

DYNAMIC LOSS-BASED SAMPLE REWEIGHTING FOR IMPROVED LARGE LANGUAGE MODEL PRETRAINING

Anonymous authors

Paper under double-blind review

ABSTRACT

Pretraining large language models (LLMs) on vast and heterogeneous datasets is crucial for achieving state-of-the-art performance across diverse downstream tasks. However, current training paradigms treat all samples equally, overlooking the importance or relevance of individual samples throughout the training process. Existing reweighting strategies, which primarily focus on group-level data importance, fail to leverage *fine-grained* instance-level information and do not adapt dynamically to individual sample importance as training progresses. In this paper, we introduce novel algorithms for dynamic, instance-level data reweighting aimed at improving both the efficiency and effectiveness of LLM pretraining. Our methods adjust the weight of each training sample based on its loss value in an online fashion, allowing the model to dynamically focus on more informative or important samples at the current training stage. In particular, our framework allows us to systematically devise reweighting strategies *deprioritizing redundant or uninformative data*, which we find tend to work best. Furthermore, we develop a new theoretical framework for analyzing the impact of loss-based reweighting on the convergence of gradient-based optimization, providing the first formal characterization of how these strategies affect convergence bounds. We empirically validate our approach across a spectrum of tasks, from large-scale LLM pretraining to smaller-scale linear regression problems, demonstrating that loss-based reweighting can lead to faster convergence and improved performance.

1 INTRODUCTION

The rapid advancement of large language models (LLMs) (Brown et al., 2020; Raffel et al., 2020; Touvron et al., 2023; Chowdhery et al., 2023; Achiam et al., 2023; Dubey et al., 2024) has revolutionized natural language processing and artificial intelligence (AI) capabilities. These models, trained on billions or trillions of tokens, exhibit remarkable generalization capabilities across a wide range of downstream tasks. As datasets grow to web-scale proportions and models become increasingly large, the need for more efficient training techniques has become paramount.

Existing approaches to LLM pretraining predominantly involve two phases: heavy data curation (Longpre et al., 2023; Dubey et al., 2024; Wettig et al., 2024; Penedo et al., 2024) and training with uniform sampling on the constructed corpus. Data curation for LLM pretraining typically involves a combination of automated filtering techniques and manual quality checks. For instance, heuristic-based filters are often employed to remove low-quality content and deduplicate data. Some approaches use perplexity-based filtering or auxiliary classifiers (Gao et al., 2020; Penedo et al., 2023) to identify high-quality samples. Manual curation is then applied to refine these filtered datasets, often involving human evaluation of subsets of the data to ensure quality and relevance. However, these approaches face significant limitations due to scalability issues and the static nature of data selection. First, as the pretraining corpora grow to hundreds of billions of tokens, manual curation becomes increasingly infeasible. The sheer volume of data makes it impractical for humans to review even a small fraction of the dataset. Second, data curation methods cannot adapt to the changing importance of samples during the training process, e.g., as the model improves over time certain samples that were useful early in training can become irrelevant in later training stages.

In addition, conventional gradient-based training pipelines often treat all data points as equally informative regardless of their individual importance at current training stage. This uniform sampling

strategy, while simple, overlooks the nuanced diversity within large-scale datasets and potentially wastes computational resources on less informative samples. Recent work (Xie et al., 2023; Fan et al., 2023) has explored group-level reweighting strategies, which adjust the importance of entire domains or groups of data. While these methods have shown promise, they operate at a coarse level, do not fully leverage the fine-grained information available at the instance level, and still lack dynamicity in importance assignment during the training phase.

Given these limitations, this paper seeks to address the following fundamental question:

How can we dynamically leverage instance-level information to accelerate training and improve model performance without incurring significant computational overhead while potentially reducing the need for extensive data curation?

Answering this question presents several intertwined technical challenges, particularly in the context of LLM pretraining. First, unlike in computer vision, where samples are seen repeatedly over the course of training (often hundreds of times), in LLM pretraining, individual samples are typically encountered only once due to the vast size of the datasets. This renders dynamic reweighting approaches primarily devised for multi-epoch training or relying on historical statistics, such as previous loss or gradient norms (Loshchilov & Hutter, 2015; Jiang et al., 2019), inadequate within the context of LLM pretraining. Second, storing previous statistics for individual samples becomes computationally infeasible when working with web-scale datasets, and methods requiring the persistent storage of sample-level data would significantly inflate resource requirements. Furthermore, for a general-purpose LLM, leveraging small hold-out validation sets to compute importance weights through, e.g., bilevel optimization as in Grangier et al. (2023), is ineffective.

Addressing these challenges, our work makes the following significant contributions:

Instance-Level Loss-Based Reweighting Strategies. We introduce and systematically study a variety of instance-level, loss-based reweighting strategies for improving both the efficiency and effectiveness of ML training, especially within the context of LLM pretraining. Each strategy is meticulously designed to achieve specific goals, such as focusing the learning dynamics on different parts of the loss distribution. This builds on recent work that emphasizes the importance of data diversity and sample importance in pretraining and extends previous works on group-level reweighting for LLM pretraining, such as DoReMi Xie et al. (2023), DoGE Fan et al. (2023). It also adds a new dimension by incorporating more fine-grained per-sample dynamics rather than domain-level adjustments. In fact, combining our reweighting schemes with DoGE/DoReMi achieves better or comparable performance on various few-shot reasoning benchmarks when we train on the SlimPajama dataset. Moreover, our study reveals that, in general, strategies down-weighting the importance of low-loss samples tend to consistently yield performance improvements across various scenarios.

New Theoretical Framework. We develop a new theoretical framework for analyzing the effects of loss-based reweighting on training acceleration. To the best of our knowledge, this represents the first explicit characterization of loss reweighting effects within the convergence bounds of gradient methods under widely adopted loss geometries. Our derived convergence bounds provide theoretical justification for the empirical success of downweighting low-loss samples, demonstrating faster convergence under this strategy.

Empirical Validation. We conduct extensive experiments that corroborate our claims and support our theoretical findings. Notably, we demonstrate that the advantages of downweighting low-loss samples are observed across a spectrum of problem scales and complexities: (i) in complex, large-scale problems such as LLMs pretraining, where our approach leads to improved performance and faster convergence; (ii) in simple, small-scale problems like linear regression, highlighting the fundamental nature of our findings.

2 RELATED WORK

Training Data Re-weighting/Selection for LLMs. Several recent studies (Xie et al., 2023; Chen et al., 2023; Fan et al., 2023; Thakkar et al., 2023) have explored various reweighting techniques to enhance the generalization and efficiency of language models pretraining. For instance, Xie et al. (2023) and Fan et al. (2023) optimize the composition of pretraining corpora to achieve better performance across pretraining domains or for out-of-domain generalization. Chen et al. (2023) introduces a framework for ordered skill learning, optimizing data selection based on how

effectively it teaches interdependent skills for continual pretraining and fine-tuning regimes. Although effective, these techniques operate at the group level, whereas our work explores reweighting at the instance level, offering finer control over how individual samples are treated based on their loss values. Furthermore, we demonstrate that combining domain-level methods such as DoReMi (Xie et al., 2023) or DoGE (Fan et al., 2023) with our instance-level reweighting methods results in improved performance across multiple domains. Instance-level reweighting has been used in post-training settings of LLMs (Chen et al., 2024; Jiang et al., 2024). Jiang et al. (2024) boost the self-improvement abilities of LLMs by employing sample reweighting to filter out self-generated data that have correct answers but exhibit high distribution shift. Chen et al. (2024) reweight individual samples during continual training/instruction-tuning to focus on medium loss samples. In contrast, our work systematically studies the effects of various sample-level, loss-based reweighting strategies on the efficiency and effectiveness of LLMs *pretraining*. The approach in Fan & Jaggi (2023) offers a curriculum learning framework that prioritizes samples with a higher learnability score, which is precomputed using another auxiliary model similar to DoReMi and DoGE. While we do not explicitly address curriculum learning in this work, our re-weighting mechanisms naturally allow for implementing a form of loss-based curriculum learning algorithms without the need to train and store additional proxy models as in Fan & Jaggi (2023).

Sample-level Re-weighting as generic ML solution. Sample-level reweighting has been extensively explored in other machine learning areas. For image classification, Loshchilov & Hutter (2015), Jiang et al. (2019), and Katharopoulos & Fleuret (2018) are pioneering approaches that prioritize samples based on loss values or gradient norms with the goal of accelerating training speed. Although these methods emphasize the importance of selecting high-loss samples during training, they typically require additional forward/backward passes on each training sample and are primarily designed for multi-epoch training, limiting their applicability with the vast pretraining corpus used for LLMs. Using bilevel optimization, Grangier et al. (2023) adapt the data distribution during training to focus more on relevant samples for a target data distribution. Similarly, Ren et al. (2018) employ a meta-learning approach to adjust sample weights based on validation set performance, which excels in noisy and imbalanced datasets. In this work, we introduce and investigate various light-weight, loss-based reweighting techniques that add little to no computational overhead compared to uniform sampling and do not require any nested optimization routines, often arising with bilevel optimization or/and meta-learning. Sample-level reweighting has also been explored in the context of adversarial machine learning (Zhang et al., 2020; Liu et al., 2021; Zeng et al., 2021; Sow et al., 2023), domain adaptation (Jiang & Zhai, 2007; Fang et al., 2020), data augmentation (Yi et al., 2021), and imbalanced classification (Qi et al., 2021; Ren et al., 2018). Our reweighting mechanisms also have the potential to be studied under these contexts, which we leave as future work.

3 PRELIMINARIES: AUTOREGRESSIVE LANGUAGE MODELING & GOAL

Large language models (LLMs) are typically trained using autoregressive language modeling, where the objective is to predict the next token in a sequence given the previous tokens. Formally, given a sample sequence of tokens $\mathbf{x} = (x_1, x_2, \dots, x_T)$, where x_t represents the t -th token in the sequence, the LLM model parameterized by $\theta \in \mathbb{R}^d$ is trained to maximize the likelihood of the sequence:

$$P_{\text{LLM}}(\mathbf{x}; \theta) = \prod_{t=1}^T P(x_t | x_1, \dots, x_{t-1}; \theta).$$

where $P(x_t | x_1, \dots, x_{t-1}; \theta)$ denotes the conditional probability of generating token x_t given θ and all previously seen tokens x_1, \dots, x_{t-1} . In practice, autoregressive LLMs typically compute the loss $f(\mathbf{x}; \theta)$ for sample \mathbf{x} using the negative log-likelihood (NLL) of the predicted tokens.

$$f(\mathbf{x}; \theta) = -\log P_{\text{LLM}}(\mathbf{x}; \theta) = -\sum_{t=1}^T \log P(x_t | x_1, \dots, x_{t-1}; \theta).$$

Based on these sample losses, conventional SGD-like algorithms will then update the model parameters with the average gradient over all samples in a batch \mathcal{B} at each training step t , i.e.,

$$\theta^{t+1} = \theta^t - \frac{\eta}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \nabla f(\mathbf{x}_i; \theta^t), \quad (1)$$

which results in each sample contributing equally to the overall training update. While this approach enjoys simplicity, it inherently treats all data points as equally informative, regardless of their individual relevance or difficulty. This uniform treatment overlooks the nuanced diversity within the data, which is particularly critical for current LLMs trained on vast and heterogeneous web-scale datasets. In such scenarios, the ability to dynamically differentiate between high-value and lesser-value data samples at different training stages is paramount for optimizing training efficiency and model performance.

In contrast, our work introduces and benchmarks a variety of simple, lightweight reweighting strategies that can be seamlessly integrated into existing training pipelines with little to no computational overhead. These strategies aim to address the inherent limitations of the standard uniform averaging method by assigning dynamic importance to individual data samples based on their loss values, thus providing a more refined control over the training process. Ultimately, our objective is twofold: (i) to improve the efficiency and effectiveness of LLM training by allowing the model to dynamically focus on more relevant or informative data at different training stages and (ii) to potentially provide a scalable, automated approach to online data selection during pretraining. This ambitiously seeks to reduce the need for extensive manual data curation and selection efforts, which are increasingly becoming infeasible given the growing size and diversity of pretraining corpora. We envision these reweighting strategies as a foundation for more efficient and adaptive data utilization in LLM pretraining, ultimately reducing the time, cost, and complexity associated with model pretraining.

4 APPROACH

4.1 TECHNICAL CHALLENGE: DATA REWEIGHTING FOR LLMs MUST BE DYNAMIC AND FULLY ONLINE

The core of our framework is to *dynamically* reweight individual training samples based on their importance at different training stages. Although sample importance seems static—intuitively, “garbage data is garbage data” regardless of model or training stage—data reweighting should be dynamic. For instance, similar or duplicated samples in training corpora may be useful early in training, but should be deprioritized once the model captures their patterns. Similarly, a hard but useful sample could be assigned a larger weight, however, as the model learns it could be beneficial to reduce its weight. Based on these observations, we update the model at each step t using

$$\theta^{t+1} = \theta^t - \eta_t \sum_{i \in \mathcal{B}} w(\mathbf{x}_i; \theta^t) \nabla f(\mathbf{x}_i; \theta^t), \quad (2)$$

where $\sum_{i \in \mathcal{B}} w(\mathbf{x}_i; \theta^t) = 1$ and the weight $w(\mathbf{x}_i; \theta^t)$ is dynamically assigned based on sample \mathbf{x}_i and current model θ^t . However, designing effective dynamic reweighting methods for LLM pretraining involves several inherent complexities. In LLM pretraining, each sample is typically seen only once due to the vast size of the datasets, which contrasts with problems where repeated exposure to data is common, such as in computer vision problems. This makes methods that depend on accumulating historical data, such as previous loss or gradient norms (Loshchilov & Hutter, 2015; Jiang et al., 2019), unsuitable. Furthermore, the sheer scale of these large corpus makes it computationally prohibitive to track and store sample-specific statistics, as it would significantly increase resource requirements. Additionally, approaches that rely on hold-out validation sets to compute sample weights, as in Grangier et al. (2023), are also irrelevant for pretraining a general purpose LLM. These unique technical challenges necessitate the development of a fully online, lightweight reweighting approach that scales efficiently with the size of modern LLM training corpus, which we discuss next.

