

Low-rank Momentum Factorization for Memory Efficient Training

Pouria Mahdavinia

*Department of Computer Science and Engineering
The Pennsylvania State University*

prm5426@psu.edu

Mehrdad Mahdavi

*Department of Computer Science and Engineering
The Pennsylvania State University*

mzm616@psu.edu

Reviewed on OpenReview: <https://openreview.net/forum?id=W3D3TVo9a3>

Abstract

Fine-tuning large foundation models presents significant memory challenges due to stateful optimizers like AdamW, often requiring several times more GPU memory than inference. While memory-efficient methods like parameter-efficient fine-tuning (e.g., LoRA) and optimizer state compression exist, recent approaches like GaLore bridge these by using low-rank gradient projections and subspace moment accumulation. However, such methods may struggle with fixed subspaces or computationally costly offline resampling (e.g., requiring full-matrix SVDs). We propose Momentum Factorized SGD (MoFaSGD), which maintains a dynamically updated low-rank SVD representation of the first-order momentum, closely approximating its full-rank counterpart throughout training. This factorization enables a memory-efficient fine-tuning method that adaptively updates the optimization subspace at each iteration. Crucially, MoFaSGD leverages the computed low-rank momentum factors to perform efficient spectrally normalized updates, offering an alternative to subspace moment accumulation. We establish theoretical convergence guarantees for MoFaSGD, proving it achieves an optimal rate for non-convex stochastic optimization under standard assumptions. Empirically, we demonstrate MoFaSGD’s effectiveness on large language model alignment benchmarks, achieving a competitive trade-off between memory reduction (comparable to LoRA) and performance compared to state-of-the-art low-rank optimization methods. Our implementation is available at <https://github.com/pmahdavi/MoFaSGD>.

1 Introduction

Advancements in AI have been propelled by initial scaling laws, which boost pre-training performance via ever-larger models and datasets (Kaplan et al., 2020). However, adapting these large foundation models to downstream tasks, such as instruction or preference tuning (Ouyang et al., 2022), often requires several times more GPU memory than inference. This memory burden largely stems from storing optimizer states (e.g., momentum terms) required by ubiquitous methods like AdamW. Various strategies have emerged to alleviate these costs.

Parameter-Efficient Fine-Tuning (PEFT) methods provide an architectural solution, introducing a small set of new, trainable parameters called adapters. For example, the popular LoRA technique (Hu et al., 2021) restricts model updates to these low-rank adapters, drastically reducing memory overhead by training only a small fraction of parameters. Variants such as DoRA (Liu et al., 2024), AdaLoRA (Zhang et al., 2023), VeRA (Kopiczko et al., 2023) and rsLoRA (Kalajdzievski, 2023) have further refined this approach to low-rank parameterization. Other approaches focus on compressing optimizer states directly; methods like AdaFactor (Shazeer & Stern, 2018), SM3 (Anil et al., 2019), and quantization techniques (Modoranu

et al., 2024; Feinberg et al., 2024) factorize or reduce the precision of moments. Alternatively, stateless methods like signSGD (Bernstein et al., 2018) and SWAN (Ma et al., 2024) avoid momentum accumulation by carefully normalizing gradients each iteration.

A complementary line of research, **low-rank subspace optimization**, projects gradients onto a smaller, dynamic subspace to reduce memory. This allows full-parameter updates while operating in a lower-dimensional space. Key examples include GaLore (Zhao et al., 2024a), Flora (Hao et al., 2024), and ReLoRA (Lialin et al., 2023). However, applying these dynamic subspace techniques effectively presents difficulties.

Challenges in Online Subspace Updates. While promising, dynamically updating the optimization subspace online, as in GaLore, faces key obstacles. First, managing optimizer states (particularly momentum) during abrupt subspace changes is non-trivial; existing methods may reset states (Lialin et al., 2023), transform them (Hao et al., 2024), or leave them unchanged (Hao et al., 2024). Second, determining the new subspace, often based on gradients, can incur high computational costs (e.g., full SVD on large gradient matrices (Zhao et al., 2024a)), hindering scalability. These limitations highlight the need for more efficient, iteration-level subspace adaptation strategies.

Contributions

Motivated by the success of gradient-based techniques for memory-efficient fine-tuning, the advantages of dynamic projection subspaces over static ones, and the potential for online subspace updates to better track full-rank optimization dynamics, we address three research questions:

- (Q1) Can we mitigate the challenges of online subspace resampling through a computationally efficient update strategy while rigorously tracking projection residuals?
- (Q2) Can we directly estimate full-rank adaptive optimizer trajectories using low-rank factors, instead of relying on compressed accumulation of subspace moments?
- (Q3) Can we match the theoretical per-iteration complexity of standard adaptive optimizers while maintaining LoRA-level memory savings?

We answer these questions affirmatively by introducing MoFaSGD, which maintains a dynamically updated low-rank SVD factorization of the first-order momentum. We define this factorization as:

$$\hat{\mathbf{M}}_t = \mathbf{U}_{t+1} \text{Diag}(\boldsymbol{\sigma}_{t+1}) \mathbf{V}_{t+1}^\top \approx \beta \hat{\mathbf{M}}_{t-1} + (1 - \beta) \mathbf{G}_t \quad (\text{First-Momentum Factorization}) \quad (1)$$

Here, $\hat{\mathbf{M}}_t$ denotes the approximate first-order momentum at iteration t , \mathbf{G}_t is the gradient at iteration t , and β is the momentum decay factor. The low-rank SVD factors $\mathbf{U}_{t+1} \in \mathbb{R}^{m \times r}$ (left singular vectors), $\boldsymbol{\sigma}_{t+1} \in \mathbb{R}^r$ (singular values), and $\mathbf{V}_{t+1} \in \mathbb{R}^{n \times r}$ (right singular vectors) serve as the optimizer state in our algorithm.

Motivated by our observation that the exponential moving average (EMA) of gradients, defined as $\sum_{i=1}^t \beta^{t-i} \mathbf{G}_i$, often exhibit low-rank structure, MoFaSGD applies a specific low-rank approximation strategy: maintaining an online, low-rank SVD factorization of the *first-order momentum*. This differs from state compression methods that typically target second moments (Shazeer & Stern, 2018) and avoids the complex factored approximations of curvature used in methods like Shampoo or KFAC (Gupta et al., 2018; Martens & Grosse, 2015). The factorization is updated efficiently at each iteration using tangent space projections, bypassing costly offline resampling (Zhao et al., 2024a). Crucially, MoFaSGD leverages the *same* continuously updated low-rank momentum factors ($\mathbf{U}_{t+1}, \mathbf{V}_{t+1}$) to perform spectral normalization ($\mathbf{W}_{t+1} \leftarrow \mathbf{W}_t - \eta \mathbf{U}_{t+1} \mathbf{V}_{t+1}^\top$). This integrated approach, inspired by methods like Shampoo (Gupta et al., 2018) and Muon (Jordan et al., 2024b), provides adaptive step directions without requiring separate, computationally intensive matrix operations (e.g., SVD, matrix roots, or Newton-Schulz iterations) at each step.

MoFaSGD achieves LoRA-like memory savings through a unified low-rank momentum representation, while also enabling efficient, adaptive full-parameter updates via spectral normalization. This positions our method as a low-rank variant of Muon. Unlike MoFaSGD, Muon maintains full-rank momentum buffers ($\mathcal{O}(mn)$ memory) and applies full-rank updates after normalization.

2 Related Work

To improve memory efficiency for training large-scale neural networks, various strategies have emerged, including parameter-efficient fine-tuning, subspace optimization, optimizer state compression, stateless approaches, and partial update techniques. Moreover, second-order and preconditioning methods such as Shampoo (Gupta et al., 2018) have been gaining popularity due to their faster convergence compared to AdamW, albeit with even higher memory requirements, as studied in detail in Kasimbeg et al. (2025).

Below, we detail relevant methods like subspace optimization, optimizer state compression, stateless approaches, and gradient spectral normalization; other approaches are deferred to Appendix A.

Subspace Optimization

Rather than introducing additional parameters through PEFT adapters, subspace optimization methods focus on reducing the memory footprint of gradients and optimizer states directly by projecting gradients onto a low-rank subspace and performing moment accumulation in the subspace as well (Zhao et al., 2024a; Gur-Ari et al., 2018; Gressmann et al., 2020; Yang et al., 2023; Vogels et al., 2020). GaLore (Zhao et al., 2024a) projects gradients onto a subspace defined by either the left or right singular vectors of the gradients, computed via SVD, and accumulates optimizer states within this subspace, thereby reducing the required optimizer memory. Flora (Hao et al., 2024) periodically resamples the subspace projection matrix using a multivariate Gaussian distribution for low-rank gradient projections.

AdaRankGrad (Refael et al., 2024) dynamically adjusts gradient rank during training. LDAdam (Robert et al., 2024) and APOLLO (Zhu et al., 2024) maintain optimizer states in low-dimensional representations, adopting similar strategy to subspace moment accumulation as in GaLore (Zhao et al., 2024a). Subspace management is a key challenge. Computing it can require costly operations like full SVD (Zhao et al., 2024a), and its update frequency creates a trade-off: infrequent updates can lead to stale information, while frequent updates can disrupt optimizer state accumulation. Managing optimizer states across subspace changes also requires careful strategies (Lialin et al., 2023; Hao et al., 2024).

Optimizer State Compression. This approach directly reduces the memory footprint of existing optimizer states, primarily the second moments used in adaptive methods. AdaFactor (Shazeer & Stern, 2018) and SM3 (Anil et al., 2019) achieve this through factorization techniques, reducing memory complexity from $\mathcal{O}(mn)$ to $\mathcal{O}(m+n)$. Quantization methods like 8-bit Adam (Dettmers et al., 2021) and 4-bit variants (Li et al., 2023) reduce the precision of stored moments, achieving near full-precision performance with significant memory savings. CAME (Luo et al., 2023) and Adapprox (Zhao et al., 2024b) refine these compression techniques for potentially better accuracy.

Stateless and Gradient Normalization Methods. Eliminating optimizer states entirely provides maximal memory efficiency. signSGD (Bernstein et al., 2018) achieves this by using only the sign of the gradient. Lion (Chen et al., 2024) similarly reduces the memory footprint of AdamW by discarding the second-moment estimate. Gradient normalization methods such as SWAN (Ma et al., 2024) and Muon (Jordan et al., 2024b) emulate the behavior of adaptive optimizers without storing second-moment statistics. Muon, in particular, applies momentum followed by approximate orthogonalization of the momentum matrix ($\mathbf{U}_M \mathbf{V}_M^\top$, where $\mathbf{M} = \mathbf{U}_M \mathbf{\Sigma}_M \mathbf{V}_M^\top$) using efficient Newton-Schulz iterations (Jordan et al., 2024b), drawing connections to accumulation-free variants of Shampoo (Jordan et al., 2024b; Bernstein & Newhouse, 2024b). However, Muon still retains the Nesterov momentum, and thus is not considered fully stateless.

Positioning MoFaSGD. The landscape of memory-efficient optimization presents a spectrum of trade-offs. Architectural-based solutions such as PEFT methods offer high memory savings but restrict fine-tuning updates (Hu et al., 2021) to a low-rank space. Subspace methods enable full-rank parameter updates but encounter challenges such as costly or potentially disruptive subspace resampling and moment accumulation (Zhao et al., 2024a; Hao et al., 2024); moreover, subspace techniques are mostly applied to AdamW, while their application to non-diagonal-based preconditioning methods (such as Shampoo and its variants (Gupta et al., 2018; George et al., 2018; Jordan et al., 2024b)) remains largely unexplored. Optimizer state compression techniques focus on reducing the storage cost of second moments (Shazeer & Stern, 2018; Dettmers

et al., 2021). Stateless methods are memory-efficient but discard historical information (Bernstein et al., 2018; Ma et al., 2024). This work introduces MoFaSGD to navigate these challenges.

3 Preliminaries

This section introduces notation and provides background on subspace and adaptive optimization methods.

Notation and Conventions. We denote matrices with bold capital letters (e.g., \mathbf{X}), vectors with bold lowercase letters (e.g., \mathbf{x}), and scalars with non-bold letters (e.g., x). The Frobenius inner product for matrices is denoted as $\langle \cdot, \cdot \rangle$ (e.g., $\langle \mathbf{X}, \mathbf{Y} \rangle = \text{Tr}(\mathbf{X}^\top \mathbf{Y})$). Norms are denoted by $\| \cdot \|$, with specific norms indicated using subscripts (e.g., $\| \cdot \|_F$ for the Frobenius norm). The loss function is represented by $\mathcal{L}(\cdot)$. Without loss of generality, we denote the model parameters as $\mathbf{W}_t \in \mathbb{R}^{m \times n}$, and the full-rank gradient at \mathbf{W}_t is denoted as $\mathbf{G}_t = \frac{\partial \mathcal{L}(\mathbf{W})}{\partial \mathbf{W}}|_{\mathbf{W}_t} \in \mathbb{R}^{m \times n}$. The subscript index t denotes the iteration step of any optimizer for each weight matrix. Gradients are defined based on the specific batches used in the corresponding iteration. To simplify notation, we use the same symbols as previously defined where applicable. For any matrix $\mathbf{X} \in \mathbb{R}^{m_1 \times m_2}$, the vectorization operator $\text{Vec}(\cdot)$ stacks its columns into a single vector, denoted as $\mathbf{x} = \text{Vec}(\mathbf{X}) \in \mathbb{R}^{m_1 m_2 \times 1}$. The vectorized version of a matrix is always denoted by the corresponding bold lowercase letter. The Kronecker product is denoted as \otimes , such that for $\mathbf{X} \in \mathbb{R}^{m_1 \times n_1}$ and $\mathbf{Y} \in \mathbb{R}^{m_2 \times n_2}$, the Kronecker product $\mathbf{X} \otimes \mathbf{Y} \in \mathbb{R}^{m_1 m_2 \times n_1 n_2}$. The symbol \odot denotes the element-wise (Hadamard) product between two matrices of the same dimensions. See Appendix D.2 for further notational details.

Adaptive methods with Switched-off Momentums

Without first-order momentum, adaptive methods can be framed as a form of online mirror descent:

$$\mathbf{w}_{t+1} = \arg \min_{\mathbf{w}} \mathcal{L}(\mathbf{w}_t) + \langle \mathbf{w} - \mathbf{w}_t, \mathbf{g}_t \rangle + \frac{1}{2\eta} \|\mathbf{w} - \mathbf{w}_t\|_{\mathbf{P}_t} = \mathbf{w}_t - \eta \mathbf{P}_t^{-1} \mathbf{g}_t \quad (2)$$

where the preconditioner \mathbf{P}_t can be any positive semidefinite matrix, and $\mathcal{L}(\mathbf{w}_t) + \langle \mathbf{w}_{t+1} - \mathbf{w}_t, \mathbf{g}_t \rangle$ is the local approximation of the loss function. For instance, if we let $\mathbf{P}_{t, \text{AdaGrad}} = \left[\sum_{i=1}^t \mathbf{g}_i \mathbf{g}_i^\top \right]^{\frac{1}{2}}$, we exactly recover full-matrix Adagrad (Duchi et al., 2011).

Similarly, if we define $\mathbf{L}_t = \sum_{i=1}^t \mathbf{G}_i \mathbf{G}_i^\top$ and $\mathbf{R}_t = \sum_{i=1}^t \mathbf{G}_i^\top \mathbf{G}_i$, then letting $\mathbf{P}_{t, \text{Shampoo}} = (\mathbf{R}_t \otimes \mathbf{L}_t)^{\frac{1}{4}}$ recovers the Shampoo updates for matrices (Gupta et al., 2018). Furthermore, if we consider that $\mathbf{g}_t \mathbf{g}_t^\top$ closely approximates the Gauss-Newton components of the true Hessian (Morwani et al., 2024), this motivates switching off covariance momentum (second moment) and considering the following preconditioners based on Adagrad and Shampoo: $\mathbf{P}_{t,1} = [\mathbf{g}_t \mathbf{g}_t^\top]^{\frac{1}{2}}$ and $\mathbf{P}_{t,2} = (\mathbf{G}_t^\top \mathbf{G}_t \otimes \mathbf{G}_t \mathbf{G}_t^\top)^{\frac{1}{4}}$.

A simple derivation shows that using the diagonal version of $\mathbf{P}_{t,1}$ in Equation 2 recovers Signed-SGD (Bernstein et al., 2018). Moreover, if we consider the reduced SVD of the gradient $\mathbf{G}_t = \mathbf{U}_{\mathbf{G}_t} \Sigma_{\mathbf{G}_t} \mathbf{V}_{\mathbf{G}_t}^\top$, then we can write $\mathbf{P}_{t,2} = (\mathbf{V}_{\mathbf{G}_t} \Sigma_{\mathbf{G}_t}^2 \mathbf{V}_{\mathbf{G}_t}^\top \otimes \mathbf{U}_{\mathbf{G}_t} \Sigma_{\mathbf{G}_t}^2 \mathbf{U}_{\mathbf{G}_t}^\top)^{\frac{1}{4}}$. Using Lemma D.1 in the appendix, we have:

$$\mathbf{P}_{t,2} = (\mathbf{V}_{\mathbf{G}_t} \Sigma_{\mathbf{G}_t}^2 \mathbf{V}_{\mathbf{G}_t}^\top \otimes \mathbf{U}_{\mathbf{G}_t} \Sigma_{\mathbf{G}_t}^2 \mathbf{U}_{\mathbf{G}_t}^\top)^{\frac{1}{4}} = (\mathbf{V}_{\mathbf{G}_t} \otimes \mathbf{U}_{\mathbf{G}_t}) (\Sigma_{\mathbf{G}_t}^{\frac{1}{2}} \otimes \Sigma_{\mathbf{G}_t}^{\frac{1}{2}}) (\mathbf{V}_{\mathbf{G}_t} \otimes \mathbf{U}_{\mathbf{G}_t})^\top \quad (3)$$

Furthermore, we can write $\mathbf{g}_t = \text{Vec}(\mathbf{G}_t) = (\mathbf{V}_{\mathbf{G}_t} \otimes \mathbf{U}_{\mathbf{G}_t}) \text{Vec}(\Sigma_{\mathbf{G}_t})$. Substituting the preconditioner $\mathbf{P}_{t,2}$ and the vectorized gradient into Equation 2 yields:

$$\begin{aligned} \mathbf{w}_{t+1} &= \mathbf{w}_t - \eta (\mathbf{V}_{\mathbf{G}_t} \otimes \mathbf{U}_{\mathbf{G}_t}) (\Sigma_{\mathbf{G}_t}^{-\frac{1}{2}} \otimes \Sigma_{\mathbf{G}_t}^{-\frac{1}{2}}) (\mathbf{V}_{\mathbf{G}_t} \otimes \mathbf{U}_{\mathbf{G}_t})^\top \mathbf{g}_t \\ &= \mathbf{w}_t - \eta (\mathbf{V}_{\mathbf{G}_t} \otimes \mathbf{U}_{\mathbf{G}_t}) (\Sigma_{\mathbf{G}_t}^{-\frac{1}{2}} \otimes \Sigma_{\mathbf{G}_t}^{-\frac{1}{2}}) \text{Vec}(\Sigma_{\mathbf{G}_t}) \\ &= \mathbf{w}_t - \eta (\mathbf{V}_{\mathbf{G}_t} \otimes \mathbf{U}_{\mathbf{G}_t}) \text{Vec}(\mathbf{I}_r) = \mathbf{w}_t - \eta \mathbf{U}_{\mathbf{G}_t} \mathbf{V}_{\mathbf{G}_t}^\top \end{aligned} \quad (4)$$

Thus, using $\mathbf{P}_{t,2}$ recovers spectrally normalized updates, sometimes referred to as gradient whitening (Jordan et al., 2024b; Ma et al., 2024). For a similar derivation, see Bernstein & Newhouse (2024b). Additional background on the historical development of adaptive optimization methods is provided in Appendix B.

