{E}\left[\cdot \,|\mathcal{B}_1, \cdots, \mathcal{B}_{t-1}\right]$, we notice that A_{t-1}, b_{t-1}, G_t and w_t are known given $\mathcal{B}_1, \cdots, \mathcal{B}_{t-1}$. Then, we have

$$\begin{aligned}
\mathbb{E}_t\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] &= (1-\delta_t)^2 \left\|A_{t-1}w_t + b_{t-1} - G_t\right\|_2^2 + \delta_t^2 \mathbb{E}_t\left[\left\|g_t - G_t\right\|_2^2\right] \\
&\quad +2(1-\delta_t)\delta_t \left(A_{t-1}w_t + b_{t-1} - G_t\right)^\top \mathbb{E}_t\left[g_t - G_t\right] \\
&= (1-\delta_t)^2 \left\|A_{t-1}w_t + b_{t-1} - G_t\right\|_2^2 + \delta_t^2 \mathbb{E}_t\left[\left\|g_t - G_t\right\|_2^2\right] \\
&\leq (1-\delta_t)^2 \left\|A_{t-1}w_t + b_{t-1} - G_t\right\|_2^2 + D_V \frac{\delta^2}{(t-1+\mu)^2} \\
&= (1-\delta_t)^2 \underbrace{\left\|\widehat{G}_{t-1} - G_{t-1} + A_{t-1}(w_t - w_{t-1}) - (G_t - G_{t-1})\right\|_2^2}_{T_1} + D_V \frac{\delta^2}{(t-1+\mu)^2},
\end{aligned}$$

where the second equality holds because $\mathbb{E}_t g_t = G_t$.

For three arbitrary vectors a_1, a_2, a_3, the Cauchy inequality implies that for any $\lambda > 0$,

$$\begin{aligned}
\|a_1 + a_2 + a_3\|_2^2 &= \|a_1\|_2^2 + \|a_2\|_2^2 + \|a_3\|_2^2 + 2a_1^\top a_2 + 2a_2^\top a_3 + 2a_3^\top a_1 \\
&\leq \|a_1\|_2^2 + \|a_2\|_2^2 + \|a_3\|_2^2 + \frac{1}{2\lambda}\|a_1\|_2^2 + 2\lambda\|a_2\|_2^2 + \frac{1}{2\lambda}\|a_1\|_2^2 + 2\lambda\|a_3\|_2^2 + \|a_2\|_2^2 + \|a_3\|_2^2 \\
&= (1 + \frac{1}{\lambda})\|a_1\|_2^2 + (2 + 2\lambda)\|a_2\|_2^2 + (2 + 2\lambda)\|a_3\|_2^2.
\end{aligned}$$

Applying this inequality to T_1 and replacing λ with $\lambda/(t-1+\mu)$, we have

$$\begin{aligned}
T_1 &\leq \left(1 + \frac{\lambda}{t-1+\mu}\right)\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + 2\left(1 + \frac{t-1+\mu}{\lambda}\right)\left(\|A_{t-1}(w_t - w_{t-1})\|_2^2 + \|G_t - G_{t-1}\|_2^2\right) \\
&\leq \left(1 + \frac{\lambda}{t-1+\mu}\right)\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + 2\left(1 + \frac{t-1+\mu}{\lambda}\right)\left(D_A^2 + L^2\right)\|w_t - w_{t-1}\|_2^2 \\
&\leq \left(1 + \frac{\lambda}{t-1+\mu}\right)\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + 2\left(1 + \frac{t-1+\mu}{\lambda}\right)\left(D_A^2 + L^2\right)\gamma_{t-1}^2\|\widehat{G}_{t-1}\|_2^2 \\
&\leq \left(1 + \frac{\lambda}{t-1+\mu}\right)\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + 4\left(1 + \frac{t-1+\mu}{\lambda}\right)\left(D_A^2 + L^2\right)\gamma_{t-1}^2(\|\widehat{G}_{t-1} - G_{t-1}\|_2^2 + \|G_{t-1}\|_2^2) \\
&= \left(1 + \frac{\lambda}{t-1+\mu} + 4\left(1 + \frac{t-1+\mu}{\lambda}\right)\left(D_A^2 + L^2\right)\gamma_{t-1}^2\right)\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 \\
&\quad +4\left(1 + \frac{t-1+\mu}{\lambda}\right)\left(D_A^2 + L^2\right)\gamma_{t-1}^2\|G_{t-1}\|_2^2,
\end{aligned}$$