4.2 LOSS-BASED REWEIGHTING STRATEGIES

To address the aforementioned challenges, we propose a novel framework of loss-based sample reweighting strategies, motivated by the following theorem. Theorem 1 characterizes the effects of loss-based reweighting on the convergence bound of reweighted full gradient descent when the sample losses $f(\mathbf{x}_i; \theta^t)$ are convex. The extension to the stochastic setting and analysis for other types of loss geometries is deferred to Section 5. Our full statement of Theorem 1 and detailed proofs for all theoretical statements can be found in the appendix.

Theorem 1. Consider M data points and let each loss $f(\mathbf{x}_i; \cdot)$ be convex. Further, assume the interpolation regime holds, i.e., $\exists \theta^* \in \mathbb{R}^d$ such that $\theta^* \in \arg \min_{\theta \in \mathbb{R}^d} f(\mathbf{x}_i; \theta) \forall i$. Then, for a

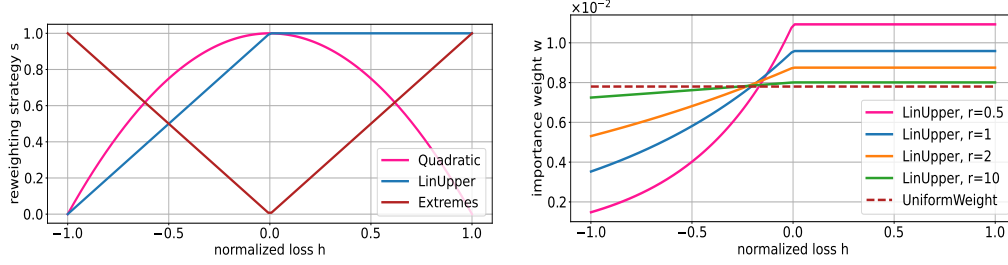


Figure 1: **Left:** Geometric curves of the different reweighting functions. **Right:** Shape of the **LinUpper** strategy after applying Eq. (4) on top of it for different values of r . These plots are obtained for a batch of 128 uniformly drawn losses. As r increases, **LinUpper** converges to the uniform averaging method.

reweighting scheme that satisfies $\max_{i \in \{1, \dots, M\}} w(\mathbf{x}_i; \theta^t) \leq 2/M$, we have

$$\frac{1}{M} \sum_{i=1}^M f(\mathbf{x}_i; \bar{\theta}^T) - \frac{1}{M} \sum_{i=1}^M f(\mathbf{x}_i; \theta^*) \leq \mathcal{O} \left(\frac{\|\theta^0 - \theta^*\|^2}{T} \right) + \frac{1}{T} \sum_{t=0}^{T-1} \delta^t, \quad (3)$$

where $\delta^t = \sum_{i=1}^M \left(\frac{1}{M} - w(\mathbf{x}_i; \theta^t) \right) (f(\mathbf{x}_i; \theta^t) - f(\mathbf{x}_i; \theta^))$ and $\bar{\theta}^T = \frac{1}{T} \sum_{t=0}^{T-1} \theta^t$.*

Theorem 1 illustrates how setting the importance weight $w(\mathbf{x}_i; \theta^t)$ based on the loss gap $f(\mathbf{x}_i; \theta^t) - f(\mathbf{x}_i; \theta^*)$ influences the convergence bound. Specifically, we make the following key observations: (i) fixing $w(\mathbf{x}_i; \theta^t) = \frac{1}{M} \forall i$, leads to $\delta_t = 0$ and we recover the traditional convergence bound of gradient descent for convex functions; (ii) assigning larger weights to the smaller loss gaps leads to $\delta_t \geq 0$, which yields a looser convergence bound; (iii) to achieve a better bound ($\delta_t \leq 0$), one should down-weight the importance of samples with low loss gaps. Further, note that Eq. (3) critically holds only for reweighting schemes satisfying $w(\mathbf{x}_i; \theta^t) \leq 2/M$, which puts an upper bound on the maximum weight that can be assigned to a single data point, and hence, eliminates worst-case importance assignment strategies—such as those based on distributionally robust optimization (Qian et al., 2019; Qi et al., 2021; Kumar et al., 2023)—which are prone to overfitting to outliers, a particularly problematic issue when dealing with the large-scale, noisy corpora used in LLM pretraining.

Based on the insights from Theorem 1 on the relationship between loss and importance weight, we study a series of reweighting strategies, each carefully designed to focus the learning dynamics on different parts of the loss distribution during training. Let w_i be the importance weight for sample \mathbf{x}_i where we drop the dependency on θ^t for clarity. We first normalize the individual sample losses into a bounded range $[-\delta, \delta]$. This step crucially ensures the loss values are scaled consistently and amenable to the subsequent weighting functions. We then apply the following analytical functions on the normalized loss h_i . Figure 1 shows the geometric curves of these different functions.

- **Linear Upper-Bound Strategy (LinUpper).** In this strategy, the weight is proportional to the normalized loss but is capped at a predefined δ value, ensuring that outliers do not dominate the training process. The functional form is $s_i := \min\{h_i + \delta, \delta\}$.
- **Quadratic Strategy (Quadratic).** This strategy prioritizes samples with moderate loss values while down-weighting both low-loss (easy or repetitive) and high-loss (potential outliers) samples by applying a quadratic function $s_i := \delta \left(1 - \frac{h_i^2}{\delta^2} \right)$.
- **Extremes-Based Strategy (Extremes).** This scheme emphasizes both the hardest (high-loss) and the easiest (low-loss) samples by applying $s_i := |h_i|$, ensuring that samples at the extremes of the loss distribution receive higher importance weights.

To further enhance training stability and effectiveness, we add a curriculum-based adjustment mechanism on top of the strategic sample weight s_i by controlling the sharpness of the importance weighting using a temperature parameter r . The final weight w_i is computed as:

$$w_i = \frac{e^{s_i/r}}{\sum_j e^{s_j/r}}, \quad (4)$$

Algorithm 1 Fully Online Instance Reweighting

```

1: Input: stepsize  $\eta$ , initial model parameters  $\theta_0 \in \mathbb{R}^d$ , temperature  $\{r_t\}_{t=0}^T$ ,  $f_{\min}$  and  $f_{\max}$ .
2: for  $t = 0, 1, 2, \dots, T - 1$  do
3:   Draw a minibatch of data samples  $\mathcal{B} = \{\mathbf{x}_i\}_{i=1}^b$ 
4:   Run forward pass to compute losses  $\{f_{i,t}\}_{i=1}^b$  for all samples in  $\mathcal{B}$ 
5:   Normalize losses into interval  $[-\delta, \delta]$  using linear transform:  $h_{i,t} = \frac{2(f_{i,t} - f_{\min})}{f_{\max} - f_{\min}} - 1$ 
6:   Apply reweighting strategy to normalized losses to obtain  $\{s_{i,t}\}_{i=1}^b$  // {LinUpper, Quadratic, Extremes}
7:   Apply curriculum adjustment using Eq. (4) to obtain sample weights  $\{w_{i,t}\}_{i=1}^b$ 
8:   Update  $\theta^{t+1} = \theta^t - \eta \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_{i,t}$ 
9: end for

```

where r is annealed during training. Early in training, when the model is still learning fundamental patterns and losses are less informative, we employ a more uniform weighting scheme to ensure diverse feature learning by adopting a large r (see Figure 1). As training progresses, we decrease the value of r to enact full use of our reweighting strategies. Algorithm 1 depicts our approach. It is important to note that our reweighting approach adds nearly zero computational overhead. During the forward pass, we already compute the sample losses, which are then used to derive the importance weights through simple analytical transformations. Since no extra loss or gradient computations are required, our reweighting strategies can be seamlessly integrated into standard training pipelines without incurring significant additional costs in terms of time or resources.

Optimal strategy minimizing δ^t . We next demonstrate that applying Eq. (4) on top of our strategy **LinUpper**, in fact, corresponds to the optimal strategy that minimizes a KL-divergence regularized version of δ^t in Theorem 1.

Proposition 1. *The optimal strategy that minimizes the KL-divergence regularized δ^t , i.e., $\delta^t + r \sum_{i=1}^M w_i \log(Mw_i)$, is given by*

$$w_i = C \min \left\{ \exp \left(\frac{h_i}{r} \right), \frac{2}{M} \right\}, \quad (5)$$

where C is a normalizing constant that ensures $\sum_i w_i = 1$.

Proof. Proof relegated to Appendix C.2. \square

As shown in Theorem 1 and Proposition 1, our experiments confirm that the **LinUpper** reweighting scheme indeed achieves faster convergence and improved performance across a spectrum of LLMs pretraining scales. The incorporation of KL-divergence regularization in the optimization of δ^t prevents extreme solutions that focus only on the highest-loss samples. Without this regularization term, the optimal solution becomes trivial: it assigns a weight of $\frac{2}{M}$ to the $\frac{M}{2}$ samples with the highest losses and zero weight to the remaining samples. However, such a strategy discards half of the dataset and performs poorly in practice. In contrast, the **LinUpper** method offers a balance between focusing on high-loss samples and maintaining the diversity of data utilized during training.

5 ANALYSIS OF LOSS-BASED REWEIGHTING ON CONVERGENCE BOUND

In this section, we analyse the effects of loss-based reweighting on the convergence bound of stochastic gradient methods under convex and nonconvex loss geometries. Consider the problem of minimizing a finite sum of M objective functions:

$$\min_{\theta \in \mathbb{R}^d} \left\{ f(\theta) := \frac{1}{M} \sum_{i=1}^M f_i(\theta) \right\} \quad (6)$$

where each $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ corresponds to a sample loss function $f(\mathbf{x}_i; \cdot)$. We make the following widely adopted assumptions:

Assumption 1 (Convexity). *Each function f_i is convex, i.e., for all $\theta, \theta' \in \mathbb{R}^d$ and $i \in 1, \dots, M$:*

$$f_i(\theta') \geq f_i(\theta) + \langle \nabla f_i(\theta), \theta' - \theta \rangle. \quad (7)$$

Assumption 2 (L -smoothness). Each function f_i is L_i -smooth. Formally, for all $\theta, \theta' \in \mathbb{R}^d$ and $i \in 1, \dots, M$:

$$\|\nabla f_i(\theta') - \nabla f_i(\theta)\| \leq L_i \|\theta' - \theta\|, \quad (8)$$

and $L = \max_{i \in [M]} L_i$.

Assumption 3 (Interpolation condition). There exists a global optimum $\theta^* \in \mathbb{R}^d$ such that:

$$\theta^* \in \arg \min_{\theta \in \mathbb{R}^d} f_i(\theta), \quad \forall i \in \{1, \dots, M\} \quad (9)$$

Assumption 3 is widely adopted in modern ML settings (Dar et al., 2021; Radhakrishnan et al., 2020; Jhunjhunwala et al., 2023), particularly when analyzing overparameterized deep neural networks. This ensures that all individual sample objectives can be minimized simultaneously, effectively eliminating adversarial trade-offs where minimizing one objective could worsen another.

Next, we characterize the effects of loss-based reweighting on the convergence bound of reweighted minibatch SGD when the sample losses f_i are convex. For ease of notation, we denote $w(\mathbf{x}_i; \theta^t)$ by $w_{i,t}$ in this section and the corresponding proofs in the Appendix.

Theorem 2. Let Assumptions 1, 2, and 3 hold. Consider a minibatch of size $|\mathcal{B}| = b$ with reweighting scheme satisfying $\max_{i \in \mathcal{B}} w_{i,t} \leq 2/b$. Then,

- **minibatch SGD:** for $\eta = \frac{1}{8L\sqrt{T}}$, we have

$$\mathbb{E}[f(\bar{\theta}^T) - f(\theta^*)] \leq \frac{8L\|\theta^0 - \theta^*\|^2}{\sqrt{T}} + \frac{1}{T} \sum_{t=0}^{T-1} \delta_t, \quad (10)$$

- **minibatch SGD with momentum:** for $\eta = \frac{1}{8L\sqrt{T+1}}$, $\lambda_t > 0$, we have

$$\mathbb{E}[f(\theta^T) - f(\theta^*)] \leq \frac{8L\|\theta^0 - \theta^*\|^2}{\sqrt{T+1}} + \frac{2}{T+1} \sum_{t=0}^{T-1} \delta_t + \frac{2}{T+1} \sum_{t=0}^{T-1} \lambda_t \mu_t, \quad (11)$$

where $\delta_t = \mathbb{E}[\sum_{i \in \mathcal{B}} (\frac{1}{b} - w_{i,t}) (f_i(\theta^t) - f_i(\theta^*))]$, $\mu_t = \mathbb{E}[\sum_{i \in \mathcal{B}} (\frac{1}{b} - w_{i,t}) (f_i(\theta^t) - f_i(\theta^{t-1}))]$, and $\bar{\theta}^T = \frac{1}{T} \sum_{t=0}^{T-1} \theta^t$.