Subspace Optimization Methods

Many studies support the conjecture that gradients during the training of deep neural networks exhibit a low-rank structure, lying in a low-dimensional subspace (Gur-Ari et al., 2018; Gressmann et al., 2020; Yang et al., 2023). This property has been studied both theoretically and empirically and has been leveraged to improve optimization algorithms ranging from communication-efficient distributed training (Vogels et al., 2020) to efficient fine-tuning (Hu et al., 2021; Lialin et al., 2023; Hao et al., 2024; Zhao et al., 2024a). Zhao et al. (2024a) aim to explicitly leverage this property to perform subspace training, where gradients are projected and accumulated in a low-rank subspace, leading to significant memory footprint reduction.

We briefly clarify subspace methods mathematically, following the formulation of GaLore (Zhao et al., 2024a). Let $\mathbf{W}_t \in \mathbb{R}^{m \times n}$ represent the model parameters, and let $\mathbf{Q}_t \in \mathbb{R}^{m \times r}$ be the projection matrix defining the subspace at iteration t . Then, GaLore performs the following update:

$$\begin{aligned}
 \mathbf{G}_{t,r} &= \mathbf{Q}_t^\top \mathbf{G}_t \in \mathbb{R}^{r \times n} && \text{(Subspace projection)} \\
 \mathbf{M}_{t,r} &= \beta_1 \mathbf{M}_{t-1,r} + (1 - \beta_1) \mathbf{G}_{t,r} && \text{(First subspace moment)} \\
 \mathbf{V}_{t,r} &= \beta_2 \mathbf{V}_{t-1,r} + (1 - \beta_2) (\mathbf{G}_{t,r} \odot \mathbf{G}_{t,r}) && \text{(Second subspace moment)} \\
 \mathbf{W}_{t+1} &= \mathbf{W}_t - \eta_t \mathbf{Q}_t \left(\frac{\mathbf{M}_{t,r}}{\sqrt{\mathbf{V}_{t,r}}} \right) && \text{(Project back and update)}
 \end{aligned} \tag{5}$$

Moreover, by switching off subspace momentum accumulations, the GaLore update can simply be seen as projected gradient descent: $\mathbf{W}_{t+1} = \mathbf{W}_t - \eta_t \mathbf{Q}_t \mathbf{Q}_t^\top \mathbf{G}_t$. In addition to subspace accumulation, GaLore performs offline updates of the subspace \mathbf{Q}_t by updating it as $\mathbf{U}_{\mathbf{G}_t}^{1:r}$, the top- r left singular vectors of \mathbf{G}_t , obtained via an SVD operation at predetermined intervals.

4 Momentum Factorized SGD with Spectral Normalization

This section introduces MoFaSGD, our memory-efficient optimization method, and its theoretical underpinnings. We begin in Subsection 4.1 by detailing the core algorithmic ideas behind MoFaSGD. This includes the motivation derived from the low-rank structure observed in optimizers, the process of maintaining and efficiently updating low-rank momentum factors using tangent space projection, and the subsequent use of these factors for spectrally normalized parameter updates (Algorithm 1). Following the algorithmic description, Subsection 4.2 presents the convergence properties and theoretical analysis of MoFaSGD, justifying our design choices and establishing formal performance guarantees.

4.1 The MoFaSGD Algorithm

Motivation: Low-Rank Momentum Factors

Feinberg et al. (2024) show that the EMA of gradient covariance $\sum_{i=1}^t \beta^{t-i} \mathbf{G}_i \mathbf{G}_i^\top$ and $\sum_{i=1}^t \beta^{t-i} \mathbf{G}_i^\top \mathbf{G}_i$ maintain low-rank structure and spectral decay properties throughout training, and Zhao et al. (2024a) argue that the gradients themselves become low-rank during fine-tuning. Building on these observations, we conjecture that the gradient EMAs exhibit low-rank properties. We experimentally evaluate this conjecture in Section 5.3 and show that the mass of the gradient EMA is largely centered on its top few singular values. Low-rank structure has also been widely leveraged for LLM fine-tuning, e.g., LoRA (Hu et al., 2021) adapts low-rank matrices during training, while GaLore (Zhao et al., 2024a) leverages the low-rank structure of gradients. Thus, leveraging the low-rank property of the gradient EMA connects the implicit assumptions underlying GaLore and LoRA.

Building on this conjecture, we propose to maintain a low-rank SVD factorized representation of the first momentum. We highlight that this factorization plays two major roles in our algorithm design. First, it is leveraged for online subspace sampling, and second, it provides a low-rank estimation of the first momentum (gradient EMA). Formally, we define the low-rank moment factors as:

$$\hat{\mathbf{M}}_t \triangleq \mathbf{U}_{t+1} \Sigma_{t+1} \mathbf{V}_{t+1}^\top \approx \sum_{i=1}^t \beta^{t-i} \mathbf{G}_i \tag{Low-rank Moment Factors}$$

Algorithm 1 MoFaSGD: Momentum Factorized Stochastic Gradient Descent**Require:** Step size η , decay β , rank r **Ensure:** Optimized weights \mathbf{W}_t

```

1: Initialize:  $\mathbf{W}_0$ 
2:  $\mathbf{G}_0 \leftarrow \nabla_{\mathbf{W}} \mathcal{L}(\mathbf{W}_0)$ 
3: Initialize moment factors:  $(\mathbf{U}_0, \mathbf{\Sigma}_0, \mathbf{V}_0) \leftarrow \text{SVD}_r(\mathbf{G}_0)$ 
4:  $t \leftarrow 0$ 
5: repeat
6:    $\mathbf{G}_t \leftarrow \nabla_{\mathbf{W}} \mathcal{L}(\mathbf{W}_t)$ 
7:    $(\mathbf{U}_{t+1}, \mathbf{\Sigma}_{t+1}, \mathbf{V}_{t+1}) \leftarrow \text{UMF}(\mathbf{G}_t, \mathbf{U}_t, \mathbf{\Sigma}_t, \mathbf{V}_t, \beta)$ 
8:    $\mathbf{W}_{t+1} \leftarrow \mathbf{W}_t - \eta \mathbf{U}_{t+1} \mathbf{V}_{t+1}^\top$ 
9:    $t \leftarrow t + 1$ 
10: until convergence criterion is met
11: return  $\mathbf{W}_t$ 

```

function **UMF**($\mathbf{G}_t, \mathbf{U}_t, \mathbf{\Sigma}_t, \mathbf{V}_t, \beta$)

```

1: Compute subspace projections:
2:    $\mathbf{G}_t \mathbf{V}_t, \mathbf{U}_t^\top \mathbf{G}_t, \mathbf{U}_t^\top \mathbf{G}_t \mathbf{V}_t$ 
3: Compute QR factors:
4:    $(\mathbf{U}'_t, \mathbf{R}_{\mathbf{U}_t}) = \text{QR}([\mathbf{U}_t \quad \mathbf{G}_t \mathbf{V}_t])$ 
5:    $(\mathbf{V}'_t, \mathbf{R}_{\mathbf{V}_t}) = \text{QR}([\mathbf{V}_t \quad \mathbf{G}_t^\top \mathbf{U}_t])$ 
6: Construct  $2r \times 2r$  matrix:
7:    $\mathbf{S}_t = \mathbf{R}_{\mathbf{U}_t} \begin{bmatrix} \beta \mathbf{\Sigma}_t - \mathbf{U}_t^\top \mathbf{G}_t \mathbf{V}_t & \mathbf{I}_r \\ \mathbf{I}_r & \mathbf{0}_r \end{bmatrix} \mathbf{R}_{\mathbf{V}_t}^\top$ 
8: Compute rank- $r$  SVD:
9:    $\mathbf{U}''_t \mathbf{\Sigma}''_t (\mathbf{V}''_t)^\top \leftarrow \text{SVD}_r(\mathbf{S}_t)$ 
10:  $\mathbf{U}_{t+1} \leftarrow \mathbf{U}'_t \mathbf{U}''_t$ 
11:  $\mathbf{V}_{t+1} \leftarrow \mathbf{V}'_t \mathbf{V}''_t$ 
12:  $\mathbf{\Sigma}_{t+1} \leftarrow \mathbf{\Sigma}''_t$ 
13: return  $(\mathbf{U}_{t+1}, \mathbf{\Sigma}_{t+1}, \mathbf{V}_{t+1})$ 

```

where $\{\mathbf{G}_i\}_{i=1}^t$ are the observed gradients until iteration t . The left moment factor $\mathbf{U}_{t+1} \in \mathbb{R}^{m \times r}$ and the right moment factor $\mathbf{V}_{t+1} \in \mathbb{R}^{n \times r}$ are orthogonal matrices, while $\mathbf{\Sigma}_{t+1} \in \mathbb{R}^{r \times r}$ is diagonal.

Our method maps the full-rank gradient \mathbf{G}_t to the tangent space of the previous moment factor representation $(\mathbf{U}_t, \mathbf{\Sigma}_t, \mathbf{V}_t)$, to ensure a smooth adaptation of the subspace when a new gradient arrives. Formally, we leverage the tangent space of the previous iteration $\mathcal{T}_{(\mathbf{U}_t, \mathbf{\Sigma}_t, \mathbf{V}_t)}$ as the new subspace for projection, which we denote as \mathcal{T}_t :

$$\mathcal{T}_t = \{\mathbf{U}_t \mathbf{M} \mathbf{V}_t^\top + \mathbf{U}_p \mathbf{V}_t^\top + \mathbf{U}_t \mathbf{V}_p^\top \mid \mathbf{M} \in \mathbb{R}^{r \times r}, \mathbf{U}_p \in \mathbb{R}^{m \times r}, \mathbf{V}_p \in \mathbb{R}^{n \times r}, \mathbf{U}_t^\top \mathbf{U}_p = \mathbf{0}, \mathbf{V}_t^\top \mathbf{V}_p = \mathbf{0}\} \quad (6)$$

and the projection to this subspace can be derived as follows:

$$\hat{\mathbf{G}}_t \triangleq \text{Proj}_{\mathcal{T}_t}(\mathbf{G}_t) = \mathbf{U}_t \mathbf{U}_t^\top \mathbf{G}_t + \mathbf{G}_t \mathbf{V}_t \mathbf{V}_t^\top - \mathbf{U}_t \mathbf{U}_t^\top \mathbf{G}_t \mathbf{V}_t \mathbf{V}_t^\top \quad (\text{Online subspace projection})$$

where the definition of projection to the tangent space is $\text{Proj}_{\mathcal{T}_t}(\mathbf{G}_t) = \arg \min_{\mathbf{G} \in \mathcal{T}_t} \|\mathbf{G} - \mathbf{G}_t\|_F$.

Projecting onto the tangent subspace of previous gradients as shown in Theorem 4.3, results in a lower compression error compared to the left, right, or two-sided subspace projections used in GaLore (Zhao et al., 2024a).

Efficient Momentum Factor Updates

The second component of our approach is to efficiently approximate the current moment factor $\hat{\mathbf{M}}_t$ given the projected gradient $\hat{\mathbf{G}}_t$ and the previous moment factor $\hat{\mathbf{M}}_{t-1}$. A naive update, $\hat{\mathbf{M}}_t = \text{SVD}_r(\hat{\mathbf{G}}_t + \beta \hat{\mathbf{M}}_{t-1})$ where $\text{SVD}_r(\cdot)$ denotes the rank- r truncated SVD, involves a computationally expensive SVD operation, which we aim to avoid. Since both $\hat{\mathbf{M}}_{t-1}$ and $\hat{\mathbf{G}}_t$ are rank- r , their sum has a rank of at most $2r$. This observation allows us to approximate $\hat{\mathbf{M}}_t$ in $\mathcal{O}((m+n)r^2)$, which is far more efficient than a full-matrix SVD. Let $(\mathbf{U}'_t, \mathbf{R}_{\mathbf{U}_t}) = \text{QR}([\mathbf{U}_t \quad \mathbf{G}_t \mathbf{V}_t])$, and $(\mathbf{V}'_t, \mathbf{R}_{\mathbf{V}_t}) = \text{QR}([\mathbf{V}_t \quad \mathbf{G}_t^\top \mathbf{U}_t])$, where QR stands for the QR decomposition, and $\mathbf{U}'_t \in \mathbb{R}^{m \times 2r}$, $\mathbf{V}'_t \in \mathbb{R}^{n \times 2r}$ and $\mathbf{R}_{\mathbf{U}_t}, \mathbf{R}_{\mathbf{V}_t} \in \mathbb{R}^{2r \times 2r}$. Then we can write:

$$\begin{aligned} \hat{\mathbf{G}}_t + \beta \hat{\mathbf{M}}_{t-1} &= \mathbf{U}_t \mathbf{U}_t^\top \mathbf{G}_t + \mathbf{G}_t \mathbf{V}_t \mathbf{V}_t^\top + \mathbf{U}_t (\beta \mathbf{\Sigma}_t - \mathbf{U}_t^\top \mathbf{G}_t \mathbf{V}_t) \mathbf{V}_t^\top \\ &= [\mathbf{U}_t \quad \mathbf{G}_t \mathbf{V}_t] \begin{bmatrix} \beta \mathbf{\Sigma}_t - \mathbf{U}_t^\top \mathbf{G}_t \mathbf{V}_t & \mathbf{I}_r \\ \mathbf{I}_r & \mathbf{0}_r \end{bmatrix} \begin{bmatrix} \mathbf{V}_t^\top \\ \mathbf{U}_t^\top \mathbf{G}_t \end{bmatrix} = \mathbf{U}'_t (\mathbf{R}_{\mathbf{U}_t} \begin{bmatrix} \beta \mathbf{\Sigma}_t - \mathbf{U}_t^\top \mathbf{G}_t \mathbf{V}_t & \mathbf{I}_r \\ \mathbf{I}_r & \mathbf{0}_r \end{bmatrix} \mathbf{R}_{\mathbf{V}_t}^\top) \mathbf{V}_t'^\top \end{aligned} \quad (7)$$

Let $\mathbf{U}''_t \mathbf{\Sigma}''_t (\mathbf{V}''_t)^\top = \text{SVD}_r(\mathbf{R}_{\mathbf{U}_t} \begin{bmatrix} \beta \mathbf{\Sigma}_t - \mathbf{U}_t^\top \mathbf{G}_t \mathbf{V}_t & \mathbf{I}_r \\ \mathbf{I}_r & \mathbf{0}_r \end{bmatrix} \mathbf{R}_{\mathbf{V}_t}^\top)$, and note that the inner matrix has rank at most r , and $\mathbf{U}''_t, \mathbf{V}''_t \in \mathbb{R}^{2r \times r}$. We can finally write our momentum factor update rule as:

$$\mathbf{U}_{t+1} = \mathbf{U}'_t \mathbf{U}''_t \quad , \quad \mathbf{V}_{t+1} = \mathbf{V}'_t \mathbf{V}''_t \quad , \quad \mathbf{\Sigma}_{t+1} = \mathbf{\Sigma}''_t \quad (8)$$

The computation complexity of our approach includes two QR decompositions, $\mathcal{O}((m+n)r^2)$, one full SVD on a $2r \times 2r$ matrix, $\mathcal{O}(r^3)$, and hence the total complexity is $\mathcal{O}((m+n)r^2 + r^3)$.

From Momentum Factors to Spectrally Normalized Updates

Inspired by the connection between spectrally normalized gradient updates and effective non-diagonal preconditioning methods like Shampoo (Gupta et al., 2018), and motivated by the strong empirical performance of Muon (Jordan et al., 2024b), which applies spectral normalization to gradient momentum, we leverage our low-rank momentum factorization (Equation Low-rank Moment Factors) for the main optimizer step. The MoFaSGD update rule is:

$$\mathbf{W}_{t+1} = \mathbf{W}_t - \eta \mathbf{U}_{t+1} \mathbf{V}_{t+1}^\top \quad (9)$$

Here, $\mathbf{U}_{t+1} \in \mathbb{R}^{m \times r}$ and $\mathbf{V}_{t+1} \in \mathbb{R}^{n \times r}$ represent the left and right singular vectors derived from the efficiently computed low-rank approximation of the first-order momentum, $\hat{\mathbf{M}}_t$. MoFaSGD (summarized in Algorithm 1) contrasts with Muon, which operates on the full-rank momentum \mathbf{M}_t (requiring $\mathcal{O}(mn)$ memory) and uses Newton-Schulz iterations to approximate $\mathbf{U}_{\mathbf{M}_t} \mathbf{V}_{\mathbf{M}_t}^\top$ for its update step. MoFaSGD can thus be viewed as a memory-efficient, low-rank variant of Muon.

Furthermore, we highlight the key distinctions between MoFaSGD and GaLore (Zhao et al., 2024a). GaLore employs a two-stage process: 1) Projecting gradients onto a low-rank subspace defined by the singular vectors of the *gradient* itself, updated periodically (referred to as offline subspace resampling), and 2) Accumulating first and second moments (akin to Adam) within this low-rank subspace.

MoFaSGD adopts different strategies for both stages. Firstly, regarding the subspace definition, MoFaSGD performs gradient low-rank projection, but crucially, onto the *tangent space* defined by the singular vectors of the *gradient momentum* ($\hat{\mathbf{M}}_t$). This online, per-iteration subspace adaptation contrasts with GaLore’s offline resampling based on single gradients. The choice of the tangent space projection is theoretically motivated by its optimality in minimizing projection residuals (Theorem 4.3), while using the momentum’s singular vectors aims for a more stable subspace that evolves smoothly, mitigating potential noise in individual gradients.

Secondly, concerning the optimizer update, MoFaSGD deliberately avoids accumulating moments *within* the subspace, unlike GaLore, thereby avoiding potential error propagation from stale subspaces. MoFaSGD directly uses the computed momentum factors ($\mathbf{U}_{t+1}, \mathbf{V}_{t+1}$) to perform the spectrally normalized update in Equation 9. This design choice aims to circumvent potential errors arising from subspace moment accumulation, particularly when the subspace changes frequently (i.e., near-online updates, or small subspace update intervals in GaLore). As empirically supported in Section 5.3, frequent subspace updates can indeed negatively impact GaLore’s performance, suggesting that subspace moment accumulation errors might increase with the frequency of subspace changes.