where the second inequality applies Proposition 3.3 and Lipschitz continuity of ∇F, the third inequality applies Lemma A.1 and the forth inequality further applies the Cauchy inequality. We let $\lambda = 2\gamma\sqrt{D_A^2 + L^2}$. The upper bound is simplified to

$$\begin{aligned} T_1 &\leq \left(1 + \frac{4\gamma\sqrt{D_A^2+L^2}}{t-1+\mu} + \frac{4\gamma^2(D_A^2+L^2)}{(t-1+\mu)^2}\right)\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + 4\left(1 + \frac{t-1+\mu}{2\gamma\sqrt{D_A^2+L^2}}\right)\frac{\gamma^2(D_A^2+L^2)}{(t-1+\mu)^2}\|G_{t-1}\|_2^2 \\ &= \left(1 + \frac{2\gamma\sqrt{D_A^2+L^2}}{t-1+\mu}\right)^2\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + 4\left(1 + \frac{t-1+\mu}{2\gamma\sqrt{D_A^2+L^2}}\right)\frac{\gamma^2(D_A^2+L^2)}{(t-1+\mu)^2}\|G_{t-1}\|_2^2 \qquad (3) \\ &\leq \left(1 + \frac{2\gamma\sqrt{D_A^2+L^2}}{t-1+\mu}\right)^2\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + 4\left(1 + \frac{t-1+\mu}{2\gamma\sqrt{D_A^2+L^2}}\right)\frac{\gamma^2L^2(D_A^2+L^2)}{(t-1+\mu)^2}\|w_{t-1} - w^*\|_2^2 \\ &\leq \left(1 + \frac{2\gamma\sqrt{D_A^2+L^2}}{t-1+\mu}\right)^2\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + \frac{4\gamma L^2\sqrt{D_A^2+L^2}}{t-1+\mu}\|w_{t-1} - w^*\|_2^2, \end{aligned}$$

where the second inequality further applies the Lipschitz continuity of ∇F. The last inequality holds because μ is large enough such that

$$\frac{1}{t-1+\mu} \leq \frac{1}{\beta} < \frac{1}{2\gamma\sqrt{D_A^2+L^2}},$$

and thus

$$\frac{1}{t-1+\mu}\left(1 + \frac{t-1+\mu}{2\gamma\sqrt{D_A^2+L^2}}\right) < \frac{1}{\gamma\sqrt{D_A^2+L^2}}.$$

The expectation $\mathbb{E}_t[\|\widehat{G}_t - G_t\|_2^2]$ is upper bounded by

$$\mathbb{E}_t\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] \leq (1-\delta_t)^2\left(1 + \frac{2\gamma\sqrt{D_A^2+L^2}}{t-1+\mu}\right)^2\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + \frac{C_1\gamma}{t-1+\mu}\|w_{t-1} - w^*\|_2^2 + \frac{D_V\delta^2}{(t-1+\mu)^2},$$

where $C_1 = 4L^2\sqrt{D_A^2+L^2}$. The setting of parameters satisfies that $\beta_t < \delta_t < 1$ and therefore

$$\begin{aligned} (1-\delta_t)^2\left(1 + \frac{2\gamma\sqrt{D_A^2+L^2}}{t-1+\mu}\right)^2 &< \left(1 - \frac{\beta}{t-1+\mu}\right)^2\left(1 + \frac{2\gamma\sqrt{D_A^2+L^2}}{t-1+\mu}\right)^2 \\ &< \left(1 - \frac{\beta - 2\gamma\sqrt{D_A^2+L^2}}{t-1+\mu}\right)^2 < 1 - \frac{\beta - 2\gamma\sqrt{D_A^2+L^2}}{t-1+\mu} := 1 - \frac{\beta^*}{t-1+\mu}, \end{aligned}$$

where $\beta^* := \beta - 2\gamma\sqrt{D_A^2+L^2} > 0$. We obtain

$$\mathbb{E}_t\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] \leq \left(1 - \frac{\beta^*}{t-1+\mu}\right)\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + \frac{C_1\gamma}{t-1+\mu}\|w_{t-1} - w^*\|_2^2 + \frac{D_V\delta}{(t-1+\mu)^2}.$$