Proof. Proof provided in Appendix C.3. \square

We note similar observations as in Theorem 1 and again down-weighting the importance of samples with low loss gaps leads to a better bound. Typically, the minibatch gradient with uniform weights is unbiased, $\mathbb{E}[\sum_{i \in \mathcal{B}} \frac{1}{b} \nabla f_i(\theta^t)] = \nabla f(\theta^t)$. However, in our setting, the importance weights are functions of the sample functions, hence, leading to biased gradients and the analysis of Theorem 2 being more involved. Update equations with heavy ball momentum are provided in eq. (39) (see Appendix C.3). Further, we provide the analysis for minibatch SGD when the sample functions are non-convex in Appendix C.4.

6 EXPERIMENTAL EVALUATION

6.1 TRAINING AND EVALUATION SETUP

We conduct our experiments using decoder-only transformer models (Vaswani, 2017; Radford et al., 2019) with parameter sizes of 120M, 210M, and 300M, which we refer to as, respectively, GPT2-mini, GPT2-small, and GPT2-medium. We train on the SlimPajama (Soboleva et al., 2023) corpus that includes seven diverse domains: Common Crawl (CC), C4, GitHub, StackExchange, Book, Arxiv, and Wikipedia. In the first step, we pre-train the three GPT2 models on all seven domains where we compare our sample-level reweighting methods (**LinUpper**, **Quadratic**, **Extremes**) against the uniform averaging baseline in which each sample contributes equally. In a second step, we assess the impact of combining our reweighting techniques with existing domain-level reweighting methods, specifically DoGE and DoReMi, to demonstrate the potential for further performance gains when

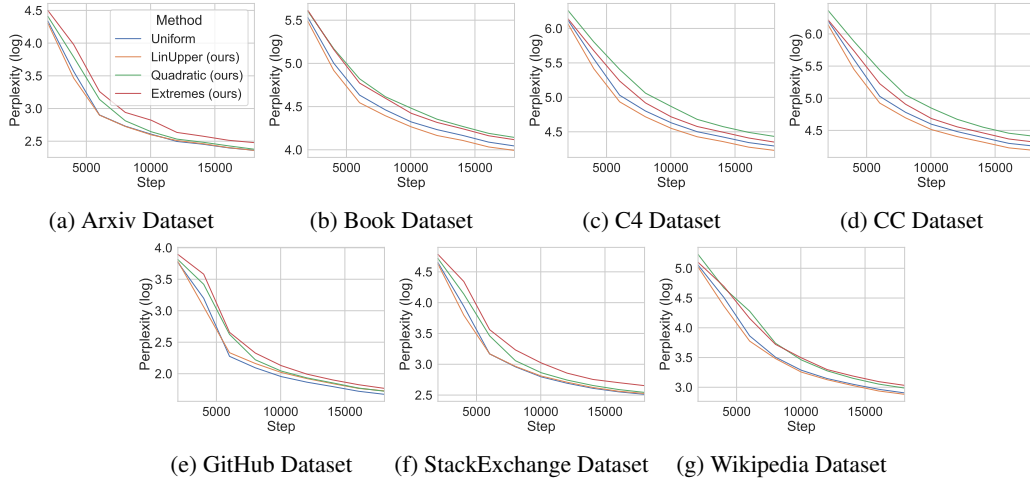


Figure 2: Per-domain perplexities on hold-out validation sets under the uniform domain sampling setting for the GPT2-medium model. Our reweighting strategy **LinUpper** strategy achieves better or at least comparable perplexity on 5 out of 7 domains.

fine-grained sample reweighting is added on top of domain-level adjustments. For the pretraining of all models, we use a minibatch size of 48, a sequence length of 512, and train for a total of 20,000 steps. For all reweighting methods, we set r to 100 for the first 1,000 steps and then reduce it to 1 for the remaining steps. Model architecture details can be found in Appendix A.2. Our codebase is publicly available.¹

We measure performance using two types of metrics:

Perplexity on Training Distribution. We evaluate per-domain perplexity and average perplexity on all training domains using hold-out validation sets. By doing so, we aim to generate insights on the sample efficiency during pretraining, i.e., how well a model can learn from a given dataset and adapt to the data distribution.

Few-Shot Reasoning Ability. To assess generalization beyond the pretraining distribution, we evaluate models on a series of diverse reasoning tasks in a 5-shot setting, including LogiQA Liu et al. (2020), LogiQA 2 Liu et al. (2023), SciQ Welbl et al. (2017), and PiQA Bisk et al. (2020).

Table 1: Per-domain perplexities on hold-out validation datasets under the *uniform* domain sampling setting for the GPT2 models. Across a wide variety of domains, our **LinUpper** method outperforms the uniform baseline.

Model	Method	Benchmark							Mean
		Arxiv	Book	C4	CC	GitHub	StackExchange	Wikipedia	
GPT2-mini	Uniform	2.48	4.21	4.47	4.44	1.84	2.66	3.14	3.32
	LinUpper (ours)	2.49	4.16	4.41	4.38	1.89	2.67	3.12	3.30
	Quadratic (ours)	2.48	4.26	4.55	4.53	1.87	2.67	3.17	3.36
	Extremes (ours)	2.54	4.23	4.47	4.44	1.90	2.74	3.19	3.36
GPT2-small	Uniform	2.36	4.04	4.29	4.26	1.67	2.51	2.90	3.15
	LinUpper (ours)	2.36	3.99	4.23	4.19	1.73	2.53	2.88	3.13
	Quadratic (ours)	2.38	4.14	4.43	4.41	1.72	2.54	2.99	3.23
	Extremes (ours)	2.48	4.12	4.35	4.32	1.77	2.65	3.03	3.24
GPT2-medium	Uniform	2.34	4.02	4.27	4.22	1.66	2.49	2.88	3.13
	LinUpper (ours)	2.34	3.96	4.20	4.15	1.72	2.50	2.85	3.10
	Quadratic (ours)	2.35	4.11	4.40	4.36	1.70	2.51	2.93	3.19
	Extremes (ours)	2.45	4.08	4.31	4.27	1.76	2.62	3.00	3.21

¹Anonymized link: <https://anonymous.4open.science/r/iclr-loss-reweighting/> or alternatively in the submission’s supplementary material.

Table 2: 5-shot Performance evaluations under *uniform* domain sampling for the GPT2-medium model. We report the accuracy for every benchmark. The highest accuracy per benchmark is highlighted in **bold**. The **LinUpper** reweighting strategy consistently yields competitive results, outperforming the other compared baselines on 3 out of 4 tasks.

Benchmark Accuracy	Method			
	Uniform	LinUpper (ours)	Quadratic (ours)	Extremes (ours)
LogiQA	25.7	27.9	25.5	26.5
LogiQA 2	27.5	27.6	28.6	27.7
SciQ	49.0	51.8	48.7	49.6
PiQA	55.5	56.2	54.6	55.2
<i>Mean</i>	39.4	40.9	39.3	39.8

6.2 COMPARISONS UNDER UNIFORM DOMAIN SAMPLING

Comparisons on Pretraining Distribution. Figure 2 shows the per-domain perplexity for the GPT2-medium model on all seven domains. We provide the final per-domain perplexity results for the GPT2-mini and GPT2-small models in Table 1. The perplexity plots for GPT2-mini and GPT2-small can be found in Appendix A.3. Figure 2 shows that our proposed **LinUpper** strategy outperforms or matches the other baseline in 5 out of 7 domains, with particularly notable improvements in CC, C4, and Book domains. **LinUpper** achieves significantly lower perplexity in these domains, demonstrating its effectiveness at handling noisier data sources. This suggests that down-weighting low-loss samples –central to **LinUpper**– helps reduce the effects of redundant or similar data, and thus accelerates convergence as supported by our theoretical findings. We also note competitive results for **LinUpper** method at the smaller scale models in Table 1, but the advantage is more significant in the larger GPT2-medium model.

Few-Shot Evaluations. Table 2 provides the 5-shot evaluations on 4 different tasks for the GPT2-medium model. The **LinUpper** reweighting strategy consistently yields competitive results, outperforming the other compared baselines on 3 out of 4 tasks. These results further validate the general effectiveness of **LinUpper** across different types of reasoning challenges, including logical inference (LogiQA), physical commonsense reasoning (PiQA), and scientific reasoning (SciQ). Notably, **Quadratic** performs best in LogiQA 2 (28.6%) , indicating that reducing the effects of both high-loss (potential outliers) and low-loss (potentially redundant data) samples can sometimes improve generalization performance even though it generally does not achieve best efficiency for learning the pretraining distribution. This suggests that the **Quadratic** scheme could be useful when training under noisier and less curated datasets. Also not that the **Extremes** strategy consistently lags behind the other methods, which further supports the advantage of reducing the importance of low-loss samples.

Table 3: 5-shot performance evaluations under *non-uniform* domain sampling for the GPT2-medium model. Our **LinUpper** instance reweighting approach creates synergies with state-of-the-art domain reweighting techniques.

Benchmark Accuracy	Method			
	DoGE	DoGE + LinUpper	DoReMi	DoReMi + LinUpper
LogiQA	27.2	28.6	27.2	27.6
LogiQA 2	27.5	28.0	27.7	27.9
SciQ	52.8	53.2	53.3	54.5
PiQA	55.8	56.3	55.7	56.1
<i>Mean</i>	40.8	41.5	41.0	41.5

6.3 COMPARISONS UNDER NON-UNIFORM DOMAIN SAMPLING

The performance of the compared methods under non-uniform domain sampling setting are presented in the Table 3. The results clearly demonstrate the value of incorporating our **LinUpper** strategy on top of existing domain-level reweighting methods such as DoGE and DoReMi. Across all tasks, **LinUpper** consistently improves the performance of both domain reweighting baselines, reinforcing the effectiveness of dynamically adjusting sample weights at the instance level in conjunction

Table 4: Llama-1.4B average performance across LUR and QA benchmarks. Overall, our **LinUpper** method improves results for 14 out of 19 tasks with an average performance gain of 1.74% among LUR tasks and 1.72% among QA tasks. Please see Tables 8 and 10 in Appendix A.3 for performance on each task.

Benchmark Category	Uniform	LinUpper (ours)	Difference
<i>LUR Tasks Mean Performance</i>	46.65	48.40	1.74
<i>QA Tasks Mean Performance</i>	42.35	44.07	1.72

Table 5: Intermediary results of the Llama-7B model (checkpoints obtained after training on 80B tokens). Overall, our **LinUpper** method improves results for 15 out of 19 tasks with an average performance gain of 4.26% among LUR tasks and 0.53% among QA tasks.

Benchmark Category	Uniform	LinUpper (ours)	Difference
<i>LUR Tasks Mean Performance</i>	49.89	54.15	4.26
<i>QA Tasks Mean Performance</i>	45.86	46.39	0.53

with domain-level reweighting. In particular, our reweighting strategy **LinUpper** provides notable improvements for LogiQA task, boosting accuracy from 27.2% to 28.6% with DoGE and from 27.2% to 27.6% with DoReMi. The gains in LogiQA 2 are similar, where our reweighting method **LinUpper** consistently improves the performance of both methods. For scientific reasoning (SciQ) and physical commonsense (PiQA) tasks, **LinUpper** strategy also enhances both baseline methods, with the most substantial improvement observed in SciQ (from 52.8% to 53.2% with DoGE and from 53.3% to 54.5% with DoReMi). These results indicate that **LinUpper** is particularly effective in tasks requiring nuanced reasoning and factual knowledge.

6.4 ADDITIONAL EXPERIMENTS ON 1.4B AND 7B LLAMA MODELS

We have conducted **new experiments to train 1.4B (billion) and 7B models with Llama architecture** on randomly sampled subsets of the FineWeb 15T dataset. For the 1.4B model, we have completed training the models on 100B tokens. For the 7B model, we are currently (at the moment of submitting this revision of our paper) at 80B tokens seen and report the benchmark results on the latest checkpoints. For these larger models, we use the QA benchmarks and additionally employ language understanding and reasoning (LUR) tasks since their capacity allows for more complex tasks. Please see **Appendix A for more details, including hyperparameters, model architectural details, and performance on the full list of benchmarks**. Tables 4 and 5 shows the average performance among LUR and QA tasks (see Tables 8 and 10 in Appendix A.3 for performance on each task) for the 1.4B and 7B model, respectively. For the Llama-1.4B, we see that our **LinUpper** strategy improves the model performance for language understanding and reasoning tasks by 1.74% on average and 1.72% on average for question-answering tasks. In particular, we note that our **LinUpper** method improves performance for 14 out of 19 total tasks (see Table 8). For the Llama-7B, our **LinUpper** method improves results for 15 out of 19 (see Table 10) tasks with an average performance gain of 4.26% among LUR tasks and 0.53% among QA tasks. These improvements are particularly noteworthy given the scale of these Llama models, highlighting the scalability and effectiveness of our method.

7 CONCLUSION

In this paper, we introduce a novel sample-level reweighting framework aimed at improving the efficiency and effectiveness of large language model (LLM) pretraining. By dynamically adjusting the importance of individual samples based on their loss values, our approach overcomes the limitations of traditional uniform averaging methods and adds a new dimension to existing domain-level reweighting methods by incorporating more fine-grained sample-level dynamics. Through extensive theoretical and empirical validation, we demonstrate that down-weighting low-loss samples accelerates convergence and improves performance. Our experiments show that our proposed **LinUpper** strategy consistently outperforms uniform sampling on common LLM reasoning and commonsense benchmarks. Similarly, our approach creates synergies with existing domain reweighting techniques, which underpins the importance of sample-level dynamics. We observe that the benefits of our method are more pronounced in larger models, while smaller models show less significant improvements on benchmarks. This could be due to the fact that smaller models may have limited capacity to fully exploit the advantages of our approach. However, our overall findings highlight the potential of loss-based reweighting strategies to optimize LLM pretraining both in training efficiency and in model evaluation performance.