MoFaSGD’s *novelty* lies in its unique combination of: 1) Projecting gradients onto the dynamically updated tangent space derived from the low-rank *momentum* factors, and 2) Utilizing these factors directly for spectrally normalized updates, thereby bypassing the potential pitfalls of subspace moment accumulation inherent in methods like GaLore.

4.2 Convergence and Theoretical Analysis

Our analysis addresses the non-convex optimization problem:

$$\min_{\mathbf{W} \in \mathbb{R}^{m \times n}} \mathcal{L}(\mathbf{W}) = \mathbb{E}_\xi [\mathcal{L}(\mathbf{W}, \xi)] \quad (10)$$

where we assume access to an unbiased, variance-bounded stochastic gradient oracle $\nabla \mathcal{L}(\mathbf{W}, \xi)$. Below, we first introduce the necessary definitions and assumptions that underpin our theoretical results.

Definitions and Assumptions

For any optimization iterate \mathbf{W}_i , we denote the full-batch gradient by $\bar{\mathbf{G}}_i = \nabla \mathcal{L}(\mathbf{W}_i)$ and the stochastic gradient by $\mathbf{G}_i = \nabla \mathcal{L}(\mathbf{W}_i, \xi_i)$. Formally, we leverage the following standard assumptions throughout our analysis:

Assumption 4.1. $\mathcal{L}(\cdot)$ is L -smooth with respect to the nuclear norm $\|\cdot\|_*$. In other words, for any two arbitrary $\mathbf{W}_1, \mathbf{W}_2 \in \mathbb{R}^{m \times n}$, we have: $\|\nabla \mathcal{L}(\mathbf{W}_1) - \nabla \mathcal{L}(\mathbf{W}_2)\|_* \leq L \|\mathbf{W}_1 - \mathbf{W}_2\|_2$

This assumption naturally generalizes the typical smoothness condition from vector optimization and has been previously utilized in the literature (Large et al., 2025; Bernstein & Newhouse, 2024b). For further details, please see Appendix D.1. Additionally, we assume the availability of a stochastic gradient oracle satisfying standard properties:

Assumption 4.2. For any model parameter \mathbf{W} , we have access to an unbiased and variance-bounded stochastic oracle, as follows: $\mathbb{E}_\xi[\nabla \mathcal{L}(\mathbf{W}, \xi)] = \nabla \mathcal{L}(\mathbf{W})$, and $\mathbb{E}_\xi[\|\nabla \mathcal{L}(\mathbf{W}, \xi) - \nabla \mathcal{L}(\mathbf{W})\|_*] \leq \sigma$

With these definitions and assumptions clarified, we now present an intuitive overview of our main theoretical results, which highlight the strengths and optimality of our proposed MoFaSGD algorithm.

Optimality of Tangent Space Projection (Theorem 4.3). We establish that projecting each gradient \mathbf{G}_t onto the tangent space defined by its singular vectors achieves the minimal projection residual error among a broad class of low-rank projection schemes, such as projection onto the left or right singular vector subspaces. **Optimal $O(1/\sqrt{T})$ Convergence Rate (Theorem 4.5).** Under Assumptions 4.1 and 4.2, MoFaSGD achieves convergence to a stationary point at the optimal $O(1/\sqrt{T})$ rate. Critically, the factorization of momentum does not degrade the asymptotic convergence rate.

Proof Outline

Our proof structure involves three primary steps. First, we decompose the momentum low-rank approximation error by defining the full-rank momentum as $\mathbf{M}_t = \sum_{i=0}^t \beta^{t-i} \mathbf{G}_i$, and breaking down the approximation error into two key components: $\|\hat{\mathbf{M}}_t - \bar{\mathbf{G}}_t\|_* \leq \|\mathbf{M}_t - \bar{\mathbf{G}}_t\|_* + \|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_*$. These terms are individually controlled via exponential averaging and tangent-space projections. Second, we rigorously prove the optimality of the tangent-space projection (Theorem 4.3), demonstrating that $\hat{\mathbf{G}}_t = \mathbf{U}_t \mathbf{U}_t^\top \mathbf{G}_t + \mathbf{G}_t \mathbf{V}_t \mathbf{V}_t^\top - \mathbf{U}_t \mathbf{U}_t^\top \mathbf{G}_t \mathbf{V}_t \mathbf{V}_t^\top$ minimizes the residual $\|\mathbf{G}_t - \hat{\mathbf{G}}_t\|_F$, thus ensuring optimal fitting into the evolving momentum subspace. Lastly, by combining the upper bounds on the aforementioned terms with a standard descent lemma under nuclear-norm smoothness, we derive our final convergence behavior. By carefully bounding the approximation terms from the prior steps and leveraging them in our derived descent lemma under nuclear-norm smoothness, we arrive at our main convergence result.

The low-rank factorization of gradient momentum as described by Equation Low-rank Moment Factors is the cornerstone of our proposed method. The quality of the momentum approximation plays a key role in the effectiveness of our approach. We provide an intuitive sketch of the theoretical analysis to bound the factorization residual $\|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_* = \|\sum_{i=1}^t \beta^{t-i} \mathbf{G}_i - \mathbf{U}_{t+1} \boldsymbol{\Sigma}_{t+1} \mathbf{V}_{t+1}^\top\|_F^2$. By recursively bounding this term, we can show that it is sufficient to bound the term $\|\mathbf{U}_{t+1} \boldsymbol{\Sigma}_{t+1} \mathbf{V}_{t+1}^\top - \beta \mathbf{U}_t \boldsymbol{\Sigma}_t \mathbf{V}_t^\top - \mathbf{G}_t\|_F$, which can itself be shown to be bounded by $\|\text{Proj}_{\mathcal{T}_t}(\mathbf{G}_t) - \mathbf{G}_t\|_F$. Thus, the quality of the factored momentum approximation is directly related to the residual of the gradient’s low-rank projection.

Results

The choice of tangent space projection is optimal in the sense of minimizing the term $\|\text{Proj}_{\mathcal{T}_t}(\mathbf{G}_t) - \mathbf{G}_t\|_F$, as demonstrated in the following theorem.

Theorem 4.3. Let $\mathbf{L} \in \mathbb{R}^{m \times r}$, $\mathbf{R} \in \mathbb{R}^{n \times r}$ be any arbitrary sketching matrices, and $(\alpha_1, \alpha_2, \alpha_3)$ be any arbitrary triple of scalars. Let $\text{Proj}_{(\mathbf{L}, \mathbf{R})}(\mathbf{G}) = \alpha_1 \mathbf{L} \mathbf{L}^\top \mathbf{G} + \alpha_2 \mathbf{G} \mathbf{R} \mathbf{R}^\top + \alpha_3 \mathbf{L} \mathbf{L}^\top \mathbf{G} \mathbf{R} \mathbf{R}^\top$. Then, the projection residual is minimized when (1) $(\alpha_1, \alpha_2, \alpha_3) = (1, 1, -1)$ and (2) $\mathbf{L}^\top \mathbf{L} = \mathbf{R}^\top \mathbf{R} = \mathbf{I}_r$. In this case, the residual norm is $\|\text{Proj}_{(\mathbf{L}, \mathbf{R})}(\mathbf{G}_t) - \mathbf{G}_t\| = \|(I - \mathbf{L} \mathbf{L}^\top) \mathbf{G}_t (I - \mathbf{R} \mathbf{R}^\top)\|$.

Remark 4.4. If we let $\mathbf{L} = \mathbf{U}_{\mathbf{G}_t}^{1:r}$ and $\mathbf{R} = \mathbf{V}_{\mathbf{G}_t}^{1:r}$, then we can conclude that the residual error would be upper bounded by $\sigma_{\mathbf{G}_t}^{r+1:\min(m,n)} = \sum_{i=r+1}^{\min(m,n)} \sigma_{\mathbf{G}_t}^i$, which, considering the low-rank property of the gradient, we expect it to be small, or even zero if the rank of the full gradient is less than r . Note that the subspace projection in GaLore (Zhao et al., 2024a) is actually equivalent to letting either $(\alpha_1, \alpha_2, \alpha_3) = (0, 0, 1)$ or $(\alpha_1, \alpha_2, \alpha_3) = (1, 0, 0)$.

We now present our main convergence bound for Algorithm 1.

Theorem 4.5. *Let Assumptions 4.1 and 4.2 hold. Moreover, assume $\text{rank}(\mathbf{G}_0) \leq r$. By letting $\beta \leq \frac{1}{3}$ and $\eta \leq 1$, the iterates of Algorithm 1 satisfy the following:*

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E}[\|\nabla \mathcal{L}(\mathbf{w}_t)\|_*] \leq \mathcal{O} \left(\frac{\mathcal{L}(\mathbf{w}_0) - \mathbb{E}[\mathcal{L}(\mathbf{w}_{T+1})]}{\eta T} + \eta L + \frac{\sigma}{\sqrt{T}} \right) \quad (11)$$

Moreover, if we set $\eta = \Theta \left(\sqrt{\frac{\mathcal{L}(\mathbf{w}_0) - \mathbb{E}[\mathcal{L}(\mathbf{w}_{T+1})]}{TL}} \right)$, we can derive the following simplified bound as:

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E}[\|\nabla \mathcal{L}(\mathbf{w}_t)\|_*] \leq \mathcal{O} \left(\frac{(\mathcal{L}(\mathbf{w}_0) - \mathbb{E}[\mathcal{L}(\mathbf{w}_{T+1})])^{\frac{1}{2}} \sqrt{L} + \sigma}{\sqrt{T}} \right) \quad (12)$$

Remark 4.6. The convergence metric used for Theorem 4.5 is the stationary point with respect to the nuclear norm, which is averaged over all gradients. When setting $\sigma = 0$, Algorithm 1 achieves the rate of $\mathcal{O}(\frac{1}{\sqrt{T}})$, which is known to be optimal in the sense of non-convex stochastic optimization under smoothness and an unbiased, bounded stochastic gradient oracle (Arjevani et al., 2023).

5 Experiments

We evaluate MoFaSGD’s effectiveness and efficiency across three large language modeling setups: pre-training, natural language understanding (NLU) fine-tuning, and instruction-tuning. These setups allow us to assess MoFaSGD’s performance across different training regimes and task complexities.

5.1 Pre-training setup: NanoGPT Speedrun

We first evaluate MoFaSGD in a pre-training context using the Modded NanoGPT benchmark (Jordan et al., 2024a). This benchmark focuses on training a GPT-2 architecture on a subset of the FineWeb dataset (Penedo et al., 2025) and measures performance using validation perplexity on a held-out partition of FineWeb. The benchmark’s default optimizer is Muon (Jordan et al., 2024b), which holds current training speed records for this task and serves as a strong, competitive baseline.

We compare MoFaSGD against full fine-tuning baselines AdamW (Loshchilov et al., 2017) and Muon, as well as the low-rank baseline GaLore (Zhao et al., 2024a). Following the standard NanoGPT speedrun setup, we use the hyperparameters tuned for Muon. We tune the learning rates for AdamW, GaLore, and MoFaSGD, along with GaLore’s SVD frequency and MoFaSGD’s momentum decay (β) via grid search, selecting the best configuration based on final validation perplexity (details in Appendix C.2). We evaluate ranks $r \in \{16, 32, 128\}$, common choices for low-rank methods in pre-training setup.

Our primary experiment uses a budget of 0.73 billion tokens from FineWeb, aligning with the budget used by Muon to reach the target perplexity of 3.27. Figure 3a shows the validation perplexity curves. MoFaSGD consistently outperforms GaLore across all tested ranks, achieving lower final perplexity. The performance advantage is particularly noticeable at lower ranks ($r = 16$). This suggests MoFaSGD’s dynamic subspace tracking is more effective at capturing important momentum directions than GaLore’s infrequent updates, especially under strict rank constraints. Both low-rank methods underperform the full-rank AdamW and Muon baselines within this specific token budget. This is likely because full-rank methods have more degrees of freedom, and the 0.73B token budget, optimized for Muon’s convergence speed, might be insufficient for low-rank methods to fully match their performance.

To assess longer-term performance, we conduct an extended run for 10,000 steps ($\sim 5.3B$ tokens) using rank $r = 32$ for both MoFaSGD and GaLore. As shown in Figure 3b, MoFaSGD maintains its performance advantage over GaLore, indicating that the benefits of its momentum factorization approach persist during longer training phases.

Ablation: Convergence vs. Efficiency

We conduct a detailed ablation study on the effect of the low-rank parameter $r \in \{16, 32, 128\}$ during NanoGPT pre-training for both GaLore and MoFaSGD. As illustrated in Figure 1 and Figure 2, higher ranks consistently improve convergence speed and final validation loss. MoFaSGD achieves smoother loss curves and stronger performance across all ranks, particularly under tight memory budgets (e.g., $r = 16$), where GaLore exhibits noticeable instability. This supports our claim that MoFaSGD’s tangent-space subspace tracking better preserves optimizer continuity under aggressive compression.

In terms of runtime (Table 1), GaLore shows minimal runtime variation across ranks due to dominant offline SVD costs. In contrast, MoFaSGD’s runtime scales more significantly with rank, reflecting its per-step online factorization cost. Nonetheless, MoFaSGD remains faster than GaLore at $r = 32$ and achieves superior final loss, highlighting a favorable trade-off between expressivity and computational efficiency.

Table 1: Comparison of MoFaSGD and GaLore across different ranks during NanoGPT pre-training. Best values per row are in **bold**.

| Rank | Final Val Loss | | Runtime (s) | | Throughput | |
|------|----------------|--------|-------------|-------------|----------------|----------------|
| | MoFaSGD | GaLore | MoFaSGD | GaLore | MoFaSGD | GaLore |
| 16 | 3.8981 | 4.0773 | 2156 | 3755 | 338,450 | 194,395 |
| 32 | 3.7208 | 3.8953 | 3972 | 3911 | 183,770 | 186,619 |
| 128 | 3.5700 | 3.6561 | 4817 | 3839 | 151,527 | 190,131 |




Figure 1: Validation loss vs. training steps for MoFaSGD (left) and GaLore (right) across ranks $r \in \{16, 32, 128\}$. MoFaSGD shows smoother and faster convergence.




Figure 2: Validation loss vs. wall-clock time across ranks. MoFaSGD scales better in convergence and runtime efficiency.

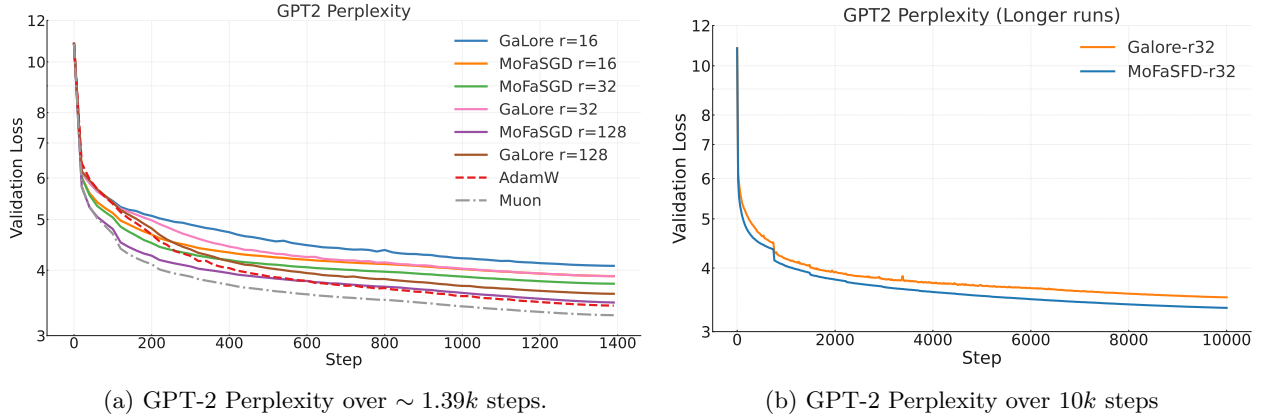


Figure 3: Comparison of optimizer performance on GPT-2 using validation perplexity loss.

5.2 Post-training Evaluation: Setups and Results

We evaluate MoFaSGD in two post-training scenarios: standard NLU fine-tuning on the GLUE benchmark and large-scale instruction-tuning on the Tulu3 dataset.

GLUE Benchmark. We follow the experimental setup detailed in Zhao et al. (2024a) for fine-tuning a RoBERTa-Base model (125M parameters) on seven diverse NLU tasks from the GLUE benchmark (Wang, 2018): MNLI, QQP, SST-2, MRPC, COLA, QNLI, and RTE. We use the same hyperparameters as Zhao et al. (2024a) for comparison baselines to ensure fairness. For MoFaSGD (with ranks $r = 4$ and $r = 8$), we tune the learning rate for each task based on validation accuracy, keeping the momentum decay β fixed at 0.95 (details in Appendix C.3). Final validation accuracy and loss are used for evaluation.

Tulu3 Instruction Tuning. To assess performance on more complex alignment tasks, we fine-tune the LLaMA-3.1 8B model on the `tulu-3-sft-mixture` dataset, a large ($\sim 900K$ samples) and diverse instruction-tuning dataset (Lambert et al., 2024). We adopt most hyperparameters from the Tulu3 setup (Lambert et al., 2024), training for one epoch with an effective batch size of 128; however, we use only a subsample of 200K examples to reduce the training budget. We perform a grid search over learning rates for each optimizer. For the low-rank methods (MoFaSGD, LoRA, GaLore), we use rank $r = 8$. Key hyperparameters for MoFaSGD (β) and the baselines (GaLore SVD frequency, LoRA alpha) are set as specified in Appendix C.4. We evaluate the final checkpoints using the OLMES evaluation framework (Gu et al., 2024), benchmarking across MMLU, TruthfulQA, BigBenchHard, GSM8K, and HumanEval.

Performance Analysis

In these post-training tasks, we compare MoFaSGD against LoRA (Hu et al., 2021) optimized with AdamW, and GaLore (Zhao et al., 2024a), which are prevalent methods for memory-efficient fine-tuning. We also include results for full-parameter fine-tuning with AdamW as a performance ceiling reference. Table 2 compares the theoretical complexities, highlighting MoFaSGD’s comparable memory footprint to LoRA alongside efficient online subspace updates, while Figure 4 comprehensively details the memory usage of MoFaSGD compared to other baselines in our Tulu3 instruction-tuning setup.