Taking full expectation on both side yields

$$\mathbb{E}\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] \leq \left(1 - \frac{\beta^*}{t-1+\mu}\right)\mathbb{E}\left[\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2\right] + \frac{C_1\gamma}{t-1+\mu}\mathbb{E}\left[\|w_{t-1} - w^*\|_2^2\right] + \frac{D_V\delta}{(t-1+\mu)^2} \qquad (4)$$

Next, we consider the distance between the iterate and the optimal point. We obtain

$$\begin{aligned} \|w_t - w^*\|_2^2 &\leq \left\|w_{t-1} - \gamma_{t-1}\widehat{G}_{t-1} - w^*\right\|_2^2 \\ &= \|w_{t-1} - w^*\|_2^2 + \gamma_{t-1}^2\left\|\widehat{G}_{t-1}\right\|_2^2 - 2\gamma_{t-1}\widehat{G}_{t-1}^\top(w_{t-1} - w^*) \\ &\leq \|w_{t-1} - w^*\|_2^2 + \gamma_{t-1}^2 D_G^2 - 2\gamma_{t-1}(\widehat{G}_{t-1} - G_{t-1})^\top(w_{t-1} - w^*) - 2\gamma_{t-1}G_{t-1}^\top(w_{t-1} - w^*) \\ &\leq (1 - 2c\gamma_{t-1})\|w_{t-1} - w^*\|_2^2 - 2\gamma_{t-1}(\widehat{G}_{t-1} - G_{t-1})^\top(w_{t-1} - w^*) + \gamma_{t-1}^2 D_G^2 \\ &\leq (1 - 2c\gamma_{t-1})\|w_{t-1} - w^*\|_2^2 + \gamma_{t-1}c\|w_{t-1} - w^*\|_2^2 + \frac{\gamma_{t-1}}{c}\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + \gamma_{t-1}^2 D_G^2 \\ &= \left(1 - \frac{c\gamma}{t-1+\mu}\right)\|w_{t-1} - w^*\|_2^2 + \frac{\gamma}{c(t-1+\mu)}\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2 + \frac{D_G^2\gamma^2}{(t-1+\mu)^2}, \end{aligned}$$

where the third inequality holds because of strong convexity of $F(w)$. Taking expectation we have

$$\mathbb{E}\left[\|w_t - w^*\|_2^2\right] \leq \left(1 - \frac{\gamma c}{t-1+\mu}\right)\mathbb{E}\left[\|w_{t-1} - w^*\|_2^2\right] + \frac{\gamma}{c(t-1+\mu)}\mathbb{E}\left[\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2\right] + \frac{D_G^2\gamma^2}{(t-1+\mu)^2}. \tag{5}$$

The setting of hyperparameters guarantees that $\frac{\beta^*-1}{C_1\gamma} > \frac{\gamma}{c(\gamma c-1)} > 0$. Applying Lemma A.2, there exist $M_1 \geq (\mu+1)\|w_1 - w^*\|_2^2$ and $M_2 \geq (\mu+1)\mathbb{E}\left[\left\|\widehat{G}_1 - G_1\right\|_2^2\right]$, such that

$$\frac{\gamma}{c(\gamma c-1)}M_2 + \frac{D_G^2\gamma^2}{\gamma c - 1} \leq M_1 \leq \frac{\beta^*-1}{C_1\gamma}M_2 - \frac{D_V\delta^2}{C_1\gamma}. \tag{6}$$