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 7432–7439, 2020.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Mayee F Chen, Nicholas Roberts, Kush Bhatia, Jue WANG, Ce Zhang, Frederic Sala, and Christopher Re. Skill-it! a data-driven skills framework for understanding and training language models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=IoizwO1NLf>.
- Xuxi Chen, Zhendong Wang, Daouda Sow, Junjie Yang, Tianlong Chen, Yingbin Liang, Mingyuan Zhou, and Zhangyang Wang. Take the bull by the horns: Hard sample-reweighted continual training improves llm generalization. *arXiv preprint arXiv:2402.14270*, 2024.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.
- Yehuda Dar, Vidya Muthukumar, and Richard G Baraniuk. A farewell to the bias-variance tradeoff? an overview of the theory of overparameterized machine learning. *arXiv preprint arXiv:2109.02355*, 2021.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Simin Fan and Martin Jaggi. Irreducible curriculum for language model pretraining. *arXiv preprint arXiv:2310.15389*, 2023.
- Simin Fan, Matteo Pagliardini, and Martin Jaggi. DOGE: Domain reweighting with generalization estimation. In *Second Agent Learning in Open-Endedness Workshop*, 2023. URL <https://openreview.net/forum?id=q1KqsqwYXm>.
- Tongtong Fang, Nan Lu, Gang Niu, and Masashi Sugiyama. Rethinking importance weighting for deep learning under distribution shift. *Advances in neural information processing systems*, 33: 11996–12007, 2020.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- David Grangier, Pierre Ablin, and Awni Hannun. Adaptive training distributions with scalable online bilevel optimization. *arXiv preprint arXiv:2311.11973*, 2023.
- Divyansh Jhunjhunwala, Shiqiang Wang, and Gauri Joshi. Fedexp: Speeding up federated averaging via extrapolation. In *The Eleventh International Conference on Learning Representations*, 2023.
- Angela H Jiang, Daniel L-K Wong, Giulio Zhou, David G Andersen, Jeffrey Dean, Gregory R Ganger, Gauri Joshi, Michael Kaminsky, Michael Kozuch, Zachary C Lipton, et al. Accelerating deep learning by focusing on the biggest losers. *arXiv preprint arXiv:1910.00762*, 2019.
- Chunyang Jiang, Chi-min Chan, Wei Xue, Qifeng Liu, and Yike Guo. Importance weighting can help large language models self-improve. *arXiv preprint arXiv:2408.09849*, 2024.

- Jing Jiang and ChengXiang Zhai. Instance weighting for domain adaptation in nlp. *ACL*, 2007.
- Angelos Katharopoulos and François Fleuret. Not all samples are created equal: Deep learning with importance sampling. In *International conference on machine learning*, pp. 2525–2534. PMLR, 2018.
- Ramnath Kumar, Kushal Majmundar, Dheeraj Nagaraj, and Arun Sai Suggala. Stochastic re-weighted gradient descent via distributionally robust optimization. *arXiv preprint arXiv:2306.09222*, 2023.
- Feng Liu, Bo Han, Tongliang Liu, Chen Gong, Gang Niu, Mingyuan Zhou, Masashi Sugiyama, et al. Probabilistic margins for instance reweighting in adversarial training. *Advances in Neural Information Processing Systems*, 34:23258–23269, 2021.
- Hanmeng Liu, Jian Liu, Leyang Cui, Zhiyang Teng, Nan Duan, Ming Zhou, and Yue Zhang. Logiqa 2.0—an improved dataset for logical reasoning in natural language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2023.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. Logiqa: A challenge dataset for machine reading comprehension with logical reasoning. *arXiv preprint arXiv:2007.08124*, 2020.
- Shayne Longpre, Gregory Yauney, Emily Reif, Katherine Lee, Adam Roberts, Barret Zoph, Denny Zhou, Jason Wei, Kevin Robinson, David Mimno, et al. A pretrainer’s guide to training data: Measuring the effects of data age, domain coverage, quality, & toxicity. *arXiv preprint arXiv:2305.13169*, 2023.
- Ilya Loshchilov and Frank Hutter. Online batch selection for faster training of neural networks. *arXiv preprint arXiv:1511.06343*, 2015.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. *arXiv preprint arXiv:2306.01116*, 2023.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben allal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. The fineweb datasets: Decanting the web for the finest text data at scale, 2024. URL <https://arxiv.org/abs/2406.17557>.
- Qi Qi, Zhishuai Guo, Yi Xu, Rong Jin, and Tianbao Yang. An online method for a class of distributionally robust optimization with non-convex objectives. *Advances in Neural Information Processing Systems*, 34:10067–10080, 2021.
- Qi Qian, Shenghuo Zhu, Jiasheng Tang, Rong Jin, Baigui Sun, and Hao Li. Robust optimization over multiple domains. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 4739–4746, 2019.
- 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.
- Adityanarayanan Radhakrishnan, Mikhail Belkin, and Caroline Uhler. Overparameterized neural networks implement associative memory. *Proceedings of the National Academy of Sciences*, 117(44):27162–27170, 2020.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- Mengye Ren, Wenyan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International conference on machine learning*, pp. 4334–4343. PMLR, 2018.
- Othmane Sebbouh, Robert M Gower, and Aaron Defazio. Almost sure convergence rates for stochastic gradient descent and stochastic heavy ball. In *Conference on Learning Theory*, pp. 3935–3971. PMLR, 2021.

- Daria Soboleva, Faisal Al-Khateeb, Robert Myers, Jacob R Steeves, Joel Hestness, and Nolan Dey. SlimPajama: A 627B token cleaned and deduplicated version of RedPajama. <https://cerebras.ai/blog/slimpajama-a-627b-token-cleaned-and-deduplicated-version-of-redpajama>, 2023. URL <https://huggingface.co/datasets/cerebras/SlimPajama-627B>.
- Daouda Sow, Sen Lin, Zhangyang Wang, and Yingbin Liang. Doubly robust instance-reweighted adversarial training. *arXiv preprint arXiv:2308.00311*, 2023.
- Megh Thakkar, Tolga Bolukbasi, Sriram Ganapathy, Shikhar Vashishth, Sarath Chandar, and Partha Talukdar. Self-influence guided data reweighting for language model pre-training. *arXiv preprint arXiv:2311.00913*, 2023.
- 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.
- A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- Johannes Welbl, Nelson F. Liu, and Matt Gardner. Crowdsourcing multiple choice science questions. In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pp. 94–106, Copenhagen, Denmark, September 2017. Association for Computational Linguistics.
- Alexander Wettig, Aatmik Gupta, Saumya Malik, and Danqi Chen. QuRating: Selecting high-quality data for training language models. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 52915–52971. PMLR, 21–27 Jul 2024. URL <https://proceedings.mlr.press/v235/wettig24a.html>.
- Sang Michael Xie, Hieu Pham, Xuanyi Dong, Nan Du, Hanxiao Liu, Yifeng Lu, Percy Liang, Quoc V Le, Tengyu Ma, and Adams Wei Yu. Doremi: Optimizing data mixtures speeds up language model pretraining. *arXiv preprint arXiv:2305.10429*, 2023.
- Mingyang Yi, Lu Hou, Lifeng Shang, Xin Jiang, Qun Liu, and Zhi-Ming Ma. Reweighting augmented samples by minimizing the maximal expected loss. *arXiv preprint arXiv:2103.08933*, 2021.
- Huimin Zeng, Chen Zhu, Tom Goldstein, and Furong Huang. Are adversarial examples created equal? a learnable weighted minimax risk for robustness under non-uniform attacks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 10815–10823, 2021.
- Jingfeng Zhang, Jianing Zhu, Gang Niu, Bo Han, Masashi Sugiyama, and Mohan Kankanhalli. Geometry-aware instance-reweighted adversarial training. *arXiv preprint arXiv:2010.01736*, 2020.

Appendix

A	Experimental Details	15
A.1	Reference Implementation & Reproducibility	15
A.2	Training & Evaluation Details	16
A.3	Additional Experimental Results	17
A.4	Sensitivity to the value of r	17
B	Toy Regression Example	19
B.1	Toy Regression Problem Setup	19
B.2	Comparisons	20
C	Proofs	21
C.1	Full Gradient	22
C.2	Optimal Weight Strategy	23
C.3	Convex Minibatch SGD	24
C.4	Non-convex Minibatch SGD	27

A EXPERIMENTAL DETAILS

For full reproducibility of our work we provide insights on extending existing training routines with our minimal-invasive loss reweighting scheme.

A.1 REFERENCE IMPLEMENTATION & REPRODUCIBILITY

To integrate our loss reweighting method with a PyTorch training routine, we have to add two plug-and-play components. First, we need to define four **utility functions** to ensure the proper functionality of our approach with reference implementations provided in Listing 1:

1. `apply_strategy(...)`: Implements different reweighting strategies such as focusing on medium losses (quadratic), capping high losses (linupper), or emphasizing extreme values (extremes).
2. `scale_losses(...)`: Applies exponential scaling to adjust the importance of losses based on a temperature parameter r .
3. `normalize_losses(...)`: Normalizes losses to a bounded interval $[-\delta, \delta]$ for consistency in subsequent reweighting.
4. `get_batch_loss_from_logits(...)`: Computes per-sample losses using cross-entropy, ensuring padding tokens do not affect the results.

These four functions are then employed inside the training step. We need to capture the unprocessed raw losses and apply our reweighting scheme. We provide a reference implementation for PyTorch FSDP training in Listing 2 to enable multi-GPU training.

For a full reference, we provide a ready-to-use implementation on GitHub: <https://anonymous.4open.science/r/iclr-loss-reweighting/>.

Listing 1: Utility functions

```

784 1  import torch
785 2
786 3  # Exponential scaling function
787 4  def scale_losses(losses, r):
788 5      return torch.exp(losses / r)
789 6
790 7  # Normalization function for losses
791 8  def normalize_losses(losses, delta=1., l_min=0., l_max=1.):
792 9      return 2. * delta * losses / max(l_max - l_min, 1e-6) - delta * (
793 10         l_max + l_min) / max(l_max - l_min, 1e-6)
794 11
795 12 # Reweighting strategies
796 13 def apply_strategy(losses, delta=1.0, strategy="linupper"):
797 14     if strategy == "linupper":
798 15         return torch.minimum(losses + delta, delta * torch.ones_like(
799 16             losses))
800 17     elif strategy == "uniform":
801 18         # We do not reweight here. This is our baseline.
802 19         return losses
803 20     elif strategy == "quadratic":
804 21         return 1 - losses**2 / delta**2
805 22     elif strategy == "extremes":
806 23         return torch.abs(losses)
807 24     else:
808 25         raise NotImplementedError
809 26
810 27 # Compute batch losses from logits
811 28 def get_batch_loss_from_logits(logits, labels):
812 29     ignore_index = -100
813 30     shift_logits = logits[..., :-1, :].contiguous()
814     shift_labels = labels[..., 1:].contiguous()
815     num_active = (shift_labels != ignore_index).sum(dim=1)
```

```

810 31     loss_fct = torch.nn.CrossEntropyLoss(reduction='none')
811 32     loss = loss_fct(shift_logits.view(-1, logits.size(-1)),
812     shift_labels.view(-1).long())
813 33     return loss.view(logits.size(0), -1).sum(dim=1) / num_active,
814     num_active

```

Listing 2: Integration of reweighting strategies into a multi-GPU training loop.

```

819 1     import torch
820 2     import torch.distributed as dist
821 3
822 4     # Define reweighting function mapping
823 5     STRATEGY = "linupper"
824 6
825 7     # This has to be placed inside a training step.
826 8     ...
827 9
828 10    # Assume model outputs (logits) and labels are already computed
829 11    device_losses, len_norms = get_batch_loss_from_logits(outputs, labels
830 12    )
831 13
832 14    # Initialize placeholder to store the losses from all GPUs
833 15    gathered_losses = torch.zeros(
834 16        dist.get_world_size(),
835 17        len(device_losses),
836 18        device=device_losses.device,
837 19        dtype=device_losses.dtype
838 20    )
839 21
840 22    # Gather losses across all GPUs into tensor
841 23    dist.all_gather_into_tensor(gathered_losses, device_losses.detach())
842 24
843 25    r = r_scheduler(step, cfg.initial_r)
844 26    # Compute sample weights
845 27    with torch.no_grad():
846 28        min_loss = gathered_losses.min().item()
847 29        max_loss = gathered_losses.max().item()
848 30        normalized_losses = normalize_losses(gathered_losses.view(-1),
849 31            delta=1., l_min=min_loss, l_max=max_loss)
850 32        reweighted_losses = apply_strategy(normalized_losses, delta=1.,
851 33            strategy=STRATEGY)
852 34        scaled_losses = scale_losses(reweighted_losses -
853 35            reweighted_losses.max().item(), l=r)
854 36        weights = scaled_losses / scaled_losses.sum()
855 37
856 38        # for instance local_rank can be obtained with dist.get_rank()
857 39        device_weights = weights.view(dist.get_world_size(), -1)[
858 40            local_rank, :]
859 41
860 42    # Reweight losses and scale appropriately
861 43    loss = torch.sum(device_weights * device_losses) * dist.
862 44        get_world_size()
863 45    loss.backward()

```

A.2 TRAINING & EVALUATION DETAILS

In the following, we outline the hyper-parameter configuration and hardware requirements to fully reproduce our experiments.