To illustrate MoFaSGD’s optimization efficiency during instruction tuning on the Tulu3 benchmark, we compare its validation loss trajectory against GaLore and LoRA across both training epochs and wall-clock time. As shown in Figure 5a and Figure 5b, MoFaSGD consistently achieves lower validation loss over the course of training, indicating superior sample efficiency. Notably, when measured against real-world wall-clock time, MoFaSGD converges faster than both GaLore and LoRA, demonstrating improved practical efficiency in addition to theoretical gains. These trends are further supported by our throughput analysis: MoFaSGD reaches 4206 tokens/sec, outperforming GaLore (3214 tokens/sec) and approaching the high throughput of LoRA (4536 tokens/sec). These findings reinforce our earlier conclusion that MoFaSGD’s

| Optimizer | Memory Complexity | Subspace Resampling |
|-----------|--------------------|----------------------------------|
| GaLore | $mn + mr + 2nr$ | $\mathcal{O}(m^2n)$ (offline) |
| LoRA | $mn + 3mr + 3nr$ | — |
| MoFaSGD | $mn + mr + nr + r$ | $\mathcal{O}((m+n)r^3)$ (online) |

Table 2: Comparison of memory and subspace resampling complexity for low-rank optimizers. Let $\mathbf{W} \in \mathbb{R}^{m \times n}$ represent model parameters, and r is the rank of low-rank optimizers (w.l.o.g assume $m \leq n$). Note that memory complexity includes model parameters and optimizer states.


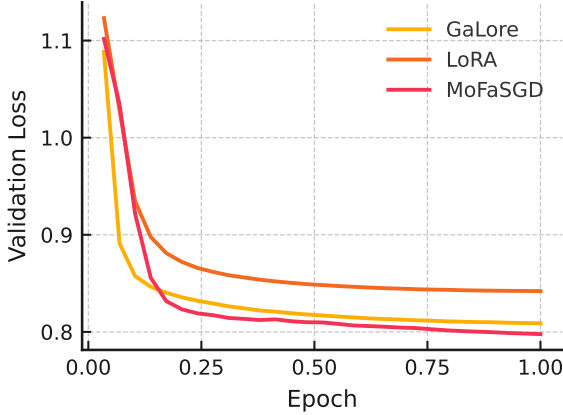



Figure 4: Empirical memory breakdown (GB) for LLaMA3.1-8B using different optimizers.

spectrally normalized updates and dynamic momentum factorization enable more effective fine-tuning under memory constraints. We have also included the training loss curves in Appendix C.5, showcasing MoFaSGD’s convergence behavior on both GLUE and Tulu3 setups.



(a) Validation Loss vs. Epoch



(b) Validation Loss vs. Wall-clock Time

Figure 5: MoFaSGD demonstrates superior sample efficiency and faster wall-clock convergence, achieving lower validation loss than both GaLore and LoRA throughout LLaMA3.1-8B instruction tuning on the Tulu3 benchmark.

Table 3 summarizes the final validation accuracies on the GLUE tasks (ranks $r=4$ and $r=8$). MoFaSGD achieves performance comparable to, and on average slightly better (+0.5% for $r=4$, +0.47% for $r=8$) than both LoRA and GaLore, while using slightly less estimated memory. This demonstrates MoFaSGD’s competitiveness and memory efficiency on standard NLU benchmarks.

Table 4 presents the results on the challenging instruction-tuning benchmarks using the Tulu3 setup (rank $r=8$). MoFaSGD outperforms both GaLore (+0.8% avg.) and LoRA (+2.3% avg.) on the average score across the five benchmarks. This highlights MoFaSGD’s potential advantage in complex tasks where accurately capturing the training dynamics over long sequences and diverse instructions is crucial. However, consistent with prior work (Wang et al., 2023), all low-rank methods exhibit a performance gap compared to full fine-tuning with AdamW (MoFaSGD is -4.2% avg. vs AdamW). This gap underscores the inherent trade-off between memory efficiency and achievable performance, particularly as task complexity increases and may require capturing subtle, high-rank parameter updates that low-rank approximations inherently miss.

Table 3: Comparison of final validation accuracies (%) on seven GLUE tasks (MNLI, QQP, SST-2, MRPC, CoLA, QNLI, RTE) when fine-tuning a RoBERTa-base model using different optimizers. For GaLore, LoRA, and MoFaSGD, we report results with rank $r \in \{4, 8\}$. Memory usage is estimated for each method, and the final column shows the average accuracy. Note that for memory measurement, we include only the parameters and the optimizer states for a fair comparison.

| Optimizer | MNLI | QQP | SST-2 | MRPC | CoLA | QNLI | RTE | Memory | Avg. |
|---------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------|-------|
| AdamW (Full-Rank) | 86.8 | 91.98 | 94.48 | 90.90 | 62.25 | 93.15 | 79.41 | 747M | 85.57 |
| GaLore ($r = 4$) | 85.23 | 89.62 | 94.17 | 90.72 | 60.33 | 93.20 | 77.22 | 253M | 84.36 |
| LoRA ($r = 4$) | 84.25 | 89.73 | 93.59 | 90.53 | 60.42 | 92.91 | 78.59 | 257M | 84.29 |
| MoFaSGD ($r = 4$) | 85.12 | 89.85 | 94.15 | 90.78 | 61.91 | 93.10 | 79.08 | 251M | 84.86 |
| GaLore ($r = 8$) | 86.01 | 89.65 | 94.04 | 90.65 | 59.96 | 93.16 | 78.14 | 257M | 84.52 |
| LoRA ($r = 8$) | 85.18 | 90.15 | 93.87 | 90.82 | 61.11 | 93.05 | 78.77 | 264M | 84.71 |
| MoFaSGD ($r = 8$) | 86.32 | 90.26 | 94.36 | 90.75 | 62.16 | 93.12 | 79.28 | 253M | 85.18 |

Table 4: Final scores of Llama-3.1 8B on the Tulu3-SFT-mixture dataset using four different optimizers. The table reports performance on MMLU, TruthfulQA, BigBenchHard, GSM8K, and HumanEval, along with the average of these five benchmarks (Avg.).

| Optimizer | MMLU | TruthfulQA | BigBenchHard | GSM8K | HumanEval | Avg. |
|-------------------|-------------|-------------|--------------|-------------|-------------|-------------|
| AdamW (Full-Rank) | 62.8 | 46.5 | 66.7 | 72.7 | 81.0 | 65.9 |
| GaLore | 58.9 | 44.2 | 57.6 | 68.4 | 75.2 | 60.9 |
| LoRA | 56.1 | 42.6 | 56.5 | 67.9 | 73.9 | 59.4 |
| MoFaSGD | 59.4 | 45.8 | 58.3 | 68.4 | 76.8 | 61.7 |

5.3 Ablations

Momentum Spectral Analysis

MoFaSGD is motivated by the conjecture that the first moment (the EMA of gradients) preserves a low-rank structure throughout training. This conjecture stems from the GaLore hypothesis on the low-rankness of the gradients themselves (Zhao et al., 2024a), and, more importantly, from the observation in Feinberg et al. (2024), which shows a fast spectral decay in the EMA of the gradient covariance, as discussed in more detail in Section 4.1. To investigate our conjecture, we analyze the first-moment buffer \mathbf{M}_t from the AdamW optimizer states generated during the Tulu3 instruction-tuning setup.

For each relevant parameter matrix \mathbf{M}_t , we perform SVD and compute the energy ratio captured by the top- r singular values as $\frac{\sum_{i=1}^r \sigma_{i, \mathbf{M}_t}^2}{\|\mathbf{M}_t\|_F^2}$. This ratio represents the percentage of the momentum’s total energy contained within the top- r subspace. We compute the average of this ratio across all 2D weight matrices in the model at various training steps. Figure 6a shows the average energy ratio for $r = 16$ and $r = 32$ throughout training. We observe that the top-32 singular values consistently capture around 80% of the momentum’s energy, while the top-16 capture approximately 75%. This persistent and significant concentration of energy in a low-rank subspace strongly supports our hypothesis.

Impact of GaLore’s Subspace Update Frequency

We analyze the impact of subspace update frequency in methods like GaLore compared to MoFaSGD’s implicit per-iteration adaptation. MoFaSGD updates its momentum factors at every step, while GaLore performs explicit, costly subspace updates (e.g., full SVD on the gradient) at intervals τ . To investigate whether simply increasing GaLore’s update frequency (decreasing τ) matches the benefits of MoFaSGD’s online subspace adaptation, we conduct an ablation study on τ using the NanoGPT pre-training setup (

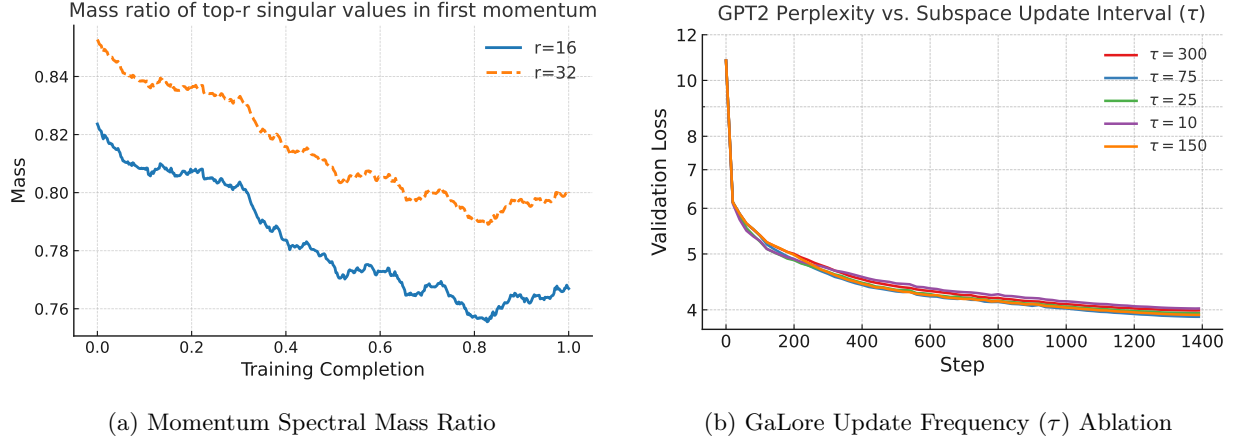


Figure 6: (a) Average mass ratio of the AdamW first moment captured by top- r singular vectors during Tulu3 fine-tuning ($r = 16, 32$). (b) Validation perplexity vs. subspace update interval (τ) for GaLore ($r = 32$) on NanoGPT.

0.73B token budget with rank $r = 32$). We vary τ across $\{10, 25, 75, 150, 300\}$ steps. Figure 6b shows the validation perplexity curves.

Very frequent updates ($\tau = 10$ or $\tau = 25$) do not yield the best performance and are slightly worse than less frequent updates (e.g., $\tau = 150$). This aligns with GaLore’s findings (Zhao et al., 2024a) and suggests that overly frequent subspace changes in GaLore can disrupt optimizer state accumulation. This finding underscores the challenge of online subspace changes. This finding suggests that MoFaSGD’s approach of avoiding abrupt subspace changes offers a more stable and efficient way to leverage low-rank structures.

5.4 Memory Usage Breakdown and Profiling

We assess MoFaSGD’s memory efficiency by decomposing its GPU memory usage across five categories: parameters, optimizer states, gradients, activations, and adapters. Figure 4 shows a comparative breakdown of memory consumption across six optimizer setups on LLaMA3.1-8B. MoFaSGD achieves total memory usage of 29.4 GB, competitive with fused GaLore and LoRA, while enabling full-parameter updates. In contrast, AdamW exceeds 70 GB due to high-cost full-rank momentum buffers and persistent gradient accumulation. SWAN and GaLore (non-fused) similarly suffer from gradient buffers, which dominate their memory footprints.

These savings arise from three key design elements: (i) eliminating second-moment buffers entirely, (ii) maintaining a low-rank SVD factorization of first-order momentum, and (iii) fusing gradient projection and zeroing operations during backpropagation, which prevents gradient accumulation from persisting across steps.

To further validate these results, Figure 7 shows the memory trace during MoFaSGD training. We observe a clean separation between parameter storage, a narrow and persistent optimizer state band, and tightly bounded gradient memory. Compared to AdamW (Appendix C.6), which shows 16 GB persistent gradient buffers and 32 GB optimizer states, MoFaSGD significantly reduces runtime memory pressure. A full quantitative table with GB-level memory usage across all optimizers is provided in Appendix C.6.

5.5 Implementation Details

Initialization. To initialize the momentum factors in MoFaSGD, we perform a full singular value decomposition (SVD) once at the beginning of training, using the gradient from the first step. Specifically, we set $U_0 = U_{G_0}^{1:r}$, $V_0 = V_{G_0}^{1:r}$, and $\Sigma_0 = \Sigma_{G_0}^{1:r}$, corresponding to the top- r components. We apply MoFaSGD exclusively to the linear layers of transformer blocks, as spectral normalization is particularly effective for




Figure 7: GPU memory trace during LLaMA3.1-8B training with MoFaSGD ($r = 8$).

these layers. Notably, both GaLore (Zhao et al., 2024a) and Muon (Jordan et al., 2024b) adopt similar approach: applying their custom optimizers only to transformer linear layers while using AdamW for embedding weights and 1D layers. We follow this convention and use AdamW (in `bf16`) for those remaining layers. Consequently, the GPU memory usage for optimizer states—including in fused-GaLore and MoFaSGD—is approximately 4.2 GB.

Gradient Accumulation and Fused Implementation

Gradient accumulation is essential for training large models with limited hardware memory. For methods such as GaLore (Zhao et al., 2024a) and MoFaSGD, promptly clearing gradient buffers via backward hooks is crucial; otherwise, memory savings akin to LoRA are not realized (see Figure 14). To address this, GaLore performs a fused optimizer step, where the gradient is used to update the corresponding parameter immediately during the backward pass, after which the gradient buffer is cleared.

However, this approach is incompatible with gradient accumulation over multiple micro-batches, since the optimizer step must occur only after all gradients are accumulated. To resolve this, we introduce low-rank gradient buffers into the optimizer state and register a backward hook that accumulates low-rank projected gradients instead of performing the optimizer step immediately. For MoFaSGD, updating the momentum factors only requires the low-rank projections $\mathbf{G}_t \mathbf{V}_t$, $\mathbf{U}_t^\top \mathbf{G}_t$, and $\mathbf{U}_t^\top \mathbf{G}_t \mathbf{V}_t$. We thus implement a temporary low-rank gradient buffer, enabling gradient accumulation while avoiding the need for a persistent full-rank buffer (see memory usage in Figure 7).

A similar strategy is used for GaLore: we implement a gradient-accumulation-friendly fused version by storing full-rank gradients in low-rank form. Since GaLore only needs the projection $\mathbf{Q}_t^\top \mathbf{G}_t$ to update its subspace momentum, low-rank accumulation suffices.

Stateless optimizers

Recent stateless optimizers such as SWAN (Ma et al., 2024) and SinkGD (Scetbon et al., 2025) currently do not have open-source implementations, and we encountered challenges in achieving stable convergence with our preliminary implementations. Nevertheless, given the structural similarity between SWAN and Muon (Jordan et al., 2024b) without momentum, we approximate the memory usage of such approaches by profiling Muon with its momentum buffer disabled. These results are reported in Figure 4 under the label "SWAN" as a representative proxy for stateless optimizers.

One technical consideration with these methods is their compatibility with gradient accumulation, which is a common practice in low-resource settings. In such scenarios, the fused backward strategy may not be applicable, and persistent gradient buffers are still required in memory (see Figure 11). While stateless optimizers are promising from a conceptual standpoint, this constraint currently poses practical challenges for achieving memory efficiency on par with parameter-efficient fine-tuning approaches like LoRA, or subspace-based approaches like fused GaLore and MoFaSGD.

6 Conclusion

We introduced MoFaSGD, a memory-efficient optimizer that uses memory comparable to LoRA-like methods while achieving strong convergence. The core idea is to maintain a low-rank factorization of momentum as the optimizer state and leverage this representation to directly update model parameters at each iteration using spectrally normalized updates, building upon the success of similar approaches in full-training scenarios. We provide a comprehensive theoretical analysis and establish an upper bound on convergence for stochastic non-convex optimization, matching existing lower bounds. Moreover, we empirically demonstrate the effectiveness of our method compared to standard low-rank optimization approaches in both pre-training and post-training setups.

Limitations and Future Directions. MoFaSGD shows strong empirical performance, but several limitations suggest promising directions for future work. First, as a low-rank subspace optimizer, MoFaSGD may underperform compared to full-rank methods on more complex tasks. Exploring adaptive or dynamic rank selection strategies, beyond the fixed-rank settings used in our study, could help mitigate this gap; for instance, future work could investigate monitoring projection residuals to allocate more rank to layers with higher approximation error or implementing a budget-aware allocation scheme to optimally balance performance and memory. Second, although our theoretical results establish convergence guarantees, they rely on assumptions such as nuclear-norm smoothness. The practical relevance and limitations of such assumptions in deep learning setups remain an open area for study.

Acknowledgment

This work was partially supported by NSF CAREER Award #2239374.

References

- Rohan Anil, Vineet Gupta, Tomer Koren, and Yoram Singer. Memory efficient adaptive optimization. *Advances in Neural Information Processing Systems*, 32, 2019.
- Yossi Arjevani, Yair Carmon, John C Duchi, Dylan J Foster, Nathan Srebro, and Blake Woodworth. Lower bounds for non-convex stochastic optimization. *Mathematical Programming*, 199(1):165–214, 2023.
- Jeremy Bernstein and Laker Newhouse. Modular duality in deep learning. *arXiv preprint arXiv:2410.21265*, 2024a.
- Jeremy Bernstein and Laker Newhouse. Old optimizer, new norm: An anthology. *arXiv.org*, 2024b. doi: 10.48550/arxiv.2409.20325.
- Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, and Animashree Anandkumar. signsgd: Compressed optimisation for non-convex problems. In *International Conference on Machine Learning*, pp. 560–569. PMLR, 2018.
- Jeremy Bernstein, Chris Mingard, Kevin Huang, Navid Azizan, and Yisong Yue. Automatic gradient descent: Deep learning without hyperparameters. *arXiv preprint arXiv:2304.05187*, 2023.
- Xiangning Chen, Chen Liang, Da Huang, Esteban Real, Kaiyuan Wang, Hieu Pham, Xuanyi Dong, Thang Luong, Cho-Jui Hsieh, Yifeng Lu, et al. Symbolic discovery of optimization algorithms. *Advances in neural information processing systems*, 36, 2024.
- Tim Dettmers, Mike Lewis, Sam Shleifer, and Luke Zettlemoyer. 8-bit optimizers via block-wise quantization. *arXiv preprint arXiv:2110.02861*, 2021.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36, 2024.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.