Apparently,

$$\begin{aligned}\mathbb{E}\left[\|w_1 - w^*\|_2^2\right] &\leq \frac{M_1}{1+\mu},\\ \mathbb{E}\left[\left\|\widehat{G}_1 - G_1\right\|_2^2\right] &\leq \frac{M_2}{1+\mu}.\end{aligned}$$

We use induction. Assuming

$$\begin{aligned}\mathbb{E}\left[\|w_{t-1} - w^*\|_2^2\right] &\leq \frac{M_1}{t-1+\mu},\\ \mathbb{E}\left[\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2\right] &\leq \frac{M_2}{t-1+\mu},\end{aligned}$$

and according to (4) and (5), we obtain

$$\begin{aligned}\mathbb{E}\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] &\leq \left(1 - \frac{\beta^*}{t-1+\mu}\right)\frac{M_2}{t-1+\mu} + \frac{C_1\gamma}{(t-1+\mu)^2}M_1 + \frac{D_V\delta^2}{(t-1+\mu)^2}\\ &= \frac{t-2+\mu}{(t-1+\mu)^2}M_2 - \frac{1}{(t-1+\mu)^2}\left((\beta^*-1)M_2 - C_1\gamma M_1 - D_V\delta^2\right)\\ &\leq \frac{M_2}{t+\mu},\end{aligned}$$

where the last inequality holds because $(t+\mu)(t-2+\mu) \leq (t-1+\mu)^2$ and the second term is guaranteed by (6) to be negative. We also have

$$\begin{aligned}\mathbb{E}\left[\|w_t - w^*\|_2^2\right] &\leq \left(1 - \frac{\gamma c}{t-1+\mu}\right)\frac{M_1}{t-1+\mu} + \frac{\gamma}{c(t-1+\mu)^2}M_2 + \frac{D_G^2\gamma^2}{(t-1+\mu)^2}\\ &= \frac{t-2+\mu}{(t-1+\mu)^2}M_1 - \frac{1}{(t-1+\mu)^2}\left((\gamma c-1)M_1 - \frac{\gamma}{c}M_2 - D_G^2\gamma^2\right)\\ &\leq \frac{M_1}{t+\mu}.\end{aligned}$$

This completes the proof. □

A.5 Proof to Theorem 3.5

Proof. Equation (3) also holds for non-convex loss functions. Similar to the proof to Theorem 3.4, we further obtain

$$\mathbb{E}\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] \leq \left(1 - \frac{\beta^*}{t-1+\mu}\right)\mathbb{E}\left[\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2\right] + \frac{C_2\gamma}{t-1+\mu}\mathbb{E}\left[\|G_{t-1}\|_2^2\right] + \frac{D_V\delta}{(t-1+\mu)^2}, \tag{7}$$

where $\beta^* = \beta - 2\gamma\sqrt{D_A^2 + L^2}$ and $C_2 = 4\sqrt{D_A^2 + L^2}$. We apply the Lipschitz continuity of ∇F to obtain

$$\begin{aligned} F(w_t) - F(w_{t-1}) &\leq F(w_{t-1} - \gamma_{t-1}\widehat{G}_{t-1}) - F(w_{t-1}) \\ &\leq -\gamma_{t-1} G_{t-1}^\top \widehat{G}_{t-1} + \frac{L\gamma_{t-1}^2 \left\|\widehat{G}_{t-1}\right\|_2^2}{2} \\ &= \frac{\gamma_{t-1}}{2}\|G_{t-1} - \widehat{G}_{t-1}\|_2^2 - \frac{\gamma_{t-1}}{2}\|G_{t-1}\|_2^2 + \frac{\gamma_{t-1}(L\gamma_{t-1} - 1)}{2}\|\widehat{G}_{t-1}\|_2^2 \\ &\leq \frac{\gamma_{t-1}}{2}\|G_{t-1} - \widehat{G}_{t-1}\|_2^2 - \frac{\gamma_{t-1}}{2}\|G_{t-1}\|_2^2, \end{aligned}$$

where the first inequality holds because of Lemma A.3 and the last inequality holds because $\gamma_t L < 1$. This implies

$$\frac{\gamma}{t-1+\mu}\mathbb{E}\left[\|G_{t-1}\|_2^2\right] \leq \frac{\gamma}{t-1+\mu}\mathbb{E}\left[\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2\right] + 2\mathbb{E}\left[F(w_{t-1})\right] - 2\mathbb{E}\left[F(w_t)\right]. \quad (8)$$