Table 6: Model Architecture

Architecture	Name	Parameters	Layers	Attention Heads	Embedding Dimensions	Hidden Dimensions	Max. Sequence Length
GPT-2	Mini	124M	12	12	768	3072	512
	Small	210M	24	16	786	3072	512
	Medium	300M	36	24	786	3072	512
LLAMA	1.4B	1400M	24	16	2048	2048	8192

Models. We employ three different *GPT-2 models* to study the effectiveness of our reweighting approach on various model sizes. We report the architectural details in Table 6. Furthermore, to demonstrate the performance of our reweighting approach on state-of-the-art models, we employ a 1.4B parameter model with *Llama* architecture.

Table 7: Training Hyperparameters for our benchmark evaluations

Architecture	Name	Minibatch Size	Learning Rate (LR)	Weight Decay	LR Schedule	LR End	Warmup Steps	Max. Grad Norm	r	Training Steps	Total Documents Seen
GPT-2	Mini	32	0.0005	0.01	Linear Warmup Cosine	0.0001	500	1.0	0.4	20,000	640,000
	Small	48	0.0005	0.01	Linear Warmup Cosine	0.0001	500	1.0	0.4	20,000	960,000
	Medium	48	0.0005	0.01	Linear Warmup Cosine	0.0001	500	1.0	0.4	20,000	960,000
LLAMA	1.4B	128	0.0003	0.01	Linear Warmup Cosine	0.00003	2,000	1.0	1.0	100,000	12,800,000

Datasets. As a dataset for training our GPT-2 models, we employ the SlimPajama-6B dataset, consisting of 7 domain partitions (ArXiv, Book, CC, C4, GitHub, StackExchange, and Wikipedia). The exact number of documents seen is reported in Table 7. We use the FineWeb 15T dataset to train the larger Llama-1.4B and Llama-7B models. For the Llama-1.4B model, we have completed training on 100B randomly sampled tokens. For the 7B model, we are currently (at the moment of submitting this revision of our paper) at 80B tokens seen and report the benchmark results on the latest checkpoints.

Pretraining procedure. We pre-train all models with masked language modeling and do not apply any instruction tuning routine for subsequent evaluations. The GPT-2 models are trained for 20,000 steps in total, and the Llama-1.4B model is trained for 100,000 steps with the hyperparameters provided in Table 7. We use the Huggingface `trainer` interface and the `accelerate` library for distributed training. We provide a ready-to-use execution script in our publicly available code base.

Benchmarks. We use the LM Eval Harness to generate our benchmark results. We evaluate the GPT-2 models every 2000 steps and the Llama model every 5000 steps. For the GPT-2 models, we find that only question-answering tasks work well at this scale, which is why we report 5-shot results on LogiQA, LogiQA-2, SciQ, and PiQA. All other benchmarks yielded random results at this scale, i.e., they were not useful for model evaluation. The reason is the limited capacity of the GPT-2 models compared to larger models like the Llama-1.4B or 7B models. For the Llama architecture models, we employ 13 language understanding and reasoning (LUR) and 6 question-answering (QA) tasks to provide a holistic picture.

A.3 ADDITIONAL EXPERIMENTAL RESULTS

Perplexity. We present additional perplexity visualizations of our GPT2 pretraining runs. Figure 3 shows the results for GPT2-mini and Figure 4 shows the results for GPT2-small.

Benchmark results. The experiments with the 1.4B parameter Llama model show notable performance improvement of 2.1% for language understanding / reasoning Table 8 and question answering tasks ??, respectively. Overall, with our LinUpper strategy, we improve the performance of 10 out of 16 total language understanding and reasoning benchmarks. For question answering, we perform better on 8 out of 10 tasks.

A.4 SENSITIVITY TO THE VALUE OF R

We conduct new experiments using different values of r to understand the sensitivity of our methods to the value of r . The perplexity plots for our method LinUpper with different values of r are provided in fig. 5. Figure 5 shows that when the value of r is large (e.g. $r = 1$) then, as expected, the

Table 8: Llama-1.4B language understanding and reasoning benchmark results. Overall, our LinUpper method improves results for 6 out of 13 tasks with an average performance gain of 0.66%.

Benchmark Name	Uniform	LinUpper (ours)	Difference
ARC Challenge	30.91	31.67	0.76
ARC Easy	61.92	61.32	-0.60
COPA	76.00	76.40	0.40
Lambada (OpenAI)	47.28	45.57	-1.71
Lambada (Standard)	42.29	39.22	-3.07
MMLU	26.42	26.27	-0.15
MNLI	33.00	36.46	3.46
MNLI (mismatch)	32.85	36.08	3.23
RTE	49.10	50.97	1.88
SST-2	56.94	69.01	12.08
TinyARC (norm.)	37.72	39.21	1.49
TinyWinoGrande	54.67	59.36	4.70
WinoGrande	57.42	57.60	0.18
<i>Mean</i>	<i>46.65</i>	<i>48.40</i>	<i>1.74</i>

Table 9: Llama-1.4B language question answering benchmark results. Overall, our LinUpper method improves results for 5 out of 6 tasks with an average performance gain of 1.72%.

Benchmark Name	Uniform	LinUpper (ours)	Difference
BoolQ	61.59	65.47	3.88
LogiQA	20.28	23.66	3.38
LogiQA2	26.02	27.48	1.46
SciQ	91.30	91.90	0.60
SocialQA	43.86	45.19	1.33
TriviaQA (Exact Match)	11.05	10.73	-0.32
<i>Mean</i>	<i>42.35</i>	<i>44.07</i>	<i>1.72</i>

Table 10: Intermediary results of the Llama-7B model on language understanding and reasoning benchmark tasks. Overall, our LinUpper method improves results for 12 out of 13 tasks with an average performance gain of 4.26%.

Benchmark Name	Uniform	LinUpper (ours)	Difference
ARC Challenge	36.77	39.33	2.56
ARC Easy	68.06	70.66	2.61
COPA	81.00	84.00	3.00
Lambada (OpenAI)	56.70	58.43	1.73
Lambada (Standard)	51.39	54.90	3.51
MMLU	26.06	26.78	0.72
MNLI	35.15	37.27	2.12
MNLI (mismatch)	35.74	38.60	2.86
RTE	49.82	57.40	7.58
SST-2	62.04	88.3	26.26
TinyARC	34.46	31.06	-3.40
TinyWinoGrande	46.18	51.24	5.06
WinoGrande	65.19	65.9	0.71
<i>Mean</i>	<i>49.89</i>	<i>54.15</i>	<i>4.26</i>

Table 11: Intermediary results of the Llama-7B model on question answering benchmark tasks. Overall, our LinUpper method improves results for 3 out of 6 tasks with an average performance gain of 0.53%.

Benchmark Name	Uniform	LinUpper (ours)	Difference
BoolQ	67.25	66.82	-0.43
LogiQA	24.12	23.04	-1.08
LogiQA2	26.91	28.44	1.53
SciQ	93.80	94.90	1.10
SocialQA	46.57	46.57	0.00
TriviaQA (Exact Match)	16.55	18.59	2.04
<i>Mean</i>	<i>45.86</i>	<i>46.39</i>	<i>0.53</i>

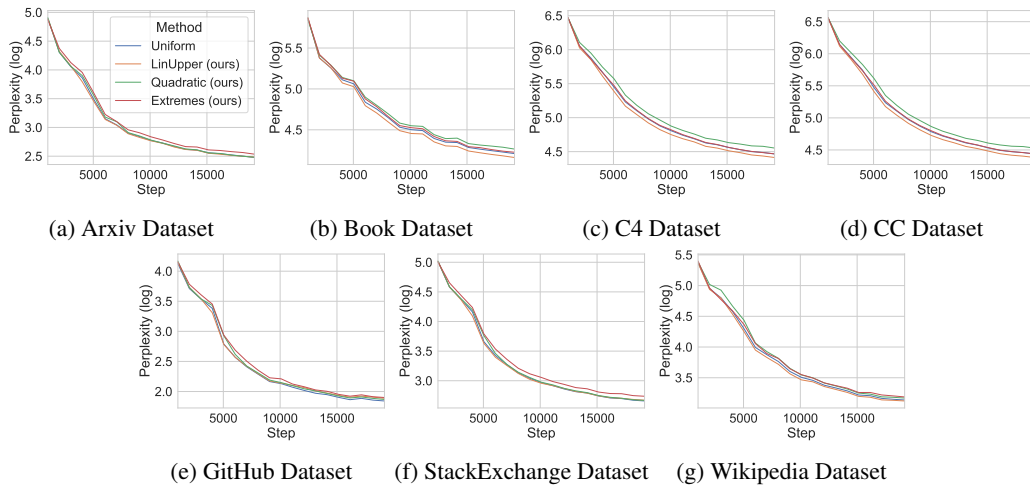


Figure 3: Per-domain perplexities on hold-out validation sets under the uniform domain sampling setting for the GPT2-mini model. Our reweighting strategy **LinUpper** strategy achieves better or at least comparable perplexity on 6 out of 7 domains.

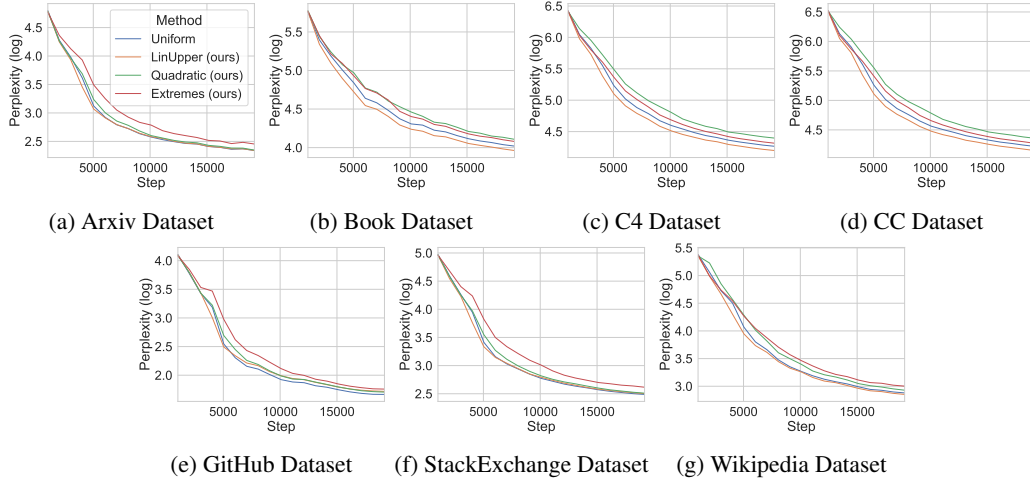


Figure 4: Per-domain perplexities on hold-out validation sets under the uniform domain sampling setting for the GPT2-small model. Our reweighting strategy **LinUpper** strategy achieves better or at least comparable perplexity on 6 out of 7 domains.

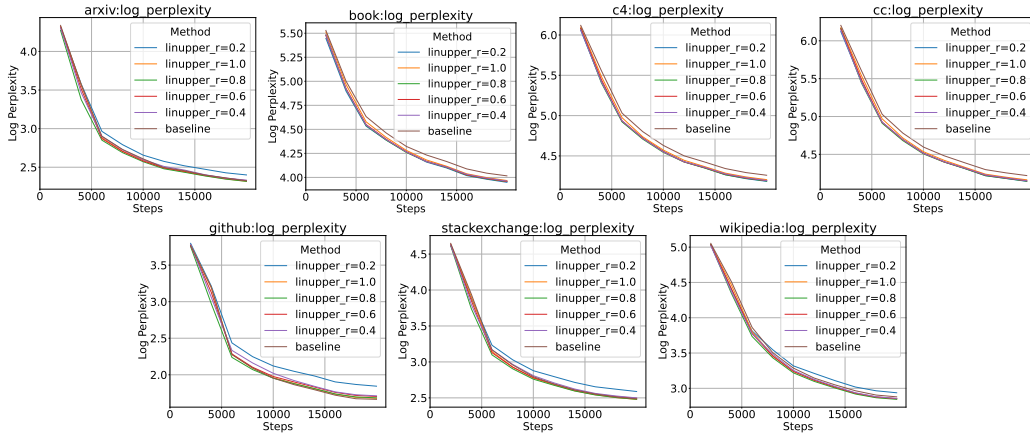


Figure 5: Sensitivity of our method LinUpper to the value of r . When the r is large (e.g. $r = 1$) the performance of our method becomes closer to that of the uniform baseline. Decreasing the value of r leads to diminished effect of low-loss samples, but can also have a negative effect on the performance when r is too small (eg. $r = 0.2$)

performance of our method becomes closer to that of the uniform baseline. Decreasing the value of r leads to diminished effect of low-loss samples, but this can also have a negative effect on the performance when r is too small (eg. $r = 0.2$), potentially due to data wastage by overfiltering the low-loss samples.

B TOY REGRESSION EXAMPLE

B.1 TOY REGRESSION PROBLEM SETUP

To further illustrate the effectiveness of our reweighting strategies on simple ML problems, we apply them to a regression problem where the goal is to learn a linear mapping from input features to output labels, while dealing with noisy data and outliers. We generate a synthetic dataset as follows.