- Sai Surya Duvvuri, Fnu Devvrit, Rohan Anil, Cho-Jui Hsieh, and Inderjit S Dhillon. Combining axes preconditioners through kronecker approximation for deep learning. In *The Twelfth International Conference on Learning Representations*, 2024.
- Vladimir Feinberg, Xinyi Chen, Y Jennifer Sun, Rohan Anil, and Elad Hazan. Sketchy: Memory-efficient adaptive regularization with frequent directions. *Advances in Neural Information Processing Systems*, 36, 2024.
- Thomas George, César Laurent, Xavier Bouthillier, Nicolas Ballas, and Pascal Vincent. Fast approximate natural gradient descent in a kronecker factored eigenbasis. *Advances in neural information processing systems*, 31, 2018.
- Frithjof Gressmann, Zach Eaton-Rosen, and Carlo Luschi. Improving neural network training in low dimensional random bases. *Advances in Neural Information Processing Systems*, 33:12140–12150, 2020.
- Roger Grosse and James Martens. A kronecker-factored approximate fisher matrix for convolution layers. In *International Conference on Machine Learning*, pp. 573–582. PMLR, 2016.
- Yuling Gu, Oyvind Tafjord, Bailey Kuehl, Dany Haddad, Jesse Dodge, and Hannaneh Hajishirzi. Olmes: A standard for language model evaluations. *arXiv preprint arXiv:2406.08446*, 2024.
- Vineet Gupta, Tomer Koren, and Yoram Singer. Shampoo: Preconditioned stochastic tensor optimization. In *International Conference on Machine Learning*, pp. 1842–1850. PMLR, 2018.
- Guy Gur-Ari, Daniel A Roberts, and Ethan Dyer. Gradient descent happens in a tiny subspace. *arXiv preprint arXiv:1812.04754*, 2018.
- Yongchang Hao, Yanshuai Cao, and Lili Mou. Flora: Low-rank adapters are secretly gradient compressors. *arXiv preprint arXiv:2402.03293*, 2024.
- Roger A Horn and Charles R Johnson. *Topics in matrix analysis*. Cambridge university press, 1994.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pp. 2790–2799. PMLR, 2019.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Keller Jordan, Jeremy Bernstein, Brendan Rappazzo, @fernbear.bsky.social, Boza Vlado, You Jiacheng, Franz Cesista, Braden Koszarsky, and @Grad62304977. modded-nanogpt: Speedrunning the nanogpt baseline, 2024a. URL <https://github.com/KellerJordan/modded-nanogpt>.
- Keller Jordan, Yuchen Jin, Vlado Boza, You Jiacheng, Franz Cecista, Laker Newhouse, and Jeremy Bernstein. Muon: An optimizer for hidden layers in neural networks, 2024b. URL <https://kellerjordan.github.io/posts/muon/>.
- Damjan Kalajdzievski. A rank stabilization scaling factor for fine-tuning with lora. *arXiv preprint arXiv:2312.03732*, 2023.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- Priya Kasimbeg, Frank Schneider, Runa Eschenhagen, Juhan Bae, Chandramouli Shama Sastry, Mark Saroufim, Boyuan Feng, Less Wright, Edward Z Yang, Zachary Nado, et al. Accelerating neural network training: An analysis of the algoperf competition. *arXiv preprint arXiv:2502.15015*, 2025.

- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations (ICLR)*, 2015.
- Dawid J Kopiczko, Tijmen Blankevoort, and Yuki M Asano. Vera: Vector-based random matrix adaptation. *arXiv preprint arXiv:2310.11454*, 2023.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, et al. Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124*, 2024.
- Tim Large, Yang Liu, Jacob Huh, Hyojin Bahng, Phillip Isola, and Jeremy Bernstein. Scalable optimization in the modular norm. *Advances in Neural Information Processing Systems*, 37:73501–73548, 2025.
- Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021.
- Bingrui Li, Jianfei Chen, and Jun Zhu. Memory efficient optimizers with 4-bit states. *Advances in Neural Information Processing Systems*, 36:15136–15171, 2023.
- Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021.
- Vladislav Lialin, Sherin Muckatira, Namrata Shivagunde, and Anna Rumshisky. Relora: High-rank training through low-rank updates. In *The Twelfth International Conference on Learning Representations*, 2023.
- Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1):503–528, 1989.
- Jingyuan Liu, Jianlin Su, Xingcheng Yao, Zhejun Jiang, Guokun Lai, Yulun Du, Yidao Qin, Weixin Xu, Enzhe Lu, Junjie Yan, et al. Muon is scalable for llm training. *arXiv preprint arXiv:2502.16982*, 2025.
- Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. Dora: Weight-decomposed low-rank adaptation. *arXiv preprint arXiv:2402.09353*, 2024.
- Ilya Loshchilov, Frank Hutter, et al. Fixing weight decay regularization in adam. *arXiv preprint arXiv:1711.05101*, 5, 2017.
- Qijun Luo, Hengxu Yu, and Xiao Li. Badam: A memory efficient full parameter training method for large language models. *arXiv preprint arXiv:2404.02827*, 2024.
- Yang Luo, Xiaozhe Ren, Zangwei Zheng, Zhuo Jiang, Xin Jiang, and Yang You. Came: Confidence-guided adaptive memory efficient optimization. *arXiv preprint arXiv:2307.02047*, 2023.
- Kai Lv, Hang Yan, Qipeng Guo, Haijun Lv, and Xipeng Qiu. Adalomo: Low-memory optimization with adaptive learning rate. *arXiv preprint arXiv:2310.10195*, 2023a.
- Kai Lv, Yuqing Yang, Tengxiao Liu, Qinghui Gao, Qipeng Guo, and Xipeng Qiu. Full parameter fine-tuning for large language models with limited resources. *arXiv preprint arXiv:2306.09782*, 2023b.
- Chao Ma, Wenbo Gong, Meyer Scetbon, and Edward Meeds. Swan: Preprocessing sgd enables adam-level performance on llm training with significant memory reduction. *arXiv preprint arXiv:2412.13148*, 2024.
- James Martens. New insights and perspectives on the natural gradient method. *Journal of Machine Learning Research*, 21(146):1–76, 2020.
- James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International conference on machine learning*, pp. 2408–2417. PMLR, 2015.

- Ionut-Vlad Modoranu, Mher Safaryan, Grigory Malinovsky, Eldar Kurtic, Thomas Robert, Peter Richtárik, and Dan Alistarh. Microadam: Accurate adaptive optimization with low space overhead and provable convergence. *arXiv.org*, 2024. doi: 10.48550/arxiv.2405.15593.
- Depen Morwani, Itai Shapira, Nikhil Vyas, Eran Malach, Sham Kakade, and Lucas Janson. A new perspective on shampoo’s preconditioner. *arXiv preprint arXiv:2406.17748*, 2024.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Guilherme Penedo, Hynek Kydliček, Anton Lozhkov, Margaret Mitchell, Colin A Raffel, Leandro Von Werra, Thomas Wolf, et al. The fineweb datasets: Decanting the web for the finest text data at scale. *Advances in Neural Information Processing Systems*, 37:30811–30849, 2025.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–16. IEEE, 2020.
- Yehonathan Refael, Jonathan Svirsky, Boris Shustin, Wasim Huleihel, and Ofir Lindenbaum. Adarankgrad: Adaptive gradient-rank and moments for memory-efficient llms training and fine-tuning. *arXiv preprint arXiv:2410.17881*, 2024.
- Thomas Robert, Mher Safaryan, Ionut-Vlad Modoranu, and Dan Alistarh. Ldadam: Adaptive optimization from low-dimensional gradient statistics. *arXiv.org*, 2024.
- Meyer Scetbon, Chao Ma, Wenbo Gong, and Edward Meeds. Gradient multi-normalization for stateless and scalable llm training. *arXiv preprint arXiv:2502.06742*, 2025.
- Noam Shazeer and Mitchell Stern. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pp. 4596–4604. PMLR, 2018.
- T. Tieleman and G. Hinton. Lecture 6.5 - rmsprop: Divide the gradient by a running average of its recent magnitude. Coursera: Neural Networks for Machine Learning, 2012.
- Thijs Vogels, Sai Praneeth Karimireddy, and Martin Jaggi. Practical low-rank communication compression in decentralized deep learning. *Advances in Neural Information Processing Systems*, 33:14171–14181, 2020.
- Nikhil Vyas, Depen Morwani, Rosie Zhao, Mujin Kwun, Itai Shapira, David Brandfonbrener, Lucas Janson, and Sham Kakade. Soap: Improving and stabilizing shampoo using adam. *arXiv preprint arXiv:2409.11321*, 2024.
- Alex Wang. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- Yizhong Wang, Hamish Ivison, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Chandu, David Wadden, Kelsey MacMillan, Noah A Smith, Iz Beltagy, et al. How far can camels go? exploring the state of instruction tuning on open resources. *Advances in Neural Information Processing Systems*, 36:74764–74786, 2023.
- Greg Yang, James B Simon, and Jeremy Bernstein. A spectral condition for feature learning. *arXiv preprint arXiv:2310.17813*, 2023.
- Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*, 2021.
- Qingru Zhang, Minshuo Chen, Alexander Bukharin, Nikos Karampatziakis, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. Adalora: Adaptive budget allocation for parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.10512*, 2023.

Jiawei Zhao, Zhenyu Zhang, Beidi Chen, Zhangyang Wang, Anima Anandkumar, and Yuandong Tian. Galore: Memory-efficient llm training by gradient low-rank projection. *arXiv preprint arXiv:2403.03507*, 2024a.

Pengxiang Zhao, Ping Li, Yingjie Gu, Yi Zheng, Stephan Ludger Kölker, Zhefeng Wang, and Xiaoming Yuan. Adapprox: Adaptive approximation in adam optimization via randomized low-rank matrices. *arXiv preprint arXiv:2403.14958*, 2024b.

Hanqing Zhu, Zhenyu Zhang, Wenyan Cong, Xi Liu, Sem Park, Vikas Chandra, Bo Long, David Z Pan, Zhangyang Wang, and Jinwon Lee. Apollo: Sgd-like memory, adamw-level performance. *arXiv preprint arXiv:2412.05270*, 2024.

A Additional Related Works

Parameter-Efficient Fine-Tuning (PEFT). To reduce the memory associated with full-parameter updates, PEFT methods constrain updates to a small subset of parameters or add minimal trainable components. Adapters introduce small bottleneck layers into the model (Houlsby et al., 2019), while the popular LoRA technique (Hu et al., 2021) injects trainable low-rank matrices into existing weight layers, drastically reducing memory by training only these adapters (Hu et al., 2021). Variants include AdaLoRA for adaptive rank allocation (Zhang et al., 2023), VeRA using shared low-rank matrices (Kopiczko et al., 2023), and DoRA decomposing updates into magnitude/direction (Liu et al., 2024). Other PEFT methods tune only biases (BitFit (Zaken et al., 2021)) or learn input embeddings (prompt tuning (Li & Liang, 2021; Lester et al., 2021)). QLoRA (Dettmers et al., 2024) further combines LoRA with 4-bit weight quantization. While memory-efficient, PEFT methods generally keep the base model fixed and rely on the capacity of the trained low-rank components.

Second-Order and Preconditioning Methods.

These methods aim primarily to accelerate convergence and improve optimization performance by incorporating curvature information, offering potentially more effective descent directions compared to first-order methods (George et al., 2018; Duvvuri et al., 2024; Gupta et al., 2018). Exact second-order information (e.g., using the full Hessian in Newton’s method or the full gradient covariance in full-matrix Adagrad (Duchi et al., 2011; Duvvuri et al., 2024)) is computationally infeasible for large models due to the prohibitive cost ($\mathcal{O}(d^2)$ memory, $\mathcal{O}(d^3)$ computation) of storing and manipulating the required matrices (George et al., 2018; Duvvuri et al., 2024; Gupta et al., 2018). Therefore, practical methods rely on approximations to make harnessing second-order information tractable. Quasi-Newton methods like LBFGS (Liu & Nocedal, 1989; Duvvuri et al., 2024) build implicit Hessian approximations but can still be memory-intensive (Duvvuri et al., 2024).

A major family of approximations leverages Kronecker product factorizations. KFAC (Kronecker-Factored Approximate Curvature) (Martens & Grosse, 2015) approximates the Fisher Information Matrix (an approximation of the Hessian (Martens, 2020)) block-diagonally using Kronecker products specific to network layer types (Martens & Grosse, 2015; Grosse & Martens, 2016). Shampoo (Gupta et al., 2018), motivated by full-matrix Adagrad (Gupta et al., 2018; Anil et al., 2019), uses Kronecker products of gradient statistics ($\mathbf{G}\mathbf{G}^\top, \mathbf{G}^\top\mathbf{G}$) as a preconditioner (Gupta et al., 2018). While powerful, these factored approximations require computing matrix roots or inverses, which are computationally demanding and often amortized over multiple steps, potentially using stale curvature estimates (Gupta et al., 2018; George et al., 2018).

Refinements seek to improve accuracy or efficiency. EKFac (George et al., 2018) builds on KFAC by performing cheaper diagonal updates within the Kronecker-factored eigenbasis (KFE), yielding a provably better Fisher approximation (George et al., 2018). CASPR (Duvvuri et al., 2024) uses Kronecker sums, aiming for a potentially better Adagrad approximation than Shampoo’s Kronecker product (Duvvuri et al., 2024). SOAP (Vyas et al., 2024) runs AdamW within Shampoo’s eigenbasis, aiming for improved stability and fewer hyperparameters, especially when the basis update is infrequent (Vyas et al., 2024). These methods highlight the ongoing effort to balance the power of second-order information with computational feasibility (George et al., 2018; Vyas et al., 2024).

On-the-Fly and Partial Updates. These techniques reduce memory by avoiding large buffers or distributing states. LOMO (Lv et al., 2023b) applies updates immediately, AdaLOMO (Lv et al., 2023a) adds adaptivity, BAdam (Luo et al., 2024) uses weight blocks, and ZeRO (Rajbhandari et al., 2020) shards states in distributed settings.

B Additional Preliminaries

Adaptive Optimization Methods. Adaptive optimization methods have become essential for training deep neural networks. By adjusting learning rates on a per-parameter (or per-dimension) basis using information from past gradients and the loss curvature, they often lead to significantly faster convergence than vanilla stochastic gradient descent (SGD) in practice. This line of work began with Adagrad (Duchi et al.,

2011), which proposed leveraging the gradient preconditioner $\mathbf{P}_t = \sum_{i=1}^t \mathbf{g}_i \mathbf{g}_i^\top \in \mathbb{R}^{mn \times mn}$ and performing the update rule as $\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{P}_t^{-\frac{1}{2}} \mathbf{g}_t$, where in our notation $\mathbf{g}_i = \text{Vec}(\mathbf{G}_i) \in \mathbb{R}^{mn}$ and $\mathbf{w}_t = \text{Vec}(\mathbf{W}_t)$. Note that considering only the diagonal terms of \mathbf{P}_t and using the EMA of gradient covariances (e.g., $\mathbf{P}_t = \sum_{i=1}^t \beta_2^{t-i} \mathbf{g}_i \mathbf{g}_i^\top$) led to an early variant of RMSprop (Tieleman & Hinton, 2012), and further using EMA of the gradients themselves, $\hat{\mathbf{g}}_t = \sum_{i=1}^t \beta_1^{t-i} \mathbf{g}_i$, instead of the plain \mathbf{g}_t , yields the well-known Adam (Kingma & Ba, 2015). However, the mentioned methods are limited forms of second-order methods, where the preconditioner matrix is restricted to being diagonal. Considering either the original Adagrad (Duchi et al., 2011) perspective or second-order optimization methods such as Newton’s method, moving beyond diagonal preconditioners is a natural step toward achieving even faster convergence. However, storing a non-diagonal preconditioner is often not feasible due to the size of DNNs, requiring $O(m^2 n^2)$ memory. Shampoo (Gupta et al., 2018) proposed maintaining a Kronecker approximation $\mathbf{P}_t \sim \mathbf{R}_t \otimes \mathbf{L}_t$, where $\mathbf{L}_t \in \mathbb{R}^{m \times m}$ and $\mathbf{R}_t \in \mathbb{R}^{n \times n}$, thereby reducing the number of parameters required for storing the preconditioning matrix to $O(m^2 + n^2)$. Morwani et al. (2024) showed that Shampoo (Gupta et al., 2018) closely approximates the full Adagrad preconditioner, unified all aforementioned methods, and argued that the gradient covariance $\mathbf{g}_t \mathbf{g}_t^\top$ closely approximates the Gauss-Newton components of the Hessian at \mathbf{w}_t , thereby drawing an interesting connection between Adagrad (Duchi et al., 2011) and second-order optimization methods.

C Experimental Details

This section provides detailed configurations, hyperparameters, tuning procedures, and dataset information for the experiments presented in the main paper, aiming to ensure reproducibility.

C.1 General Implementation Details

- **Software:** Experiments were implemented using standard libraries for deep learning, including PyTorch, Hugging Face Transformers, and Accelerate. Specific library versions are detailed in the code repository.
- **Hardware:** All experiments were conducted on NVIDIA A100 GPUs. The number of GPUs may have varied slightly depending on the specific experimental setup (e.g., 4 GPUs for Tulu3 tuning).
- **Code:** The implementation for MoFaSGD is available at <https://github.com/AnonCode1/MFSGD>.
git.

C.2 Pre-training: NanoGPT Speedrun

Hyperparameters and Tuning. For this benchmark, Muon hyperparameters were kept at the tuned defaults provided by Jordan et al. (2024a). Learning rates for AdamW, GaLore, and MoFaSGD were tuned via grid search over $\{1e-4, 2e-4, 3e-4, 5e-4, 8e-4, 1e-3, 3e-3, 5e-3, 8e-3, 1e-2, 2e-2, 5e-2\}$. MoFaSGD’s momentum decay β was tuned over $\{0.5, 0.85, 0.90, 0.95\}$. GaLore’s SVD frequency was tuned over $\{10, 25, 75, 150, 300\}$. The best-performing hyperparameters based on final validation perplexity were selected.

Key hyperparameters selected for the NanoGPT pre-training experiments are summarized in Table 5.

C.3 NLU Fine-tuning: GLUE Benchmark

Hyperparameters and Tuning. The experimental setup, including hyperparameters for baseline optimizers (AdamW, GaLore, LoRA+AdamW), directly follows the configuration reported in Zhao et al. (2024a) for RoBERTa-Base on GLUE. For our method, MoFaSGD, learning rates were tuned via grid search over $\{1e-5, 2e-5, 5e-5, 1e-4, 5e-4\}$ for each task and rank combination ($r = 4$ and $r = 8$), and the batch size was fixed to 16 for all experiments. The MoFaSGD momentum decay β was also kept fixed at 0.95. The learning rate yielding the best validation accuracy on each specific task was selected. The final selected hyperparameters for MoFaSGD across the evaluated GLUE tasks are presented in Table 6.