The summation over t of (7) reads

$$\begin{aligned} \sum_{t=2}^{T}\mathbb{E}\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] &\leq \sum_{t=2}^{T}\left(1 - \frac{\beta^*}{t-1+\mu}\right)\mathbb{E}\left[\left\|\widehat{G}_{t-1} - G_{t-1}\right\|_2^2\right] \\ &+ C_2\sum_{t=2}^{T}\frac{\gamma}{t-1+\mu}\mathbb{E}\left[\|G_{t-1}\|_2^2\right] + \sum_{t=2}^{T}\frac{D_V\delta}{(t-1+\mu)^2} \\ &= \sum_{t=1}^{T-1}\left(1 - \frac{\beta^*}{t+\mu}\right)\mathbb{E}\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] + C_2\sum_{t=1}^{T-1}\frac{\gamma}{t+\mu}\mathbb{E}\left[\|G_t\|_2^2\right] + \sum_{t=1}^{T-1}\frac{D_V\delta}{(t+\mu)^2} \\ &\leq \sum_{t=1}^{T-1}\left(1 - \frac{\beta^*}{t+\mu}\right)\mathbb{E}\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] + C_2\sum_{t=1}^{T-1}\frac{\gamma}{t+\mu}\mathbb{E}\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] \\ &\quad +2(\mathbb{E}[F(w_1)] - \mathbb{E}[F(w_T)]) + \sum_{t=1}^{T-1}\frac{D_V\delta}{(t+\mu)^2}. \end{aligned}$$

Given that $\beta > 6\gamma\sqrt{D_A^2 + L^2}$, this implies

$$\begin{aligned} \sum_{t=1}^{T-1}\frac{\beta^* - \gamma C_2}{t+\mu}\mathbb{E}\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] &\leq \left(\sum_{t=1}^{T-1}\mathbb{E}\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] - \sum_{t=2}^{T}\mathbb{E}\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right]\right) + 2(\mathbb{E}[F(w_1)] - \mathbb{E}[F(w_T)]) \\ &+ \sum_{t=1}^{T-1}\frac{D_V\delta}{(t+\mu)^2} \\ &= \mathbb{E}\left[\left\|\widehat{G}_1 - G_1\right\|_2^2\right] - \mathbb{E}\left[\left\|\widehat{G}_T - G_T\right\|_2^2\right] + 2(\mathbb{E}[F(w_1)] - \mathbb{E}[F(w_T)]) \\ &+ \sum_{t=1}^{T-1}\frac{D_V\delta}{(t+\mu)^2} \leq M_3, \end{aligned}$$

for some $M_3 > 0$ independent on T. Furthermore

$$\begin{aligned} \sum_{t=1}^{T-1}\frac{\gamma}{t+\mu}\mathbb{E}\left[\|G_t\|_2^2\right] &\leq \sum_{t=1}^{T-1}\frac{\gamma}{t+\mu}\mathbb{E}\left[\left\|\widehat{G}_t - G_t\right\|_2^2\right] + 2(\mathbb{E}[F(w_1)] - \mathbb{E}[F(w_T)]) \\ &\leq \frac{\gamma}{\beta^* - C_2\gamma}M_3 + 2(\mathbb{E}[F(w_1)] - \mathbb{E}[F(w_T)]) := M_4. \end{aligned}$$