The clean input features $\mathbf{X}_{\text{clean}} \in \mathbb{R}^{n \times p}$ are sampled from the standard normal distribution, where p is the dimensionality of the data, n is the number of clean data points. The target output $\mathbf{y} \in \mathbb{R}^n$ for




Figure 6: **Left:** Test loss vs training steps for the toy regression problem. Our **LinUpper** strategy achieves faster convergence compared to all the other methods. Notably, the **Quadratic** strategy also outperforms the uniform baseline, further highlighting the benefit of selective reweighting. **Right:** Shape of the different reweighting methods used with $r = 1$.

Table 12: Comparison of our **LinUpper** loss reweighting strategy with various DRO-KL configurations on the GPT2-medium experimental setup. On average, our **LinUpper** strategy outperforms the other DRO-KL configurations.

	Uniform	LinUpper ($r = 0.4$; ours)	DRO (KL = 0.4)	DRO (KL = 1.0)	DRO (KL = 2.0)	DRO (KL = 5.0)	DRO (KL = 10.0)
LogiQA	25.7	27.9	23.2	25.5	26.8	25.9	26.3
LogiQA2	27.5	27.6	24.4	28.7	27.5	28.2	27.6
SciQ	49.0	51.8	41.1	48.6	49.4	50.5	49.4
PiQA	55.5	56.2	54.0	56.2	55.7	56.5	55.4
Mean	39.4	40.9	35.7	39.8	39.9	40.3	39.7

the clean data is generated using the following ground truth linear model with added noise:

$$\mathbf{y}_{\text{clean}} = \mathbf{X}_{\text{clean}} \mathbf{W}^* + b^* + c \cdot \epsilon,$$

where $\mathbf{W}^* \in \mathbb{R}^p$ and $b^* \in \mathbb{R}$ are the true model parameters, $\epsilon \sim \mathcal{N}(0, 1)$ represents Gaussian noise, and c is a small constant controlling the noise level. For the m outlier data, we generate the input features $\mathbf{X}_{\text{ood}} \in \mathbb{R}^{m \times p}$ as:

$$\mathbf{X}_{\text{ood}} = 0.1 \cdot \mathcal{N}(0, 1) + 2.0,$$

with the corresponding target outputs \mathbf{y}_{ood} drawn randomly from a standard normal distribution, i.e., $\mathbf{y}_{\text{ood}} \sim \mathcal{N}(0, 1)$.

The full training set consists of both the clean data $(\mathbf{X}_{\text{clean}}, \mathbf{y}_{\text{clean}})$ and the outlier data $(\mathbf{X}_{\text{ood}}, \mathbf{y}_{\text{ood}})$, concatenated together:

$$\mathbf{X}_{\text{all}} = [\mathbf{X}_{\text{clean}}; \mathbf{X}_{\text{ood}}], \quad \mathbf{y}_{\text{all}} = [\mathbf{y}_{\text{clean}}; \mathbf{y}_{\text{ood}}].$$

By incorporating outliers into the training set, we test the robustness of our reweighting strategies, which should down-weight the influence of noisy or irrelevant samples, allowing the model to learn the underlying clean data distribution more effectively.

We define the objective function for this regression task as the mean squared error between the model predictions and the target outputs. Given model parameters $\mathbf{W} \in \mathbb{R}^p$ and $b \in \mathbb{R}$, the loss for a single data point (\mathbf{x}_i, y_i) computed as:

$$\ell_i(\mathbf{W}, b) = \frac{1}{2} (\mathbf{x}_i^\top \mathbf{W} + b - y_i)^2. \quad (12)$$

To evaluate model performance, we compute the mean squared error on a held-out test set $(\mathbf{X}_{\text{test}}, \mathbf{y}_{\text{test}})$, which was generated using the same process as the clean training data.

B.2 COMPARISONS

The results from Figure 6 clearly show that **LinUpper** outperforms all other reweighting strategies and the uniform baseline, achieving the fastest convergence. This supports our theoretical findings that down-weighting low-loss samples allows the model to focus on more informative examples, thus

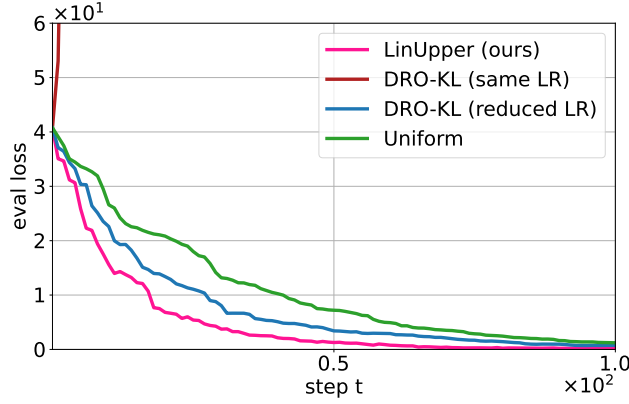


Figure 7: We created a synthetic regression problem (similar to the problem in Appendix B) with 25% of the data being outliers to compare the robustness of the different methods against outliers. The DRO-KL approach because it focuses heavily on the hard samples, it diverges when we use the same learning rate as the other 2 methods (our LinUpper & the Uniform baseline). The DRO-KL method required a smaller learning rate to converge, in which case it converged slower than our method.

accelerating the optimization. Interestingly, the **Quadratic** strategy also performs better than the uniform baseline, indicating that selectively up-weighting the moderate loss samples can improve performance. While **Quadratic** converges more slowly than **LinUpper**, it still highlights that more balanced reweighting can provide benefits over uniform sampling when dealing with noisy data.

C PROOFS

Lemma 1. *Let Assumptions 1 and 2 hold. Then we have*

$$\|\nabla f(\theta) - \nabla f(\theta')\|^2 \leq 2L(f(\theta) - f(\theta') - \langle \nabla f(\theta'), \theta - \theta' \rangle).$$

Proof. Consider the following

$$f(\theta_1) - f(\theta_2) = f(\theta_1) - f(\theta) + f(\theta) - f(\theta_2). \quad (13)$$

From the L -smoothness of f , we have for all $\theta_2, \theta \in \mathbb{R}^d$,

$$f(\theta) - f(\theta_2) \leq \langle \nabla f(\theta_2), \theta - \theta_2 \rangle + \frac{L}{2} \|\theta - \theta_2\|^2. \quad (14)$$

From the the convexity of f , we have for all $\theta_1, \theta \in \mathbb{R}^d$,

$$f(\theta_1) - f(\theta) \leq \langle \nabla f(\theta_1), \theta_1 - \theta \rangle. \quad (15)$$

Substituting Eq. (14), Eq. (15) in Eq. (13), we get

$$f(\theta_1) - f(\theta_2) \leq \langle \nabla f(\theta_1), \theta_1 - \theta \rangle + \langle \nabla f(\theta_2), \theta - \theta_2 \rangle + \frac{L}{2} \|\theta - \theta_2\|^2. \quad (16)$$

Let $\theta = \theta_2 - \frac{1}{L}(\nabla f(\theta_2) - \nabla f(\theta_1))$. Substituting the definition of θ in Eq. (16), we obtain

$$\begin{aligned} f(\theta_1) - f(\theta_2) &\leq \langle \nabla f(\theta_1), \theta_1 - \theta_2 + \frac{1}{L}(\nabla f(\theta_2) - \nabla f(\theta_1)) \rangle - \frac{1}{L} \langle \nabla f(\theta_2), \nabla f(\theta_2) - \nabla f(\theta_1) \rangle \\ &\quad + \frac{1}{2L} \|\nabla f(\theta_2) - \nabla f(\theta_1)\|^2 \\ &= \langle \nabla f(\theta_1), \theta_1 - \theta_2 \rangle - \frac{1}{L} \|\nabla f(\theta_2) - \nabla f(\theta_1)\|^2 + \frac{1}{2L} \|\nabla f(\theta_2) - \nabla f(\theta_1)\|^2 \\ &= \langle \nabla f(\theta_1), \theta_1 - \theta_2 \rangle - \frac{1}{2L} \|\nabla f(\theta_2) - \nabla f(\theta_1)\|^2. \end{aligned} \quad (17)$$

□

From lemma 1, we obtain at the optimum θ^*

$$\|\nabla f(\theta) - \nabla f(\theta^*)\|^2 \leq 2L(f(\theta) - f(\theta^*)).$$

C.1 FULL GRADIENT

We have

$$\begin{aligned} \|\theta^{t+1} - \theta^*\|^2 &= \|\theta^t - \theta^*\|^2 + \|\theta^{t+1} - \theta^t\|^2 + 2\langle \theta^{t+1} - \theta^t, \theta^t - \theta^* \rangle \\ &= \|\theta^t - \theta^*\|^2 + \|\eta_t \sum_{i=1}^M w_{i,t} \nabla f_i(\theta^t)\|^2 - 2\left\langle \eta_t \sum_{i=1}^M w_{i,t} \nabla f_i(\theta^t), \theta^t - \theta^* \right\rangle \\ &\leq \|\theta^t - \theta^*\|^2 + \eta_t^2 M \sum_{i=1}^M w_{i,t}^2 \|\nabla f_i(\theta^t)\|^2 - 2\eta_t \sum_{i=1}^M w_{i,t} \langle \nabla f_i(\theta^t), \theta^t - \theta^* \rangle \quad (18) \end{aligned}$$

Using the convexity of function f_i , we have

$$\langle \nabla f_i(\theta^t), \theta^t - \theta^* \rangle \geq f_i(\theta^t) - f_i(\theta^*). \quad (19)$$

Therefore combining, we obtain

$$\|\theta^{t+1} - \theta^*\|^2 \leq \|\theta^t - \theta^*\|^2 + \eta_t^2 M \sum_{i=1}^M w_{i,t}^2 \|\nabla f_i(\theta^t)\|^2 - 2\eta_t \sum_{i=1}^M w_{i,t} (f_i(\theta^t) - f_i(\theta^*)) \quad (20)$$

Next, we upper-bound the quantity $\|\nabla f_i(\theta^t)\|^2$.

$$\begin{aligned} \|\nabla f_i(\theta^t)\|^2 &\leq 2\|\nabla f_i(\theta^t) - \nabla f_i(\theta^*)\|^2 + 2\|\nabla f_i(\theta^*)\|^2 \\ &\leq 4L(f_i(\theta^t) - f_i(\theta^*)) + 2\sigma_*^2 \quad (21) \end{aligned}$$

where Eq. (21) follows from lemma 1.

Now, combining Eq. (20) and Eq. (21), we obtain

$$\begin{aligned} \|\theta^{t+1} - \theta^*\|^2 &\leq \|\theta^t - \theta^*\|^2 + 4ML\eta_t^2 \sum_{i=1}^M w_{i,t}^2 (f_i(\theta^t) - f_i(\theta^*)) + 2M\eta_t^2 \sigma_*^2 \sum_{i=1}^M w_{i,t}^2 \\ &\quad - 2\eta_t \sum_{i=1}^M w_{i,t} (f_i(\theta^t) - f_i(\theta^*)) \\ &= \|\theta^t - \theta^*\|^2 + \sum_{i=1}^M (4ML\eta_t w_{i,t} - 2)\eta_t w_{i,t} (f_i(\theta^t) - f_i(\theta^*)) \\ &\quad + 2M\eta_t^2 \sigma_*^2 \sum_{i=1}^M w_{i,t}^2, \quad (23) \end{aligned}$$

Selecting η_t such that $4ML\eta_t w_{i,t} - 2 \leq -1 \ \forall i$, i.e., $\eta_t \leq \frac{1}{4MLw_{\max,t}}$, yields

$$\|\theta^{t+1} - \theta^*\|^2 \leq \|\theta^t - \theta^*\|^2 - \eta_t \sum_{i=1}^M w_{i,t} (f_i(\theta^t) - f_i(\theta^*)) + 2M\eta_t^2 \sigma_*^2, \quad (24)$$

where we also used the fact that $\sum_{i=1}^M w_{i,t}^2 \leq 1$ due to $\sum_{i=1}^M w_{i,t} = 1$ and $w_{i,t} \geq 0$. Hence,

$$\begin{aligned} \|\theta^{t+1} - \theta^*\|^2 &\leq \|\theta^t - \theta^*\|^2 - \frac{\eta_t}{M} \sum_{i=1}^M (f_i(\theta^t) - f_i(\theta^*)) + \eta_t \sum_{i=1}^M \left(\frac{1}{M} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^*)) + 2M\eta_t^2 \sigma_*^2 \end{aligned}$$

$$= \|\theta^t - \theta^*\|^2 - \eta_t (f(\theta^t) - f(\theta^*)) + \eta_t \sum_{i=1}^M \left(\frac{1}{M} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^*)) + 2M\eta_t^2 \sigma_*^2. \quad (25)$$

Telescoping Eq. (25) from $t = 0$ to $T - 1$ yields

$$\begin{aligned} \|\theta^T - \theta^*\|^2 &\leq \|\theta^0 - \theta^*\|^2 - \sum_{t=0}^{T-1} \eta_t (f(\theta^t) - f(\theta^*)) + \sum_{t=0}^{T-1} \eta_t \sum_{i=1}^M \left(\frac{1}{M} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^*)) \\ &\quad + 2M\sigma_*^2 \sum_{t=0}^{T-1} \eta_t^2. \end{aligned} \quad (26)$$

For a weighting scheme such that $w_{\max,t} \leq \frac{2}{M}$ so that η_t can be fixed to $\eta = \frac{1}{8L}$, we obtain

$$\eta \sum_{t=0}^{T-1} (f(\theta^t) - f(\theta^*)) \leq \|\theta^0 - \theta^*\|^2 + \eta \sum_{t=0}^{T-1} \sum_{i=1}^M \left(\frac{1}{M} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^*)) + 2M\sigma_*^2 \eta^2 T.$$

$$f(\bar{\theta}^T) - f(\theta^*) \leq \frac{8L\|\theta^0 - \theta^*\|^2}{T} + \frac{1}{T} \sum_{t=0}^{T-1} \sum_{i=1}^M \left(\frac{1}{M} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^*)) + \frac{M\sigma_*^2}{4L},$$

where $\bar{\theta}^T = \frac{1}{T} \sum_{t=0}^{T-1} \theta^t$. Define $\delta^t = \sum_{i=1}^M \left(\frac{1}{M} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^*))$. Hence, we obtain

$$f(\bar{\theta}^T) - f(\theta^*) \leq \frac{8L\|\theta^0 - \theta^*\|^2}{T} + \frac{1}{T} \sum_{t=0}^{T-1} \delta^t + \frac{M\sigma_*^2}{4L}. \quad (27)$$

C.2 OPTIMAL WEIGHT STRATEGY

Proposition 1. *The optimal strategy that minimizes the KL-divergence regularized δ^t , i.e., $\delta^t + r \sum_{i=1}^M w_i \log(Mw_i)$, is given by*

$$w_i = C \min \left\{ \exp \left(\frac{h_i}{r} \right), \frac{2}{M} \right\}, \quad (5)$$

where C is a normalizing constant that ensures $\sum_i w_i = 1$.