Table 5: Final Selected Hyperparameters for NanoGPT Pre-training Experiment (0.73B tokens)

| Optimizer | Rank | Learning Rate (LR) | SVD Freq. | MoFaSGD β |
|---------------------|-----------|--|-----------|-----------------|
| Muon | - | 5e-2 | N/A | N/A |
| AdamW | - | 2e-3 | N/A | N/A |
| GaLore | $r = 16$ | 2e-2 | 150 | N/A |
| | $r = 32$ | 8e-3 | 75 | N/A |
| | $r = 128$ | 8e-3 | 75 | N/A |
| MoFaSGD (Ours) | $r = 16$ | 1e-3 | N/A | 0.85 |
| | $r = 32$ | 5e-4 | N/A | 0.85 |
| | $r = 128$ | 3e-4 | N/A | 0.85 |
| Batch Size (Tokens) | | 524,288 tokens | | |
| LR Schedule | | Stable then linear decay with cool-down of 0.4 | | |

Table 6: Final Selected MoFaSGD Hyperparameters for GLUE Tasks (RoBERTa-Base).

| Hyperparameter | Rank | MNLI | QQP | SST-2 | MRPC | COLA | QNLI | RTE |
|--------------------|------------|------|------|-------|------|------|------|------|
| Learning Rate (LR) | $r = 4$ | 1e-4 | 5e-5 | 5e-5 | 1e-4 | 2e-5 | 2e-5 | 5e-5 |
| | $r = 8$ | 5e-5 | 5e-5 | 2e-5 | 5e-5 | 1e-5 | 1e-5 | 5e-5 |
| MoFaSGD β | $r = 4, 8$ | | | | 0.95 | | | |
| Batch Size | $r = 4, 8$ | | | | 16 | | | |
| Epochs | $r = 4, 8$ | | | | 15 | | | |

C.4 Instruction Tuning: Tulu3

Hyperparameters and Tuning. The setup largely follows Lambert et al. (2024). Learning rates for all optimizers were selected from the grid $\{1e-5, 5e-5, 1e-4, 5e-4, 1e-3\}$. The final LR for each method was chosen based on final validation loss on a held-out subset of the tulu-3-sft mixture. We used 5% of the sampled dataset for validation. Final selected hyperparameters for the Tulu3 instruction tuning experiment are summarized in Table 7.

Table 7: Final Selected Hyperparameters for Tulu-3 Instruction Tuning

| Hyperparameter | GaLore | LoRA | MoFaSGD (Ours) |
|----------------------|--------|------|----------------|
| Learning Rate (LR) | 5e-5 | 5e-5 | 1e-4 |
| Rank (r) | 8 | 8 | 8 |
| LoRA Alpha | N/A | 16 | N/A |
| GaLore SVD Freq. | 200 | N/A | N/A |
| MoFaSGD β | N/A | N/A | 0.95 |
| Effective Batch Size | 128 | | |
| Epochs | 1 | | |

Momentum Spectral Analysis.

- **Methodology:** Analysis was performed on the Tulu3 instruction-tuning run using the default optimizer AdamW with a learning rate of $1e-5$ and a batch size of 128 for one epoch. We directly inspected the AdamW state buffer for the first-moment EMA (M_t) (`state['exp_avg']`) after each

backward pass periodically. SVD was performed on these buffers, and the energy ratio for rank r was computed as $\frac{\sum_{i=1}^r \sigma_{i, M_t}^2}{\|M_t\|_F^2}$.

C.5 Training Loss Curves

The training loss curves in Figures 8a and 8b illustrate MoFaSGD’s convergence behavior. In both the simpler GLUE fine-tuning and the complex Tulu3 instruction-tuning, MoFaSGD consistently achieves lower training loss compared to LoRA and GaLore. This suggests more effective optimization dynamics, potentially stemming from the synergistic effect of adaptive low-rank momentum factorization and the spectrally normalized updates, which might offer better preconditioning than the standard updates used in LoRA or the subspace accumulation in GaLore.

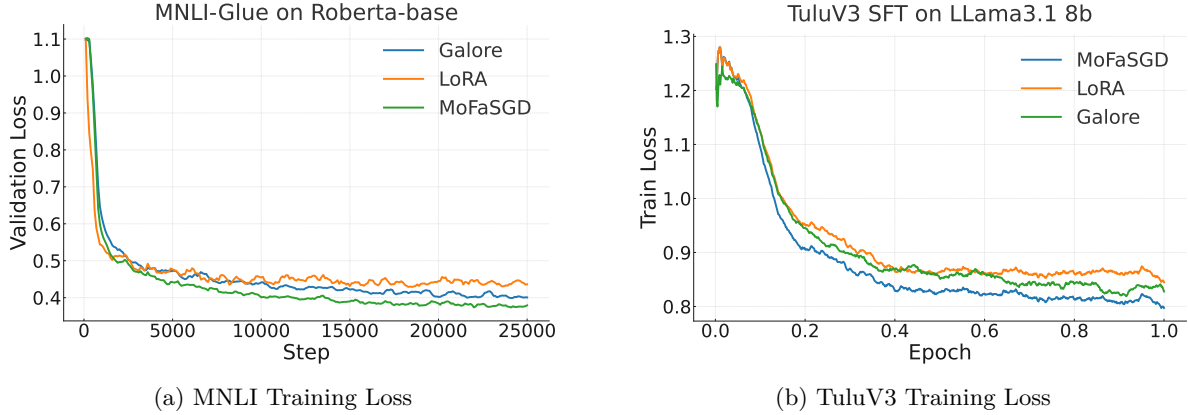


Figure 8: Training loss curves for post-training setups.

C.6 Memory Profiling Details

We provide full memory trace visualizations and a quantitative breakdown of memory usage across all optimizer configurations (LLaMA3.1-8B, BF16, no activation checkpointing, batch size 1, gradient accumulation 8). These results correspond to the experiments summarized in Figure 4 and Section 5.4.

Profiling Setup. Memory traces were collected using PyTorch’s native CUDA memory snapshot utility `torch.cuda.memory_snapshot()` at each training step. Each trace illustrates the evolution of GPU memory usage over four training steps. Manual annotations highlight key memory components.

Observations. As shown in the traces, MoFaSGD maintains a compact and stable memory profile, with minimal transient gradients and compressed optimizer state bands. AdamW and GaLore (non-fused) show persistent high memory usage due to full-rank gradient and optimizer state buffers. LoRA adds moderate adapter overhead, while fused GaLore shows improvements over its non-fused version.




Figure 9: MoFaSGD ($r = 8$): Compact optimizer states and minimal gradient spikes. Total memory ~ 29.4 GB.




Figure 10: LoRA ($r = 8$): Slightly higher activation and adapter memory, total ~ 33.6 GB.




Figure 11: SWAN: Stateless optimizer, but full-sized gradient buffers lead to higher memory usage (~ 43.9 GB).




Figure 12: AdamW (BF16): Full-rank moments and gradient buffers dominate memory (~ 70.8 GB).




Figure 13: GaLore (Fused, $r = 8$): Low gradient and optimizer state memory; total ~ 30.0 GB.

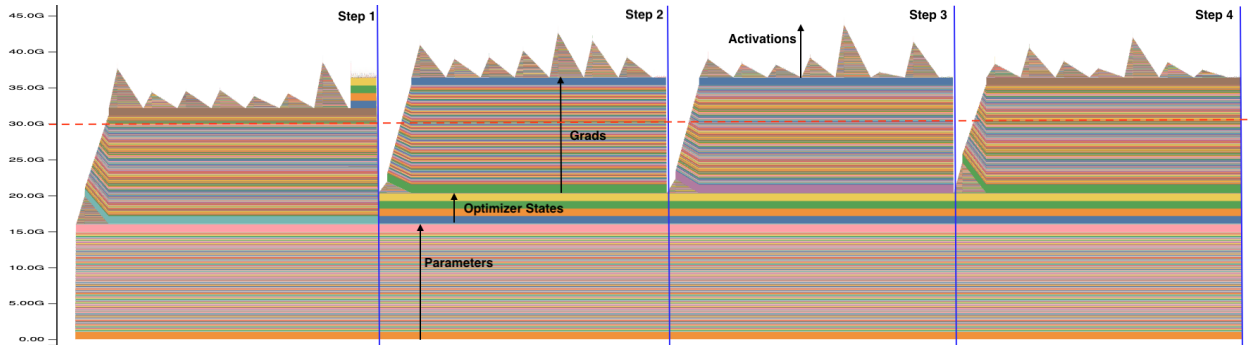


Figure 14: GaLore (Non-Fused, $r = 8$): Gradient accumulation inflates memory cost (~ 44.5 GB).

Quantitative Breakdown. The table below summarizes the memory footprint by category for all setups.

Table 8: Memory breakdown (in GB) for LLaMA3.1-8B fine-tuning.

| Optimizer | Params | Opt. States | Gradients | Activations | Adapters |
|------------------------------|--------|-------------|-----------|-------------|----------|
| MoFaSGD ($r = 8$) | 15.5 | 4.2 | 2.1 | 7.6 | 0 |
| LoRA ($r = 8$) | 15.5 | 4.2 | 2.1 | 9.8 | 2.0 |
| SWAN | 15.5 | 4.2 | 16.0 | 8.2 | 0 |
| AdamW (BF16) | 15.5 | 31.8 | 16.0 | 7.5 | 0 |
| GaLore Fused ($r = 8$) | 15.5 | 4.2 | 2.1 | 8.2 | 0 |
| GaLore Non-Fused ($r = 8$) | 15.5 | 4.2 | 16.0 | 8.8 | 0 |

D Proofs and Technical Details

D.1 On the Choice of Nuclear-norm Smoothness

One of the main assumptions behind our convergence bound in Theorem 4.5 is smoothness under the nuclear norm (Assumption 4.1). Here, we provide a detailed argument to justify this assumption. It is primarily motivated by concepts of descent in normed spaces, duality maps, and modular duality in deep learning—advancements detailed in recent works (Bernstein et al., 2023; Large et al., 2025; Bernstein & Newhouse, 2024a;b). These provide a solid foundation for spectral normalization of the descent update direction, which has been popularized in strong optimizers such as Muon (Jordan et al., 2024b; Liu et al., 2025).

First, we briefly discuss modular optimization theory, inspired by Large et al. (2025). The core idea of gradient descent is to optimize a locally linear upper bound of the loss function, as characterized by the gradient. Given a loss function $\mathcal{L} : \mathcal{W} \rightarrow \mathbb{R}$, we ideally want to find a tight upper bound on the higher-order terms, $\mathcal{R}(\mathbf{w}, \Delta\mathbf{w})$:

$$\mathcal{L}(\mathbf{w} + \Delta\mathbf{w}) \leq \mathcal{L}(\mathbf{w}) + \langle \nabla \mathcal{L}(\mathbf{w}), \Delta\mathbf{w} \rangle + \mathcal{R}(\mathbf{w}, \Delta\mathbf{w}) \quad (13)$$

For $\mathcal{W} \in \mathbb{R}^d$, classical second-order theory estimates the upper bound on higher-order terms using the maximum singular value of the Hessian. This leads to employing the ℓ_2 norm with the typical upper bound $\frac{L}{2} \|\Delta\mathbf{w}\|_2^2$, where L is the smoothness constant. However, this assumption often does not hold in deep learning practice and may yield a very loose upper bound. Modular theory (Large et al., 2025) aims to go beyond this by introducing an architecture-aware version of the upper bound based on the following conjecture:

$$\mathcal{L}(\mathbf{w} + \Delta\mathbf{w}) \leq \mathcal{L}(\mathbf{w}) + \langle \nabla \mathcal{L}(\mathbf{w}), \Delta\mathbf{w} \rangle + \frac{\lambda}{2} \|\Delta\mathbf{w}\|^2 \quad (\text{Normed Space Steepest Descent})$$

Here, \mathcal{W} can be any vector space, and $\|\cdot\| : \mathcal{W} \rightarrow \mathbb{R}$ is an arbitrary norm on it. Modular optimization theory aims to find appropriate norms on the weight space of all model parameters $\mathcal{W} = \mathcal{W}_1 \times \dots \times \mathcal{W}_L$. The modular norm on this parameter product space \mathcal{W} is defined as the scaled max over individual norms (e.g., $\max(s_1 \|\mathbf{w}_1\|_{\mathcal{W}_1}, \dots, \|\mathbf{w}_L\|_{\mathcal{W}_L})$). The specific choice of norms for each parameter is based on its role and properties. For instance, for a linear layer \mathbf{W} , we expect the normalized feature change not to be drastic, which is characterized by the induced operator norm defined as $\|\Delta\mathbf{W}\|_{\alpha \rightarrow \beta} = \max \frac{\|\mathbf{W}\mathbf{x}\|_\beta}{\|\mathbf{x}\|_\alpha}$. Empirical observations show that for typical gradient updates, the normalized feature change of a linear layer is indeed close to the RMS-RMS operator norm (Large et al., 2025).

Large et al. (2025) show that by assigning norms to atomic modules (e.g., RMS \rightarrow RMS for Linear layers and $\ell_1 \rightarrow$ RMS for embeddings), the modular norm of general architectures like Transformers satisfies the smoothness property. Informally, Proposition 5 of Large et al. (2025) shows that for a general module \mathbf{M} over $\mathcal{X}, \mathcal{Y}, \mathcal{W}$, and for common loss functions $\mathcal{L}(\mathbf{w}) = \mathbb{E}_{\mathbf{x}, \mathbf{y}}[\ell(\mathbf{M}(\mathbf{w}, \mathbf{x}))]$ such as cross-entropy, we have:

$$\|\nabla \mathcal{L}(\mathbf{w} + \Delta\mathbf{w}) - \nabla \mathcal{L}(\mathbf{w})\|_M^* \leq L_M \|\Delta\mathbf{w}\|_M \quad (14)$$

where $\|\cdot\|_M^*$ is the dual norm of the modular norm defined earlier. Since our method targets the linear layers of Transformers, and the dual norm of RMS \rightarrow RMS is the fan-in/fan-out scaled nuclear norm, we argue that assuming smoothness under the nuclear norm is more natural than under the Frobenius norm. The Frobenius norm assumption essentially treats the parameter space as \mathbb{R}^d , ignoring the matrix structure of linear layers and their architectural role.

However, note that our Assumption 4.1 is still stronger than Equation 14. Even if we use the RMS \rightarrow RMS operator norm for all linear layers and assume the remaining parameters are fixed, letting the weight space be $\mathcal{W}_{\text{linear}} = (\mathbf{W}_1, \dots, \mathbf{W}_{N_l})$, then by Equation 14, we would have:

$$\sum_{i=1}^{N_l} \|\nabla \mathcal{L}(\mathbf{W}_i + \Delta \mathbf{W}_i) - \nabla \mathcal{L}(\mathbf{W}_i)\|_* \leq L \max_{i=1}^{N_l} (\|\Delta \mathbf{W}_i\|_2),$$

where $\|\cdot\|_*$ denotes the nuclear norm, and $\|\cdot\|_2$ is the spectral norm, i.e., the $\ell_2 \rightarrow \ell_2$ operator norm.

D.2 Mathematical Tools and Notations

Kronecker Product and Vectorization. Here, we introduce the background regarding the definition of Kronecker product, Vectorization, and well-known properties of these operators that make it easier to work with them. The Kronecker product is denoted as \otimes , and for any arbitrary $\mathbf{X} \in \mathbb{R}^{m_1 \times n_1}$ and $\mathbf{Y} \in \mathbb{R}^{m_2 \times n_2}$, is defined as follows, where $x_{i,j}$ are the elements on row i and column j of the matrix \mathbf{X} :

$$\mathbf{X} \otimes \mathbf{Y} = \begin{bmatrix} x_{1,1}\mathbf{Y} & x_{1,2}\mathbf{Y} & \cdots & x_{1,n_1}\mathbf{Y} \\ x_{2,1}\mathbf{Y} & x_{2,2}\mathbf{Y} & \cdots & x_{2,n_1}\mathbf{Y} \\ \vdots & \vdots & \ddots & \vdots \\ x_{m_1,1}\mathbf{Y} & x_{m_1,2}\mathbf{Y} & \cdots & x_{m_1,n_1}\mathbf{Y} \end{bmatrix} \quad (15)$$

Moreover, we define the vectorization operator $\text{Vec}(\cdot)$ that stacks columns of the matrix as a vector. Particularly, we have:

$$\text{Vec}(\mathbf{X}) = \begin{bmatrix} x_{1,1} \\ x_{2,1} \\ \vdots \\ x_{m_1,1} \\ x_{1,2} \\ x_{2,2} \\ \vdots \\ x_{m_1,n_1} \end{bmatrix} \in \mathbb{R}^{m_1 n_1} \quad (16)$$

The following lemma summarizes basic properties of \otimes and $\text{Vec}(\cdot)$ that we leverage throughout our proofs in the paper. Since this lemma covers basic, well-known properties in Matrix Algebra, we refer the reader to Horn & Johnson (1994) for proofs and details.

Lemma D.1. *Let A , B , C and D be arbitrary matrices and assume that operations in each property are well-defined with respect to their dimensions. We have:*

1. *Scalar Multiplication:* $\mathbf{A} \otimes (\alpha \mathbf{B}) = \alpha(\mathbf{A} \otimes \mathbf{B}) = (\alpha \mathbf{A}) \otimes \mathbf{B}$
2. *Associativity:* $(\mathbf{A} \otimes \mathbf{B}) \otimes \mathbf{C} = \mathbf{A} \otimes (\mathbf{B} \otimes \mathbf{C})$
3. *Distributivity over addition:* $(\mathbf{A} + \mathbf{B}) \otimes \mathbf{C} = (\mathbf{A} \otimes \mathbf{C}) + (\mathbf{B} \otimes \mathbf{C})$, and $\mathbf{A} \otimes (\mathbf{B} + \mathbf{C}) = (\mathbf{A} \otimes \mathbf{B}) + (\mathbf{A} \otimes \mathbf{C})$
4. *Not necessarily commutative:* $(\mathbf{A} \otimes \mathbf{B}) \neq (\mathbf{B} \otimes \mathbf{A})$
5. *Transpose:* $(\mathbf{A} \otimes \mathbf{B})^\top = \mathbf{A}^\top \otimes \mathbf{B}^\top$

6. if \mathbf{A} and \mathbf{B} are positive-semidefinite, we have $(\mathbf{A} \otimes \mathbf{B})^s = \mathbf{A}^s \otimes \mathbf{B}^s$ for any positive real s , and if \mathbf{A} is positive definite, it holds for any real s .
7. Mixed-product property: $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = \mathbf{AC} \otimes \mathbf{BD}$
8. Outer product: $\mathbf{u} \otimes \mathbf{v} = \mathbf{uv}^\top$
9. Mixed Kronecker matrix-vector product: $(\mathbf{A}^\top \otimes \mathbf{B}) \text{Vec}(\mathbf{C}) = \text{Vec}(\mathbf{BCA})$

D.3 Proofs

Proof of Theorem 4.3. To prove Theorem 4.3, we first introduce the following lemma.