We apply the lower bound

$$\sum_{t=1}^{T}\frac{1}{t+\mu} \geq \int_1^{T+1}\frac{1}{t+\mu}dt = \log\left(\frac{T+1+\mu}{1+\mu}\right),$$

to obtain

$$\gamma \log\left(\frac{T+1+\mu}{1+\mu}\right) \min_{1\le t\le T} \mathbb{E}\left[\|\widehat{G}_t - G_t\|_2^2\right] \leq \min_{1\le t\le T} \mathbb{E}\left[\|\widehat{G}_t - G_t\|_2^2\right] \sum_{t=1}^{T} \gamma_t$$
$$\leq \sum_{t=1}^{T} \gamma_t \mathbb{E}\left[\|\widehat{G}_t - G_t\|_2^2\right] \leq M_3,$$

and similarly,

$$\gamma \log\left(\frac{T+1+\mu}{1+\mu}\right) \min_{1\le t\le T} \mathbb{E}\left[\|\widehat{G}_t\|_2^2\right] \leq M_4.$$

Therefore,

$$\min_{1\le t\le T} \mathbb{E}\left[\|\widehat{G}_t - G_t\|_2^2\right] \leq \mathcal{O}\left(\frac{1}{\log T}\right)$$
$$\min_{1\le t\le T} \mathbb{E}\left[\|G_t\|_2^2\right] \leq \mathcal{O}\left(\frac{1}{\log T}\right).$$

□

B Experimental Details

B.1 Values of λ

In the experiments for strongly convex losses, the l_2 regularization parameters λ are decided as described in Section 4.2. In Table 3, λ^* denotes the optimal value of λ and λ_1 to λ_5 denote the sampled values of λ when $\lambda^* = 0$.

Dataset	SmallNorb	CovType	News20	KDD	MNIST
λ^*	0.001	$5 \cdot 10^{-4}$	-	-	-
λ_1	-	-	0.460	0.970	1.600
λ_2	-	-	1.100	0.023	1.400
λ_3	-	-	0.034	0.210	0.034
λ_4	-	-	0.078	0.088	0.680
λ_5	-	-	0.212	0.390	0.460

Table 3: l_2 regularization parameters

B.2 Network Architectures

Next, we describe the FFNs used in the non-convex experiments. Each FFN has two fully connected layers with the dimensions shown in Table 4. The underlying activation function is ReLU.

Dataset	Input dimension	Hidden dimension	Output dimension
CovType	98	10	7
KDD	116	10	2

Table 4: Feedforward network structure

B.3 Additional Results

Additional experimental results are displayed next. Table 5 displays the detailed results of the strongly convex experiments. Figures 4 and 6 display the training loss curves of additional convex and non-convex experiments. Specifically, the tasks are training logistic models on SmallNorb, News20 and FFNs on CovType and KDD. The y-axes show the min-max standardized training loss values. Figure 5 shows the standard deviation of training loss over multiple runs for convex experiments. The standard deviation for KDD is not included as the difference among runs is minor. The y-axes show the min-max standardized standard deviation values.

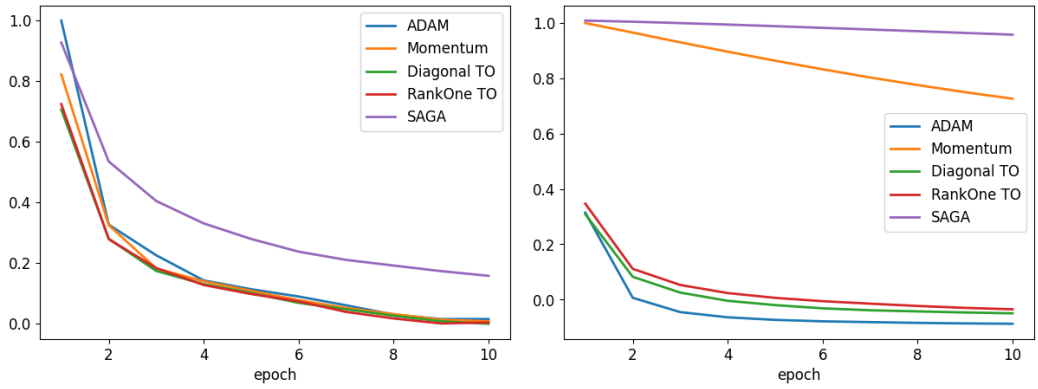

(a) SmallNorb (b) News20

Figure 4: Training loss curves for logistic regression tasks with $\lambda = 0$. Standardization is based on minimum=0.96 and maximum=1.31 in the left figure, and minimum=2.18 and maximum=2.99 in the right figure.