Proof. For the sake of clarity, we drop the index t in δ^t in the following proof.

Let $\Delta_i = f(\mathbf{x}_i; \theta^t) - f(\mathbf{x}_i; \theta^*)$ be the loss gap. After dropping all terms that do not depend on the weights w_i 's, we have the following constrained optimization problem

$$\min_{w_1, \dots, w_M} \quad - \sum_{i=1}^M w_i \Delta_i + r \sum_{i=1}^M w_i \log w_i \quad (28)$$

$$\text{subject to} \quad w_i \leq \frac{2}{M}, i \in [M], \quad (29)$$

$$\sum_{i=1}^M w_i = 1. \quad (30)$$

As we have inequality constraints in the above problem, we introduce the slack variable t and consider the Lagrangian function to solve the above problem as follows:

$$L(w_1, \dots, w_M, \nu, \mu) = - \sum_{i=1}^M w_i \Delta_i + r \sum_{i=1}^M w_i \log w_i - \nu \left(\sum_{i=1}^M w_i - 1 \right) - \mu \left(w_i - \frac{2}{M} + t^2 \right), \quad (31)$$

and we have the gradients

$$\frac{\partial L}{\partial w_i} = -\Delta_i + r + r \log w_i - \nu - \mu$$

$$\begin{aligned}\frac{\partial L}{\partial \lambda} &= - \left(\sum_{i=1}^M w_i - 1 \right) \\ \frac{\partial L}{\partial \mu} &= - \left(w_i - \frac{2}{M} + t^2 \right).\end{aligned}\quad (32)$$

We consider the complementary slackness condition. We have one slack variable t , and the corresponding Lagrange multiplier is ν . We now consider whether a slack variable is zero (which the corresponding inequality constraint is active) or the Lagrange multiplier is zero (the constraint is inactive), we obtain

- $\mu = 0, t^2 > 0$: Using $\frac{\partial L}{\partial \nu} = 0$, we get $\sum_{i=1}^M w_i = 1$. Using $\frac{\partial L}{\partial w_i} = 0, \mu = 0$, we get $-\Delta_i + r + r \log w_i - \nu = 0$. Hence, we have $w_i = \exp\left(\frac{\nu + \Delta_i - r}{r}\right)$. Therefore, we obtain $\nu = -r \ln\left(\sum_{i=1}^M \exp\left(\frac{\Delta_i - r}{r}\right)\right)$ yielding $w_i = \exp\left(\frac{-r \ln\left(\sum_{i=1}^M \exp\left(\frac{\Delta_i - r}{r}\right)\right) + \Delta_i - r}{r}\right) = C \exp\left(\frac{\Delta_i}{r}\right)$ where $C = \exp\left(-\ln\left(\sum_{i=1}^M \exp\left(\frac{\Delta_i - r}{r}\right)\right) - 1\right)$.
- $\mu \neq 0, t^2 = 0$: Using $\frac{\partial L}{\partial \mu} = 0, t = 0$, we get $w_i = \frac{2}{M}$ and $\mu + \nu = -\Delta_i + r + r \log \frac{2}{M}$.

Hence, we obtain $w_i = \min\left\{C \exp\left(\frac{\Delta_i}{r}\right), \frac{2}{M}\right\}$. \square

C.3 CONVEX MINIBATCH SGD

Theorem 2. Let Assumptions 1, 2, and 3 hold. Consider a minibatch of size $|\mathcal{B}| = b$ with reweighting scheme satisfying $\max_{i \in \mathcal{B}} w_{i,t} \leq 2/b$. Then,

- **minibatch SGD:** for $\eta = \frac{1}{8L\sqrt{T}}$, we have

$$\mathbb{E}[f(\bar{\theta}^T) - f(\theta^*)] \leq \frac{8L\|\theta^0 - \theta^*\|^2}{\sqrt{T}} + \frac{1}{T} \sum_{t=0}^{T-1} \delta_t, \quad (10)$$

- **minibatch SGD with momentum:** for $\eta = \frac{1}{8L\sqrt{T+1}}, \lambda_t > 0$, we have

$$\mathbb{E}[f(\theta^T) - f(\theta^*)] \leq \frac{8L\|\theta^0 - \theta^*\|^2}{\sqrt{T+1}} + \frac{2}{T+1} \sum_{t=0}^{T-1} \delta_t + \frac{2}{T+1} \sum_{t=0}^{T-1} \lambda_t \mu_t, \quad (11)$$

where $\delta_t = \mathbb{E}[\sum_{i \in \mathcal{B}} (\frac{1}{b} - w_{i,t}) (f_i(\theta^t) - f_i(\theta^*))]$, $\mu_t = \mathbb{E}[\sum_{i \in \mathcal{B}} (\frac{1}{b} - w_{i,t}) (f_i(\theta^t) - f_i(\theta^{t-1}))]$, and $\bar{\theta}^T = \frac{1}{T} \sum_{t=0}^{T-1} \theta^t$.

Proof. (i) **Minibatch SGD:** For the minibatch SGD based on Eq. (2), we have

$$\begin{aligned}\|\theta^{t+1} - \theta^*\|^2 &= \|\theta^t - \theta^*\|^2 + \|\theta^{t+1} - \theta^t\|^2 + 2\langle \theta^{t+1} - \theta^t, \theta^t - \theta^* \rangle \\ &= \|\theta^t - \theta^*\|^2 + \left\| \eta_t \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t) \right\|^2 - 2 \left\langle \eta_t \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t), \theta^t - \theta^* \right\rangle.\end{aligned}\quad (33)$$

Applying expectation on both sides, we get

$$\begin{aligned}\mathbb{E}\|\theta^{t+1} - \theta^*\|^2 &\leq \mathbb{E}\|\theta^t - \theta^*\|^2 + \eta_t^2 \mathbb{E}\left\| \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t) \right\|^2 - 2\eta_t \mathbb{E}\left\langle \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t), \theta^t - \theta^* \right\rangle \\ &\stackrel{(a)}{\leq} \mathbb{E}\|\theta^t - \theta^*\|^2 + \eta_t^2 \mathbb{E}\left\| \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t) \right\|^2 - 2\eta_t \mathbb{E} \sum_{i \in \mathcal{B}} w_{i,t} (f_i(\theta^t) - f_i(\theta^*))\end{aligned}$$

$$\begin{aligned}
&\leq \|\theta^t - \theta^*\|^2 + \eta_t^2 b \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t}^2 \|\nabla f_i(\theta^t)\|^2 \right] - 2\eta_t \mathbb{E} \sum_{i \in \mathcal{B}} w_{i,t} (f_i(\theta^t) - f_i(\theta^*)) \\
&\stackrel{(b)}{\leq} \|\theta^t - \theta^*\|^2 + 4L\eta_t^2 b \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t}^2 (f_i(\theta^t) - f_i(\theta^*)) \right] - 2\eta_t \mathbb{E} \sum_{i \in \mathcal{B}} w_{i,t} (f_i(\theta^t) - f_i(\theta^*)) \\
&\leq \|\theta^t - \theta^*\|^2 + \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t} (4L\eta_t^2 b w_{i,t} - 2\eta_t) (f_i(\theta^t) - f_i(\theta^*)) \right], \tag{34}
\end{aligned}$$

where (a) follows from Assumption 1, and (b) follows from lemma 1. Selecting η_t such that $4bL\eta_t w_{i,t} - 2 \leq -1 \ \forall i$, i.e., $\eta_t \leq \frac{1}{4bLw_{\max,t}}$, yields

$$\begin{aligned}
\mathbb{E} \|\theta^{t+1} - \theta^*\|^2 &\leq \|\theta^t - \theta^*\|^2 - \eta_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t} (f_i(\theta^t) - f_i(\theta^*)) \right] \\
&= \|\theta^t - \theta^*\|^2 - \eta_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} \frac{1}{b} (f_i(\theta^t) - f_i(\theta^*)) \right] \\
&\quad + \eta_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^*)) \right] \\
&= \|\theta^t - \theta^*\|^2 - \eta_t (f(\theta^t) - f(\theta^*)) + \eta_t \delta_t, \tag{35}
\end{aligned}$$

where $\delta_t = \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^*)) \right]$ and $\sum_{i \in \mathcal{S}} w_{i,t}^2 \leq 1$. Hence, telescoping Eq. (35) from $t = 0$ to $T-1$ yields

$$\mathbb{E} \|\theta^T - \theta^*\|^2 \leq \|\theta^0 - \theta^*\|^2 - \sum_{t=0}^{T-1} \eta_t (f(\theta^t) - f(\theta^*)) + \sum_{t=0}^{T-1} \eta_t \delta_t. \tag{36}$$

For a weighting scheme such that $w_{\max,t} \leq \frac{2}{b}$ so that η_t can be fixed to $\eta = \frac{1}{8L\sqrt{T}}$, we obtain

$$f(\bar{\theta}^T) - f(\theta^*) \leq \frac{8L\|\theta^0 - \theta^*\|^2}{\sqrt{T}} + \frac{1}{T} \sum_{t=0}^{T-1} \delta_t, \tag{37}$$

where $\bar{\theta}^T = \frac{1}{T} \sum_{t=0}^{T-1} \theta^t$.

The traditional convergence bound of minibatch SGD for convex functions can be recovered when $\delta_t = 0$ which occurs when the samples in the batch have uniform weights, $w_{i,t} = \frac{1}{b} \ \forall i \in \mathcal{B}$. We obtain a looser convergence bound when $\delta_t \geq 0$ resulting from the assignment of larger weights to smaller loss gaps. Moreover, down-weighting the importance of samples with low loss gaps leads to a better bound ($\delta_t \leq 0$). Note that the bound in Eq. (10) only holds for reweighting schemes that satisfy $w_{i,t} \leq 2/b$. Thus, enforcing an upper bound on the maximum weight that can be assigned to a single data point. Typically, the minibatch gradient with uniform weights is unbiased, $\mathbb{E} \left[\sum_{i \in \mathcal{B}} \frac{1}{b} \nabla f_i(\theta^t) \right] = \nabla f(\theta^t)$. However, in our setting, the importance weights are functions of the sample functions. Hence, $\mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t) \right] \neq \nabla f(\theta^t)$ leading to the analysis of Theorem 2 being more involved. Based on our analysis, we note that the step size η_t is a function of the maximum value of the weights, $w_{\max,t} = \frac{2}{b}$ which leads to the choice of $\eta = \frac{1}{8L\sqrt{T}}$ for convergence.

(ii) **Minibatch SGD with Momentum:** Given the loss-based sample reweighting strategies, Eq. (11) characterizes the effects of loss-based reweighting on the convergence bound of reweighted minibatch SGD with momentum when the sample losses f_i are convex. Given the loss-based sample reweighting strategies, we incorporate heavy ball momentum method to improve the convergence result and to characterize the effects of loss-based reweighting on the convergence bound. Hence, we have the following update equations:

$$z_{t+1} = z_t - \eta_t \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta_t), \tag{38}$$

$$\theta_{t+1} = \frac{\lambda_{t+1}}{1 + \lambda_{t+1}} \theta_t + \frac{1}{1 + \lambda_{t+1}} z_{t+1}, \quad (39)$$

where η_t is the step size and $\eta_t, \lambda_t > 0$. The above iterates θ_t are equal to the iterates of the following stochastic minibatch heavy ball momentum method Sebbouh et al. (2021).

$$\theta_{t+1} = \theta_t - \frac{\eta_t}{1 + \lambda_{t+1}} \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta_t) + \frac{\lambda_t}{1 + \lambda_{t+1}} (\theta_t - \theta_{t-1}).$$

Eq. (11) characterizes the effects of loss-based reweighting on the convergence bound of reweighted minibatch SGD with momentum when the sample losses f_i are convex.