Lemma D.2. Consider the reduced SVD decomposition of the left and right sketches as follows $\mathbf{L} = \mathbf{U}_L \mathbf{\Sigma}_L \mathbf{V}_L^\top$ and $\mathbf{R} = \mathbf{U}_R \mathbf{\Sigma}_R \mathbf{V}_R^\top$. Then, the vectorized unsketched representation of any arbitrary matrix $\mathbf{G} \in \mathbb{R}^{m \times n}$ can be decomposed as follows:

$$\text{Vec}(\mathbf{LL}^\top \mathbf{G} + \mathbf{GRR}^\top) = \mathbf{U}_{L,R} \mathbf{\Sigma}_{L,R} \mathbf{U}_{L,R}^\top \text{Vec}(\mathbf{G}) \quad (17)$$

where $\mathbf{U}_{L,R} \in \mathbb{R}^{mn \times r(m+n-r)}$ is a semi-orthogonal matrix (i.e., $\mathbf{U}_{L,R}^\top \mathbf{U}_{L,R} = \mathbf{I}_{r(m+n-r)}$), and $\mathbf{\Sigma}_{L,R} \in \mathbb{R}^{(m+n-r) \times (m+n-r)}$ is a diagonal matrix, defined as below:

$$\mathbf{U}_{L,R} = [\mathbf{U}_R \otimes \mathbf{U}_L, \mathbf{U}_R \otimes \mathbf{U}_L^\perp, \mathbf{U}_R^\perp \otimes \mathbf{U}_L] \\ \mathbf{\Sigma}_{L,R} = \begin{bmatrix} \mathbf{\Sigma}_R^2 \otimes \mathbf{I}_r + \mathbf{I}_r \otimes \mathbf{\Sigma}_L^2 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma}_R^2 \otimes \mathbf{I}_{n-r} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_{m-r} \otimes \mathbf{\Sigma}_L^2 \end{bmatrix} \quad (18)$$

Proof of Lemma D.2. Recall the left singular vectors of the left and right sketch matrices, $\mathbf{U}_L \in \mathbb{R}^{m \times r}$ and $\mathbf{U}_R \in \mathbb{R}^{n \times r}$. We can augment these reduced bases with $m-r$ and $n-r$ orthogonal vectors to get full orthonormal bases as follows: $[\mathbf{U}_L, \mathbf{U}_L^\perp] \in \mathbb{R}^{m \times m}$ and $[\mathbf{U}_R, \mathbf{U}_R^\perp] \in \mathbb{R}^{n \times n}$. Since each of these expanded bases is square orthonormal matrices, we can write $\mathbf{U}_L^\top \mathbf{U}_L = \mathbf{I}_r$, $(\mathbf{U}_L^\perp)^\top \mathbf{U}_L^\perp = \mathbf{I}_{m-r}$, and also $\mathbf{U}_L^\top \mathbf{U}_L^\perp = \mathbf{0}$. Similar properties hold for \mathbf{U}_R . Moreover, again based on the orthonormality of the augmented bases, we can write:

$$\mathbf{U}_L \mathbf{U}_L^\top + \mathbf{U}_L^\perp (\mathbf{U}_L^\perp)^\top = \mathbf{I}_m \quad , \quad \mathbf{U}_R \mathbf{U}_R^\top + \mathbf{U}_R^\perp (\mathbf{U}_R^\perp)^\top = \mathbf{I}_n \quad (19)$$

Leveraging Lemma D.1 we can write:

$$\begin{aligned} \text{Vec}(\mathbf{LL}^\top \mathbf{G} + \mathbf{GRR}^\top) &= \text{Vec}(\mathbf{LL}^\top \mathbf{G}) + \text{Vec}(\mathbf{GRR}^\top) \\ &= (\mathbf{I}_n \otimes \mathbf{LL}^\top + \mathbf{RR}^\top \otimes \mathbf{I}_m) \text{Vec}(\mathbf{G}) \\ &= (\mathbf{I}_n \otimes \mathbf{U}_L \mathbf{\Sigma}_L^2 \mathbf{U}_L^\top + \mathbf{U}_R \mathbf{\Sigma}_R^2 \mathbf{U}_R^\top \otimes \mathbf{I}_m) \text{Vec}(\mathbf{G}) \end{aligned} \quad (20)$$

Using Equation 20 and properties in Lemma D.1, we can further write:

$$\begin{aligned} \mathbf{I}_n \otimes \mathbf{U}_L \mathbf{\Sigma}_L^2 \mathbf{U}_L^\top &= (\mathbf{U}_R \mathbf{U}_R^\top + \mathbf{U}_R^\perp (\mathbf{U}_R^\perp)^\top) \otimes (\mathbf{U}_L \mathbf{\Sigma}_L^2 \mathbf{U}_L^\top) \\ &= \mathbf{U}_R \mathbf{U}_R^\top \otimes \mathbf{U}_L \mathbf{\Sigma}_L^2 \mathbf{U}_L^\top + \mathbf{U}_R^\perp (\mathbf{U}_R^\perp)^\top \otimes \mathbf{U}_L \mathbf{\Sigma}_L^2 \mathbf{U}_L^\top \\ &= (\mathbf{U}_R \otimes \mathbf{U}_L \mathbf{\Sigma}_L) (\mathbf{U}_R \otimes \mathbf{U}_L \mathbf{\Sigma}_L)^\top + (\mathbf{U}_R^\perp \otimes \mathbf{U}_L \mathbf{\Sigma}_L) (\mathbf{U}_R^\perp \otimes \mathbf{U}_L \mathbf{\Sigma}_L)^\top \\ &= (\mathbf{U}_R \otimes \mathbf{U}_L) (\mathbf{I}_r \otimes \mathbf{\Sigma}_L^2) (\mathbf{U}_R \otimes \mathbf{U}_L)^\top + (\mathbf{U}_R^\perp \otimes \mathbf{U}_L) (\mathbf{I}_{n-r} \otimes \mathbf{\Sigma}_L^2) (\mathbf{U}_R^\perp \otimes \mathbf{U}_L)^\top \end{aligned} \quad (21)$$

Using a similar derivation as in Equation 21, we obtain:

$$\mathbf{U}_R \mathbf{\Sigma}_R^2 \mathbf{U}_R^\top \otimes \mathbf{I}_m = (\mathbf{U}_R \otimes \mathbf{U}_L) (\mathbf{\Sigma}_R^2 \otimes \mathbf{I}_r) (\mathbf{U}_R \otimes \mathbf{U}_L)^\top + (\mathbf{U}_R \otimes \mathbf{U}_L^\perp) (\mathbf{\Sigma}_R^2 \otimes \mathbf{I}_{m-r}) (\mathbf{U}_R \otimes \mathbf{U}_L^\perp)^\top \quad (22)$$

Plugging Equation 21 and Equation 22 back into Equation 20, we get:

$$\begin{aligned}
\text{Vec}(\mathbf{L}\mathbf{L}^\top \mathbf{G} + \mathbf{G}\mathbf{R}\mathbf{R}^\top) &= \left[(\mathbf{U}_R \otimes \mathbf{U}_L)(\mathbf{I}_r \otimes \Sigma_L^2)(\mathbf{U}_R \otimes \mathbf{U}_L)^\top + (\mathbf{U}_R^\perp \otimes \mathbf{U}_L)(\mathbf{I}_{n-r} \otimes \Sigma_L^2)(\mathbf{U}_R^\perp \otimes \mathbf{U}_L)^\top \right. \\
&\quad + (\mathbf{U}_R \otimes \mathbf{U}_L)(\Sigma_R^2 \otimes \mathbf{I}_r)(\mathbf{U}_R \otimes \mathbf{U}_L)^\top \\
&\quad \left. + (\mathbf{U}_R \otimes \mathbf{U}_L^\perp)(\Sigma_R^2 \otimes \mathbf{I}_{m-r})(\mathbf{U}_R \otimes \mathbf{U}_L^\perp)^\top \right] \text{Vec}(\mathbf{G}) \\
&= \left[(\mathbf{U}_R \otimes \mathbf{U}_L)(\mathbf{I}_r \otimes \Sigma_L^2 + \Sigma_R^2 \otimes \mathbf{I}_r)(\mathbf{U}_R \otimes \mathbf{U}_L)^\top \right. \\
&\quad + (\mathbf{U}_R^\perp \otimes \mathbf{U}_L)(\mathbf{I}_{n-r} \otimes \Sigma_L^2)(\mathbf{U}_R^\perp \otimes \mathbf{U}_L)^\top \\
&\quad \left. + (\mathbf{U}_R \otimes \mathbf{U}_L^\perp)(\Sigma_R^2 \otimes \mathbf{I}_{m-r})(\mathbf{U}_R \otimes \mathbf{U}_L^\perp)^\top \right] \text{Vec}(\mathbf{G})
\end{aligned} \tag{23}$$

To complete the proof, it remains to show that the block matrix

$$\mathbf{U}_{L,R} = [\mathbf{U}_R \otimes \mathbf{U}_L, \mathbf{U}_R^\perp \otimes \mathbf{U}_L, \mathbf{U}_R \otimes \mathbf{U}_L^\perp] \in \mathbb{R}^{mn \times r(m+n-r)} \tag{24}$$

is semi-orthogonal, i.e. $\mathbf{U}_{L,R}^\top \mathbf{U}_{L,R} = \mathbf{I}_{r(m+n-r)}$.

First, note that pairs of sub-blocks are mutually orthogonal. For instance,

$$(\mathbf{U}_R \otimes \mathbf{U}_L)^\top (\mathbf{U}_R^\perp \otimes \mathbf{U}_L) = (\mathbf{U}_R^\top \mathbf{U}_R^\perp) \otimes (\mathbf{U}_L^\top \mathbf{U}_L) = \mathbf{0}.$$

A similar argument applies to all other sub-block pairs.

Moreover, each sub-block is orthonormal in its own right. For example,

$$(\mathbf{U}_R \otimes \mathbf{U}_L)^\top (\mathbf{U}_R \otimes \mathbf{U}_L) = (\mathbf{U}_R^\top \mathbf{U}_R) \otimes (\mathbf{U}_L^\top \mathbf{U}_L) = \mathbf{I}_r \otimes \mathbf{I}_r = \mathbf{I}_{r^2}.$$

Hence, $\mathbf{U}_{L,R}$ is a semi-orthogonal, and the result follows. \square

We turn to proving Theorem 4.3.

Let the reconstructed gradient after projection be $\tilde{\mathbf{G}} = \text{Proj}_{(\mathbf{L}, \mathbf{R})}(\mathbf{G}) = \alpha_1 \mathbf{L}\mathbf{L}^\top \mathbf{G} + \alpha_2 \mathbf{G}\mathbf{R}\mathbf{R}^\top + \alpha_3 \mathbf{L}\mathbf{L}^\top \mathbf{G}\mathbf{R}\mathbf{R}^\top$. We can write:

$$\text{Vec}(\tilde{\mathbf{G}}) = \text{Vec}(\alpha_1 \mathbf{L}\mathbf{L}^\top \mathbf{G} + \alpha_2 \mathbf{G}\mathbf{R}\mathbf{R}^\top) + \alpha_3 (\mathbf{R}\mathbf{R}^\top \otimes \mathbf{L}\mathbf{L}^\top) \text{Vec}(\mathbf{G}) \tag{25}$$

Replacing \mathbf{L} and \mathbf{R} with their corresponding reduced SVD decompositions, we have:

$$(\mathbf{R}\mathbf{R}^\top \otimes \mathbf{L}\mathbf{L}^\top) = \mathbf{U}_R \Sigma_R^2 \mathbf{U}_R^\top \otimes \mathbf{U}_L \Sigma_L^2 \mathbf{U}_L^\top = (\mathbf{U}_R \otimes \mathbf{U}_L)(\Sigma_R^2 \otimes \Sigma_L^2)(\mathbf{U}_R \otimes \mathbf{U}_L)^\top \tag{26}$$

Using Equation 23 from Lemma D.2 and Equation 26, we can rewrite:

$$\begin{aligned}
\text{Vec}(\tilde{\mathbf{G}}) &= \left[(\mathbf{U}_R \otimes \mathbf{U}_L)(\alpha_1 \mathbf{I}_r \otimes \Sigma_L^2 + \alpha_2 \Sigma_R^2 \otimes \mathbf{I}_r + \alpha_3 \Sigma_R^2 \otimes \Sigma_L^2)(\mathbf{U}_R \otimes \mathbf{U}_L)^\top \right. \\
&\quad + \alpha_2 (\mathbf{U}_R^\perp \otimes \mathbf{U}_L)(\mathbf{I}_{n-r} \otimes \Sigma_L^2)(\mathbf{U}_R^\perp \otimes \mathbf{U}_L)^\top \\
&\quad \left. + \alpha_1 (\mathbf{U}_R \otimes \mathbf{U}_L^\perp)(\Sigma_R^2 \otimes \mathbf{I}_{m-r})(\mathbf{U}_R \otimes \mathbf{U}_L^\perp)^\top \right] \text{Vec}(\mathbf{G}) \\
&\triangleq \mathbf{P}_{L,R} \text{Vec}(\mathbf{G})
\end{aligned} \tag{27}$$

where in the last equality, we define $\mathbf{P}_{L,R} \in \mathbb{R}^{mn \times mn}$ as the corresponding projection matrix of $\text{Proj}_{(\mathbf{L}, \mathbf{R})}$.

We can use Equation 19 to write:

$$\begin{aligned}
\mathbf{I}_{mn} = \mathbf{I}_n \otimes \mathbf{I}_m &= (\mathbf{U}_R \otimes \mathbf{U}_L)(\mathbf{U}_R \otimes \mathbf{U}_L)^\top + (\mathbf{U}_R^\perp \otimes \mathbf{U}_L)(\mathbf{U}_R^\perp \otimes \mathbf{U}_L)^\top + (\mathbf{U}_R \otimes \mathbf{U}_L^\perp)(\mathbf{U}_R \otimes \mathbf{U}_L^\perp)^\top \\
&\quad + (\mathbf{U}_R^\perp \otimes \mathbf{U}_L^\perp)(\mathbf{U}_R^\perp \otimes \mathbf{U}_L^\perp)^\top
\end{aligned} \tag{28}$$

We now propose our Kronecker decomposition of the subspace projection residual.

$$\begin{aligned} \text{Vec}(\tilde{\mathbf{G}} - \mathbf{G}) &= (\mathbf{I}_n \otimes \mathbf{I}_m - \mathbf{P}_{\mathbf{L}, \mathbf{R}}) \text{Vec}(\mathbf{G}) = \\ &\left[(\mathbf{U}_R \otimes \mathbf{U}_L) (\alpha_1 \mathbf{I}_r \otimes \Sigma_L^2 + \alpha_2 \Sigma_R^2 \otimes \mathbf{I}_r + \alpha_3 \Sigma_R^2 \otimes \Sigma_L^2 - \mathbf{I}_r \otimes \mathbf{I}_r) (\mathbf{U}_R \otimes \mathbf{U}_L)^\top \right. \\ &\quad + (\mathbf{U}_R^\perp \otimes \mathbf{U}_L) (\alpha_2 \mathbf{I}_{n-r} \otimes \Sigma_L^2 - \mathbf{I}_{n-r} \otimes \mathbf{I}_r) (\mathbf{U}_R^\perp \otimes \mathbf{U}_L)^\top \\ &\quad + (\mathbf{U}_R \otimes \mathbf{U}_L^\perp) (\alpha_1 \Sigma_R^2 \otimes \mathbf{I}_{m-r} - \mathbf{I}_r \otimes \mathbf{I}_{m-r}) (\mathbf{U}_R \otimes \mathbf{U}_L^\perp)^\top \\ &\quad \left. + (\mathbf{U}_R^\perp \otimes \mathbf{U}_L^\perp) (\mathbf{U}_R^\perp \otimes \mathbf{U}_L^\perp)^\top \right] \text{Vec}(\mathbf{G}). \end{aligned} \quad (29)$$

For simplicity let $\mathbf{A}_1 = (\mathbf{U}_R \otimes \mathbf{U}_L) (\alpha_1 \mathbf{I}_r \otimes \Sigma_L^2 + \alpha_2 \Sigma_R^2 \otimes \mathbf{I}_r + \alpha_3 \Sigma_R^2 \otimes \Sigma_L^2 - \mathbf{I}_r \otimes \mathbf{I}_r) (\mathbf{U}_R \otimes \mathbf{U}_L)^\top$, $\mathbf{A}_2 = (\mathbf{U}_R^\perp \otimes \mathbf{U}_L) (\alpha_2 \mathbf{I}_{n-r} \otimes \Sigma_L^2 - \mathbf{I}_{n-r} \otimes \mathbf{I}_r) (\mathbf{U}_R^\perp \otimes \mathbf{U}_L)^\top$, $\mathbf{A}_3 = (\mathbf{U}_R \otimes \mathbf{U}_L^\perp) (\alpha_1 \Sigma_R^2 \otimes \mathbf{I}_{m-r} - \mathbf{I}_r \otimes \mathbf{I}_{m-r}) (\mathbf{U}_R \otimes \mathbf{U}_L^\perp)^\top$ and $\mathbf{A}_4 = (\mathbf{U}_R^\perp \otimes \mathbf{U}_L^\perp) (\mathbf{U}_R^\perp \otimes \mathbf{U}_L^\perp)^\top$. First note that $\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3, \mathbf{A}_4 \in \mathbb{R}^{mn \times mn}$. For all distinct $(i, j) \in \{1, 2, 3, 4\}$, their matrix product is zero, $\mathbf{A}_i^\top \mathbf{A}_j = \mathbf{0}$. Without loss of generality, we show it for \mathbf{A}_2 , and \mathbf{A}_3 here.

$$\begin{aligned} \mathbf{A}_2^\top \mathbf{A}_3 &= (\mathbf{U}_R^\perp \otimes \mathbf{U}_L) (\alpha_2 \mathbf{I}_{n-r} \otimes \Sigma_L^2 - \mathbf{I}_{n-r} \otimes \mathbf{I}_r) (\mathbf{U}_R^\perp \otimes \mathbf{U}_L)^\top (\mathbf{U}_R \otimes \mathbf{U}_L^\perp) (\alpha_3 \Sigma_R^2 \otimes \mathbf{I}_{m-r} - \mathbf{I}_r \otimes \mathbf{I}_{m-r}) \\ &\quad (\mathbf{U}_R \otimes \mathbf{U}_L^\perp)^\top \\ &= (\mathbf{U}_R^\perp \otimes \mathbf{U}_L) (\alpha_2 \mathbf{I}_{n-r} \otimes \Sigma_L^2 - \mathbf{I}_{n-r} \otimes \mathbf{I}_r) (\mathbf{U}_R^{\perp\top} \mathbf{U}_R \otimes \mathbf{U}_L^\top \mathbf{U}_L^\perp)^\top (\alpha_3 \Sigma_R^2 \otimes \mathbf{I}_{m-r} - \mathbf{I}_r \otimes \mathbf{I}_{m-r}) \\ &\quad (\mathbf{U}_R \otimes \mathbf{U}_L^\perp)^\top \\ &= (\mathbf{U}_R^\perp \otimes \mathbf{U}_L) (\alpha_2 \mathbf{I}_{n-r} \otimes \Sigma_L^2 - \mathbf{I}_{n-r} \otimes \mathbf{I}_r) (\mathbf{0}_{n-r, r} \otimes \mathbf{0}_{n-r, r})^\top (\alpha_3 \Sigma_R^2 \otimes \mathbf{I}_{m-r} - \mathbf{I}_r \otimes \mathbf{I}_{m-r}) \\ &\quad (\mathbf{U}_R \otimes \mathbf{U}_L^\perp)^\top \\ &= \mathbf{0}_{mn} \end{aligned} \quad (30)$$

Using Equation 29, and above fact yields:

$$\|\tilde{\mathbf{G}} - \mathbf{G}\|_F^2 = \|\text{Vec}(\tilde{\mathbf{G}} - \mathbf{G})\|_2^2 = \|\mathbf{A}_1 \mathbf{g}\|_2^2 + \|\mathbf{A}_2 \mathbf{g}\|_2^2 + \|\mathbf{A}_3 \mathbf{g}\|_2^2 + \|\mathbf{A}_4 \mathbf{g}\|_2^2 \quad (31)$$

where $\|\cdot\|_2$ here is ℓ_2 norm of a vector, and $\mathbf{g} = \text{Vec}(\mathbf{G})$. The equality holds because $(\mathbf{A}_i \mathbf{g})^\top (\mathbf{A}_j \mathbf{g}) = \mathbf{0}_{mn}$. Based on the above equation, we have the following lower bound on residual $\|\tilde{\mathbf{G}} - \mathbf{G}\|_F^2 \geq \|\mathbf{A}_4 \mathbf{g}\|_2^2$ for any given \mathbf{L} , and \mathbf{R} . The term \mathbf{A}_4 does not depend on the choice of $(\alpha_1, \alpha_2, \alpha_3)$, also Σ_R and Σ_L . Thus, the lower bound on this residual would be tight if and only if we have $\|\mathbf{A}_1 \mathbf{g}\|_2 = \|\mathbf{A}_2 \mathbf{g}\|_2 = \|\mathbf{A}_3 \mathbf{g}\|_2 = 0$. By setting $\Sigma_L = \mathbf{I}_r$, $\Sigma_R = \mathbf{I}_r$, and then considering $(\alpha_1, \alpha_2, \alpha_3) = (1, 1, -1)$, we have $\mathbf{A}_1 = \mathbf{A}_2 = \mathbf{A}_3 = \mathbf{0}_{mn}$.