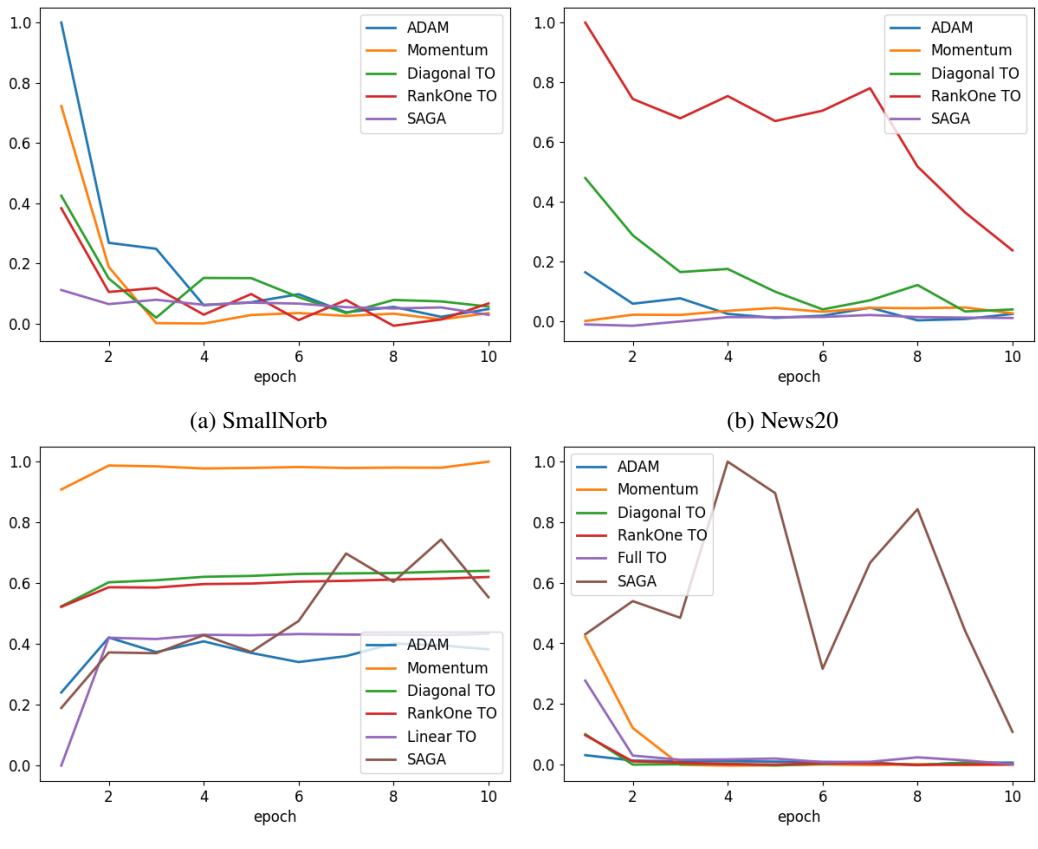

(a) SmallNorb (b) News20

(c) CovType (d) MNIST

Figure 5: Standard deviation of training loss for logistic regression tasks with $\lambda = 0$. Standardization is based on minimum=1.31×10^{-3} and maximum=4.86×10^{-2} in Figure (a), minimum=2.78×10^{-4} and maximum=1.36×10^{-2} in Figure (b), minimum=1.00×10^{-2} and maximum=2.11×10^{-2} in Figure (c), and minimum=1.48×10^{-4} and maximum=6.48×10^{-3} in Figure (d).