For the minibatch SGD with momentum based on Eq. (39), we have

$$\begin{aligned} \|z^{t+1} - \theta^*\|^2 &= \|z^t - \theta^*\|^2 + \|z^{t+1} - z^t\|^2 + 2 \langle z^{t+1} - z^t, z^t - \theta^* \rangle \\ &= \|z^t - \theta^*\|^2 + \left\| \eta_t \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t) \right\|^2 - 2 \left\langle \eta_t \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t), z^t - \theta^* \right\rangle. \end{aligned} \quad (40)$$

Applying expectation on both sides, we get

$$\begin{aligned} &\mathbb{E} \|z^{t+1} - \theta^*\|^2 \\ &\leq \|z^t - \theta^*\|^2 + \eta_t^2 \mathbb{E} \left\| \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t) \right\|^2 - 2\eta_t \mathbb{E} \left[\left\langle \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t), \theta^t - \theta^* \right\rangle \right] \\ &\quad - 2\eta_t \lambda_t \mathbb{E} \left[\left\langle \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t), \theta^t - \theta^{t-1} \right\rangle \right] \\ &\stackrel{(a)}{\leq} \|z^t - \theta^*\|^2 + \eta_t^2 \mathbb{E} \left\| \sum_{i \in \mathcal{B}} w_{i,t} \nabla f_i(\theta^t) \right\|^2 - 2\eta_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t} (f_i(\theta^t) - f_i(\theta^*)) \right] \\ &\quad - 2\eta_t \lambda_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t} (f_i(\theta^t) - f_i(\theta^{t-1})) \right] \\ &\stackrel{(b)}{\leq} \|z^t - \theta^*\|^2 + \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t} (4L\eta_t^2 b w_{i,t} - 2\eta_t) (f_i(\theta^t) - f_i(\theta^*)) \right] \\ &\quad - 2\eta_t \lambda_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t} (f_i(\theta^t) - f_i(\theta^{t-1})) \right] \\ &= \|z^t - \theta^*\|^2 + \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t} (4L\eta_t^2 b w_{i,t} - 2\eta_t - 2\eta_t \lambda_t) (f_i(\theta^t) - f_i(\theta^*)) \right] \\ &\quad + 2\eta_t \lambda_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t} (f_i(\theta^{t-1}) - f_i(\theta^*)) \right] \\ &\stackrel{(c)}{\leq} \|z^t - \theta^*\|^2 - 2\eta_t \lambda_{t+1} \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t} (f_i(\theta^t) - f_i(\theta^*)) \right] + 2\eta_t \lambda_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} w_{i,t} (f_i(\theta^{t-1}) - f_i(\theta^*)) \right] \\ &= \|z^t - \theta^*\|^2 - 2\eta_t \lambda_{t+1} \mathbb{E} \left[\sum_{i \in \mathcal{B}} \frac{1}{b} (f_i(\theta^t) - f_i(\theta^*)) \right] + 2\eta_t \lambda_{t+1} \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^*)) \right] \\ &\quad + 2\eta_t \lambda_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} \frac{1}{b} (f_i(\theta^{t-1}) - f_i(\theta^*)) \right] + 2\eta_t \lambda_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(w_{i,t} - \frac{1}{b} \right) (f_i(\theta^{t-1}) - f_i(\theta^*)) \right] \\ &= \|z^t - \theta^*\|^2 - 2\eta_t \lambda_{t+1} \mathbb{E} (f(\theta^t) - f(\theta^*)) + 2\eta_t \lambda_{t+1} \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^*)) \right] \end{aligned}$$

$$\begin{aligned}
& + 2\eta_t \lambda_t \mathbb{E} (f(\theta^{t-1}) - f(\theta^*)) + 2\eta_t \lambda_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(w_{i,t} - \frac{1}{b} \right) (f_i(\theta^{t-1}) - f_i(\theta^*)) \right] \\
& = \|z^t - \theta^*\|^2 - 2\eta_t \lambda_{t+1} \mathbb{E} (f(\theta^t) - f(\theta^*)) + \eta_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^*)) \right] \\
& + 2\eta_t \lambda_t \mathbb{E} (f(\theta^{t-1}) - f(\theta^*)) + 2\eta_t \lambda_t \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{i,t} \right) (f_i(\theta^t) - f_i(\theta^{t-1})) \right] \tag{41}
\end{aligned}$$

where (a) follows from applying Assumption 1 to the third and fourth terms, (b) follows from applying Lemma 1 to the second term, and (c) follows from setting η_t as $\eta_t \leq \frac{1}{4Lbw_{\max,t}}$ and $\lambda_{t+1} = \lambda_t + \frac{1}{2}$. Let $\delta_t = \mathbb{E} [\sum_{i \in \mathcal{B}} (\frac{1}{b} - w_{i,t}) (f_i(\theta^t) - f_i(\theta^*))]$, $\mu_t = \mathbb{E} [\sum_{i \in \mathcal{B}} (\frac{1}{b} - w_{i,t}) (f_i(\theta^t) - f_i(\theta^{t-1}))]$ and $\sum_{i \in \mathcal{B}} w_{i,t}^2 \leq 1$. Hence, telescoping Eq. (41) from $t = 0$ to $T - 1$ and noting that $z_0 = x_0$ yields

$$\mathbb{E} \|z^T - \theta^*\|^2 \leq \|\theta^0 - \theta^*\|^2 - 2\eta \lambda_{T+1} \mathbb{E} (f(\theta^T) - f(\theta^*)) + \sum_{t=0}^{T-1} 2\eta \delta_t + \sum_{t=0}^{T-1} 2\eta \lambda_t \mu_t. \tag{42}$$

With $\lambda_t = t/2$, for a weighting scheme such that $w_{\max,t} \leq \frac{2}{b}$ so that η_t can be fixed to $\eta = \frac{1}{8L\sqrt{T+1}}$, we obtain

$$\mathbb{E} [f(\theta^T) - f(\theta^*)] \leq \frac{8L\|\theta^0 - \theta^*\|^2}{\sqrt{T+1}} + \frac{2}{T+1} \sum_{t=0}^{T-1} \delta_t + \frac{2}{T+1} \sum_{t=0}^{T-1} \lambda_t \mu_t. \tag{43}$$

□

C.4 NON-CONVEX MINIBATCH SGD

Non-convex: Unlike the convex case, the importance weight $w_{i,t}$ is set based on the gradient norm instead of the loss gap $f_i(\theta^t) - f_i(\theta^*)$ when the sample functions f_i are non-convex. Theorem 3 underlines the influence of these importance weights on the convergence bound for the non-convex case. Note that computing the gradient norm requires extra computations, whereas the loss values are readily available in practice.

Assumption 4. *The norm of the difference between the sample gradient $\nabla f_i(\theta_t)$ and full gradient $\nabla f(\theta_t)$ is bounded by $\|\nabla f_i(\theta_t) - \nabla f(\theta_t)\|^2 \leq \mathcal{V}^2, \forall \theta_t, i \in [M]$.*

Theorem 3. *[Non-convex Minibatch SGD] Given that the objective functions $f_i, i \in [n]$ satisfy Assumption 2 and Assumption 4. Consider a minibatch of data points $|\mathcal{B}| = b$ with reweighting scheme $w_{i,t}$, if $i \in \mathcal{B}$ and 0 otherwise where $\max_{i \in \mathcal{B}} w_{i,t} \leq 2/b$. Then, for $\eta = \frac{1}{8L\sqrt{T}}$, we have*

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\theta^t)\|^2 \leq \frac{16L(f(\theta^0) - f^*)}{\sqrt{T}} + \frac{8}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 \right] + \frac{b\mathcal{V}^2}{4\sqrt{T}}. \tag{44}$$

Proof. Applying L -smoothness condition

$$\begin{aligned}
& f(\theta^{t+1}) \\
& \leq f(\theta^t) + \langle \nabla f(\theta^t), \theta^{t+1} - \theta^t \rangle + \frac{L}{2} \|\theta^{t+1} - \theta^t\|^2 \\
& = f(\theta^t) - \eta \left\langle \nabla f(\theta^t), \sum_{i \in \mathcal{B}} w_{it} \nabla f_i(\theta^t) \right\rangle + \frac{L\eta^2}{2} \left\| \sum_{i \in \mathcal{B}} w_{it} \nabla f_i(\theta^t) \right\|^2 \\
& = f(\theta^t) - \eta \left\langle \nabla f(\theta^t), \sum_{i \in \mathcal{B}} \frac{1}{b} \nabla f_i(\theta^t) \right\rangle + \eta \left\langle \nabla f(\theta^t), \sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \nabla f_i(\theta^t) \right\rangle + \frac{L\eta^2}{2} \left\| \sum_{i \in \mathcal{B}} w_{it} \nabla f_i(\theta^t) \right\|^2 \\
& = f(\theta^t) - \eta \left\langle \nabla f(\theta^t), \sum_{i \in \mathcal{B}} \frac{1}{b} \nabla f_i(\theta^t) \right\rangle + \eta \sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \langle \nabla f(\theta^t), \nabla f_i(\theta^t) \rangle + \frac{L\eta^2}{2} \left\| \sum_{i \in \mathcal{B}} w_{it} \nabla f_i(\theta^t) \right\|^2
\end{aligned}$$

$$\begin{aligned}
&\leq f(\theta^t) - \eta \left\langle \nabla f(\theta^t), \sum_{i \in \mathcal{B}} \frac{1}{b} \nabla f_i(\theta^t) \right\rangle + 2\eta \sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f(\theta^t)\|^2 + 2\eta \sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 \\
&+ \frac{L\eta^2}{2} \left\| \sum_{i \in \mathcal{B}} w_{it} \nabla f_i(\theta^t) \right\|^2 \\
&= f(\theta^t) - \eta \left\langle \nabla f(\theta^t), \sum_{i \in \mathcal{B}} \frac{1}{b} \nabla f_i(\theta^t) \right\rangle + 2\eta \sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 + \frac{L\eta^2}{2} \left\| \sum_{i \in \mathcal{B}} w_{it} \nabla f_i(\theta^t) \right\|^2 \\
&\leq f(\theta^t) - \eta \left\langle \nabla f(\theta^t), \sum_{i \in \mathcal{B}} \frac{1}{b} \nabla f_i(\theta^t) \right\rangle + 2\eta \sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 + \frac{Lb\eta^2}{2} \sum_{i \in \mathcal{B}} w_{it}^2 \|\nabla f_i(\theta^t)\|^2 \\
&\leq f(\theta^t) - \eta \left\langle \nabla f(\theta^t), \sum_{i \in \mathcal{B}} \frac{1}{b} \nabla f_i(\theta^t) \right\rangle + 2\eta \sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 \\
&+ Lb\eta^2 \sum_{i \in \mathcal{B}} w_{it}^2 \|\nabla f_i(\theta^t) - \nabla f(\theta^t)\|^2 + Lb\eta^2 \sum_{i \in \mathcal{B}} w_{it}^2 \|\nabla f(\theta^t)\|^2 \\
&\leq f(\theta^t) - \eta \left\langle \nabla f(\theta^t), \sum_{i \in \mathcal{B}} \frac{1}{b} \nabla f_i(\theta^t) \right\rangle + 2\eta \sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 \\
&+ Lb\eta^2 \mathcal{V}^2 + Lb\eta^2 \sum_{i \in \mathcal{B}} w_{i,t}^2 \|\nabla f(\theta^t)\|^2
\end{aligned} \tag{45}$$

where we have used the bound $\|\nabla f_i(\theta^t) - \nabla f(\theta^t)\|^2 \leq \mathcal{V}^2$ and $\sum_{i \in \mathcal{B}} w_{i,t}^2 \leq 1$. Applying expectation on both sides, we get

$$\begin{aligned}
\mathbb{E}f(\theta^{t+1}) &\leq f(\theta^t) - \eta \mathbb{E} \|\nabla f(\theta^t)\|^2 + (2\eta) \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 \right] \\
&+ Lb\eta^2 \mathcal{V}^2 + Lb\eta^2 \sum_{i \in \mathcal{B}} w_{it}^2 \mathbb{E} \|\nabla f(\theta^t)\|^2 \\
&\leq f(\theta^t) - \eta \mathbb{E} \|\nabla f(\theta^t)\|^2 + (2\eta) \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 \right] \\
&+ Lb\eta^2 \mathcal{V}^2 + Lb^2\eta^2 w_{\max,t}^2 \mathbb{E} \|\nabla f(\theta^t)\|^2 \\
&\leq f(\theta^t) - \frac{\eta}{2} \mathbb{E} \|\nabla f(\theta^t)\|^2 + 2\eta \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 \right] + Lb\eta^2 \mathcal{V}^2
\end{aligned} \tag{46}$$

where $\eta \leq \frac{1}{2Lb^2w_{\max,t}^2}$ in the last inequality. Telescoping and rearranging the terms yields

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\theta^t)\|^2 \leq \frac{2(f(\theta^0) - f^*)}{\eta T} + \frac{4}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 \right] + 2Lb\eta \mathcal{V}^2. \tag{47}$$

For a weighting scheme such that $w_{\max,t} \leq \frac{2}{b}$ so that η_t can be fixed to $\eta = \frac{1}{8L\sqrt{T}}$, we get

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\theta^t)\|^2 \leq \frac{16L(f(\theta^0) - f^*)}{\sqrt{T}} + \frac{4}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 \right] + \frac{b\mathcal{V}^2}{4\sqrt{T}}. \tag{48}$$

□

From Theorem 3, we can obtain the convergence bound of traditional minibatch SGD for non-convex functions when $w_{i,t} = \frac{1}{b} \forall i \in \mathcal{B}$. Moreover, a looser convergence bound is obtained when $\mathbb{E} \left[\sum_{i \in \mathcal{B}} \left(\frac{1}{b} - w_{it} \right) \|\nabla f_i(\theta^t)\|^2 \right] \geq 0$, for $w_{i,t} \leq \frac{1}{b}$.