Dataset	λ	Momentum	Full TO	Diagonal TO	RankOne TO
SmallNorb	λ^*	-2.3 $(-)$	N/A	$\mathbf{18.0}$ $(+)$	13.0 $(+)$
CovType	λ^*	-1.3 $(\sim)$	$\mathbf{3.3}$ $(\sim)$	-2.0 $(\sim)$	-2.1 $(\sim)$
KDD	λ_1	-85.9 $(-)$	15.0 $(+)$	$\mathbf{16.0}$ $(+)$	14.0 $(+)$
	λ_2	-2.6 $(\sim)$	6.0 $(+)$	$\mathbf{5.8}$ $(+)$	-3.0 $(\sim)$
	λ_3	-21.2 $(-)$	10.0 $(+)$	10.0 $(+)$	$\mathbf{10.2}$ $(+)$
	λ_4	-11.8 $(-)$	$\mathbf{9.0}$ $(+)$	7.6 $(+)$	-6.0 $(\sim)$
	λ_5	-40.2 $(-)$	11.0 $(+)$	$\mathbf{11.2}$ $(+)$	8.7 $(+)$
News20	λ_1	-1.4 $(-)$	N/A	$\mathbf{0.4}$ $(+)$	0.4 $(+)$
	λ_2	0.2 $(+)$	N/A	$\mathbf{0.2}$ $(+)$	0.2 $(+)$
	λ_3	-1.0 $(\sim)$	N/A	$\mathbf{4.4}$ $(+)$	4.3 $(+)$
	λ_4	2.2 $(+)$	N/A	$\mathbf{4.8}$ $(+)$	4.8 $(+)$
	λ_5	2.0 $(\sim)$	N/A	$\mathbf{16.8}$ $(+)$	16.7 $(+)$
MNIST	λ_1	-9.1 $(-)$	$\mathbf{4.2}$ $(+)$	3.9 $(+)$	-1.1 $(-)$
	λ_2	-12.1 $(-)$	$\mathbf{4.4}$ $(+)$	3.9 $(+)$	-1.0 $(-)$
	λ_3	-13.9 $(-)$	-7.1 $(-)$	$\mathbf{0.7}$ $(+)$	$\mathbf{0.7}$ $(+)$
	λ_4	$\mathbf{12.7}$ $(+)$	0.3 $(+)$	0.2 $(+)$	0.2 $(+)$
	λ_5	$\mathbf{3.1}$ $(+)$	0.7 $(+)$	0.6 $(+)$	-3.4 $(-)$

Table 5: Values of ρ for strongly convex tasks (all values $\times 10^{-3}$). Full TO denotes for full version of pseudo-linear TO in Algorithm 2. '+', '$\sim$' and '-' denote for $s < 0.05$, $0.05 \leq s \leq 0.95$, and $s > 0.95$, respectively.

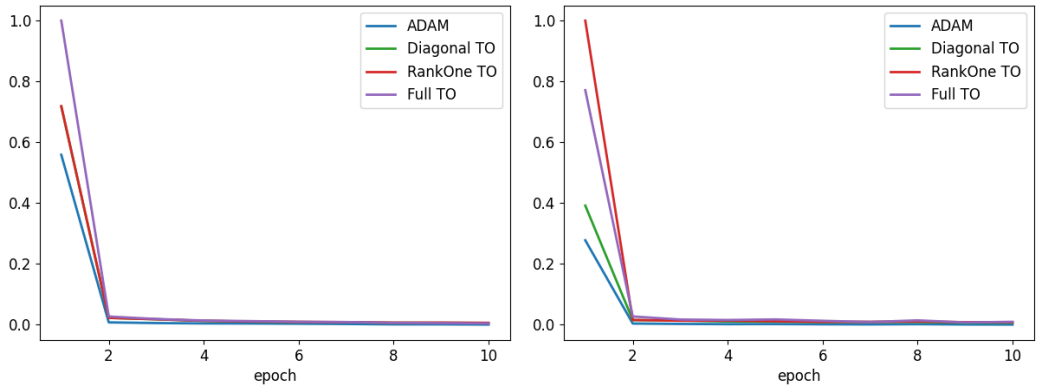

(a) FFN-CovType (b) FFN-KDD

Figure 6: Training loss curves for non-convex tasks. Standardization is based on minimum=1.48 and maximum=1.53 in the left figure, and minimum=0.32 and maximum=0.35 in the right figure.