Thus, we have shown that under conditions in Theorem 4.3, the first three terms in Equation 31 will be exactly zero, leaving the following term as the projection residual:

$$\text{Vec}(\tilde{\mathbf{G}}_{\text{optimal}} - \mathbf{G}) = (\mathbf{U}_R^\perp \otimes \mathbf{U}_L^\perp) (\mathbf{U}_R^\perp \otimes \mathbf{U}_L^\perp)^\top \text{Vec}(\mathbf{G}). \quad (32)$$

□

D.3.1 Convergence Proofs

In this section, we aim to find an upper bound on the iteration complexity of Algorithm 1 for finding an ϵ -stationary point defined by the averaged nuclear norm of gradients. First, we recall the necessary assumptions and notations.

Update Rule. Let $\{\mathbf{W}_t\}_{t=1}^T$ be the iterates of Algorithm 1. Then we have the following equivalent update rule:

$$\mathbf{W}_{t+1} = \mathbf{W}_t - \eta \mathbf{U}_{t+1} \mathbf{V}_{t+1}^\top, \quad \hat{\mathbf{M}}_t = \mathbf{U}_{t+1} \Sigma_{t+1} \mathbf{V}_{t+1}^\top, \quad \mathbf{M}_t = \sum_{i=0}^t \beta^{t-i} \mathbf{G}_i \quad (33)$$

where \mathbf{M}_t represents the full-rank momentum, and $\mathbf{G}_t = \nabla \mathcal{L}(\mathbf{W}_t, \xi_t) \in \mathbb{R}^{m \times n}$ is the stochastic gradient at parameter \mathbf{W}_t . Moreover, let $\hat{\mathbf{M}}_t$ be the low-rank momentum factorization, where $\mathbf{U}_i \in \mathbb{R}^{m \times r}$, $\mathbf{V}_i \in \mathbb{R}^{n \times r}$, and $\Sigma_i \in \mathbb{R}^{r \times r}$.

Definitions. We recall the necessary definitions. We let the optimization objective be represented as $\mathcal{L}(\mathbf{W}) = \mathbb{E}_\xi[\mathcal{L}(\mathbf{W}, \xi)]$, and moreover assume we have access to an unbiased, variance-bounded (σ) gradient oracle. For any optimization iterate \mathbf{W}_i , the full-batch gradient is denoted as $\bar{\mathbf{G}}_i = \nabla \mathcal{L}(\mathbf{W}_i)$, and the stochastic gradient is denoted as $\mathbf{G}_i = \nabla \mathcal{L}(\mathbf{W}_i, \xi_i)$. Based on assumptions on the gradient oracle, we have $\bar{\mathbf{G}}_i = \mathbb{E}[\mathbf{G}_i]$, and $\mathbb{E}[\|\mathbf{G}_i - \bar{\mathbf{G}}_i\|_*] \leq \sigma$. Moreover, we say a function $\mathcal{L}(\cdot) : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is L -smooth with respect to an arbitrary norm $\|\cdot\|$ if for any two parameters $\mathbf{W}_1, \mathbf{W}_2$, we have $\|\nabla \mathcal{L}(\mathbf{W}_1) - \nabla \mathcal{L}(\mathbf{W}_2)\|_* \leq L\|\mathbf{W}_1 - \mathbf{W}_2\|_2$.

Next, we propose our descent lemma.

Lemma D.3 (Descent lemma). *Let $\{\mathbf{W}\}_{t=0}^T$ be the iterates of Algorithm 1, optimized under a loss function that satisfies Assumption 4.1. Then, we have:*

$$\mathcal{L}(\mathbf{W}_{t+1}) \leq \mathcal{L}(\mathbf{W}_t) - \eta \|\bar{\mathbf{G}}_t\|_* + 2\eta \|\hat{\mathbf{M}}_t - \bar{\mathbf{G}}_t\|_* + \frac{\eta^2 L}{2} \quad (34)$$

Proof of Lemma D.3. First note that $\mathcal{L}(\cdot)$ is L -smooth with the nuclear norm. Therefore for any $\mathbf{W}_1, \mathbf{W}_2 \in \mathbb{R}^{m \times n}$ we have: $\|\nabla \mathcal{L}(\mathbf{W}_1) - \nabla \mathcal{L}(\mathbf{W}_2)\|_* \leq L\|\mathbf{W}_1 - \mathbf{W}_2\|_2$, where $\|\cdot\|_2$ is the spectral norm. Due to this property, we can leverage Proposition 5 in Large et al. (2025) to argue that for any $\mathbf{W}_1, \mathbf{W}_2$ we have:

$$\mathcal{L}(\mathbf{W}_2) \leq \mathcal{L}(\mathbf{W}_1) + \langle \nabla \mathcal{L}(\mathbf{W}_1), \mathbf{W}_2 - \mathbf{W}_1 \rangle_F + \frac{L}{2} \|\mathbf{W}_1 - \mathbf{W}_2\|_2^2 \quad (35)$$

Thus, considering the update rule as in Equation 33, for two consecutive iterates \mathbf{W}_t and \mathbf{W}_{t+1} , we can write:

$$\mathcal{L}(\mathbf{W}_{t+1}) \leq \mathcal{L}(\mathbf{W}_t) - \eta \langle \bar{\mathbf{G}}_t, \mathbf{U}_{t+1} \mathbf{V}_{t+1}^\top \rangle + \frac{L\eta^2}{2} \|\mathbf{U}_{t+1} \mathbf{V}_{t+1}^\top\|_2^2 \quad (36)$$

Note that if we take $\|\cdot\|_*$ as a function, its subgradient is well-known and can be derived as follows: the subgradient set at $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^\top$ is $\partial\|\mathbf{X}\|_* = \{\mathbf{U}\mathbf{V}^\top + \mathbf{H} : \mathbf{U}^\top \mathbf{H} = 0, \mathbf{V}^\top \mathbf{H} = 0, \|\mathbf{H}\|_2 \leq 1\}$. Since the nuclear norm is a convex function, for any $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^\top$ and \mathbf{Y} , we have:

$$\|\mathbf{Y}\|_* \geq \|\mathbf{X}\|_* + \langle \mathbf{U}\mathbf{V}^\top, \mathbf{Y} - \mathbf{X} \rangle \quad (37)$$

Replace $\mathbf{Y} = \hat{\mathbf{M}}_t - \bar{\mathbf{G}}_t$ and $\mathbf{X} = \hat{\mathbf{M}}_t = \mathbf{U}_{t+1} \Sigma_{t+1} \mathbf{V}_{t+1}^\top$; then we can rewrite Equation 36 as:

$$\mathcal{L}(\mathbf{W}_{t+1}) \leq \mathcal{L}(\mathbf{W}_t) - \eta \|\hat{\mathbf{M}}_t\|_* + \eta \|\hat{\mathbf{M}}_t - \bar{\mathbf{G}}_t\|_* + \frac{\eta^2}{2} L \quad (38)$$

where we also used the fact that $\|\mathbf{U}\mathbf{V}^\top\|_2 = 1$. Using the triangle inequality, we have $\|\hat{\mathbf{M}}_t\|_* \geq \|\bar{\mathbf{G}}_t\|_* - \|\hat{\mathbf{M}}_t - \bar{\mathbf{G}}_t\|_*$, and plugging this into Equation 38 completes the proof. \square

As indicated in Lemma D.3, the term $\|\hat{\mathbf{M}}_t - \bar{\mathbf{G}}_t\|_*$ is the main challenge for deriving the upper bound. Note that we can write $\|\hat{\mathbf{M}}_t - \bar{\mathbf{G}}_t\|_* \leq \|\mathbf{M}_t - \bar{\mathbf{G}}_t\|_* + \|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_*$. Thus, we aim to bound each decomposed term separately. Note that the first decomposed term $\|\mathbf{M}_t - \bar{\mathbf{G}}_t\|_*$ measures the nuclear norm difference between the full-batched gradient and the first momentum. Moreover, the second decomposed term $\|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_*$ measures the low-rank compression error of the full-rank momentum.

Lemma D.4. *Under Assumptions 4.1 and 4.1 we have:*

$$\sum_{t=1}^T \mathbb{E}[\|\mathbf{M}_t - \bar{\mathbf{G}}_t\|_*] \leq \frac{\beta}{1-\beta} \sum_{t=0}^{T-1} \mathbb{E}[\|\bar{\mathbf{G}}_t\|_*] + \frac{T\sigma}{(1-\beta)\sqrt{B}} \quad (39)$$

Proof of Lemma D.4. Let $\Delta_t = \mathbf{G}_t - \bar{\mathbf{G}}_t$. We can write $\|\mathbf{M}_t - \bar{\mathbf{G}}_t\|_* \leq \|\mathbf{M}_t - \mathbf{G}_t\|_* + \|\Delta_t\|_* = \beta\|\mathbf{M}_{t-1}\|_* + \|\Delta_t\|_*$. Moreover, we have $\|\mathbf{M}_t\|_* \leq \|\mathbf{G}_t\|_* + \beta\|\mathbf{M}_{t-1}\|_*$, which, upon unrolling, gives

$\|\mathbf{M}_t\|_* \leq \sum_{i=0}^t \beta^{t-i} \|\mathbf{G}_i\|_*$. Thus, we can write $\|\mathbf{M}_t - \mathbf{G}_t\|_* \leq \beta \sum_{i=0}^{t-1} \beta^{t-1-i} \|\mathbf{G}_i\|_*$. Finally, since $\sum_{i=0}^j \beta^i = \frac{1-\beta^{j+1}}{1-\beta} \leq \frac{1}{1-\beta}$, we have:

$$\sum_{t=1}^T \|\mathbf{M}_t - \bar{\mathbf{G}}_t\|_* \leq \frac{\beta}{1-\beta} \sum_{t=0}^{T-1} \|\bar{\mathbf{G}}_t\|_* + \frac{1}{1-\beta} \sum_{t=0}^T \|\Delta_t\|_* \quad (40)$$

Taking the expectation and leveraging Assumption 4.1, we have:

$$\sum_{t=1}^T \mathbb{E}[\|\mathbf{M}_t - \bar{\mathbf{G}}_t\|_*] \leq \frac{\beta}{1-\beta} \sum_{t=0}^{T-1} \mathbb{E}[\|\bar{\mathbf{G}}_t\|_*] + \frac{T\sigma}{(1-\beta)\sqrt{B}} \quad (41)$$

□

We now bound the second term $\|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_*$. We have the following lemma:

Lemma D.5. *Let $\|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_*$ be the compression error of the low-rank momentum estimation at iteration t . Then under Assumption 4.1 and the condition $\text{rank}(\mathbf{G}_0) \leq r$, we have the following bound:*

$$\sum_{t=1}^T \mathbb{E}[\|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_*] \leq \frac{\eta LT}{1-\beta} + \frac{2\sigma T}{\sqrt{B}(1-\beta)} \quad (42)$$

Proof of Lemma D.5. First note that we have:

$$\|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_* = \|\beta \mathbf{M}_{t-1} + \mathbf{G}_t - \hat{\mathbf{M}}_t\|_* \leq \beta \|\hat{\mathbf{M}}_{t-1} - \mathbf{M}_{t-1}\|_* + \|\mathbf{G}_t + \beta \hat{\mathbf{M}}_{t-1} - \hat{\mathbf{M}}_t\|_* \quad (43)$$

Considering the momentum factor update rule in 1, we can replace $\hat{\mathbf{M}}_t$ with $\hat{\mathbf{G}}_t + \beta \hat{\mathbf{M}}_{t-1}$, where $\hat{\mathbf{G}}_t = \mathbf{U}_t \mathbf{U}_t^\top \mathbf{G}_t + \mathbf{G}_t \mathbf{V}_t \mathbf{V}_t^\top - \mathbf{U}_t \mathbf{U}_t^\top \mathbf{G}_t \mathbf{V}_t \mathbf{V}_t^\top$. Therefore, we can write:

$$\|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_* \leq \beta \|\hat{\mathbf{M}}_{t-1} - \mathbf{M}_{t-1}\|_* + \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top) \mathbf{G}_t (\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* \quad (44)$$

The term $\|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top) \mathbf{G}_t (\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_*$ represents the low-rank compression error of gradients happening in our Algorithm 1. Thus, we aim to bound this term. We have:

$$\begin{aligned} \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top) \mathbf{G}_t (\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* &\leq \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top)(\mathbf{G}_t - \mathbf{G}_{t-1})(\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* + \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top) \mathbf{G}_{t-1} (\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* \\ &\leq \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top)(\bar{\mathbf{G}}_t - \bar{\mathbf{G}}_{t-1})(\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* + \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top) \mathbf{G}_{t-1} (\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* + \|\Delta_t - \Delta_{t-1}\|_* \\ &\leq L \|\mathbf{W}_t - \mathbf{W}_{t-1}\|_2 + \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top) \mathbf{G}_{t-1} (\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* + \|\Delta_t - \Delta_{t-1}\|_* \\ &\leq \eta L + \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top) \mathbf{G}_{t-1} (\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* + \|\Delta_t - \Delta_{t-1}\|_* \end{aligned} \quad (45)$$

where in the last inequality, we used the smoothness property in Assumption 4.1, and also the fact that $\mathbf{W}_t - \mathbf{W}_{t-1} = -\eta \mathbf{U}_t \mathbf{V}_t^\top$.

Next, we bound the term $\|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top) \mathbf{G}_{t-1} (\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_*$. First note that based on Equation 7, we can see that $\text{Range}([\mathbf{U}_{t-1} \quad \mathbf{G}_{t-1} \mathbf{V}_{t-1}]) \subseteq \text{Range}(\mathbf{U}_t)$, and therefore we have $\text{Range}(\mathbf{U}_{t-1}) \subseteq \text{Range}(\mathbf{U}_t)$. We can write:

$$\begin{aligned} \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top) \mathbf{G}_{t-1} (\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* &\leq \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top)(\mathbf{G}_{t-1} - \hat{\mathbf{G}}_{t-1})(\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* \\ &\quad + \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top) \hat{\mathbf{G}}_{t-1} (\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* \\ &\leq \|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top)(\mathbf{G}_{t-1} - \hat{\mathbf{G}}_{t-1})(\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* \\ &\leq \|(\mathbf{I} - \mathbf{U}_{t-1} \mathbf{U}_{t-1}^\top) \mathbf{G}_{t-1} (\mathbf{I} - \mathbf{V}_{t-1} \mathbf{V}_{t-1}^\top)\|_* \end{aligned} \quad (46)$$

where in the last inequality we leveraged the bound from Theorem 4.3 that upper bounds the tangent space projection error. Particularly, the fact that $\|\mathbf{G}_{t-1} - \hat{\mathbf{G}}_{t-1}\|_* = \|(\mathbf{I} - \mathbf{U}_{t-1} \mathbf{U}_{t-1}^\top) \mathbf{G}_{t-1} (\mathbf{I} - \mathbf{V}_{t-1} \mathbf{V}_{t-1}^\top)\|_*$. Plugging Equation 46 into Equation 45, we have the following recursive equation:

$$\|(\mathbf{I} - \mathbf{U}_t \mathbf{U}_t^\top) \mathbf{G}_t (\mathbf{I} - \mathbf{V}_t \mathbf{V}_t^\top)\|_* \leq \|(\mathbf{I} - \mathbf{U}_{t-1} \mathbf{U}_{t-1}^\top) \mathbf{G}_{t-1} (\mathbf{I} - \mathbf{V}_{t-1} \mathbf{V}_{t-1}^\top)\|_* + \eta L + \|\Delta_t - \Delta_{t-1}\|_* \quad (47)$$

Unrolling this recursive relation, and considering the SVD initialization at the beginning of Algorithm 1, we have:

$$\|(I - \mathbf{U}_t \mathbf{U}_t^\top) \mathbf{G}_t (I - \mathbf{V}_t \mathbf{V}_t^\top)\|_* \leq \|(I - \mathbf{U}_0 \mathbf{U}_0^\top) \mathbf{G}_0 (I - \mathbf{V}_0 \mathbf{V}_0^\top)\|_* + \eta L t + \sum_{i=1}^t \|\Delta_i - \Delta_{i-1}\|_* \quad (48)$$

Let $e_0 = \|(I - \mathbf{U}_0 \mathbf{U}_0^\top) \mathbf{G}_0 (I - \mathbf{V}_0 \mathbf{V}_0^\top)\|_*$. Note that by our initialization of \mathbf{U}_0 and \mathbf{V}_0 , e_0 is expected to be small. For instance, if we assume $\text{rank}(\mathbf{G}_0) \leq r$, then we would have $e_0 = 0$. Therefore, for simplicity we can ignore this term.

Plugging Equation 48 back into Equation 44 we have:

$$\mathbb{E}[\|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_*] \leq \beta \mathbb{E}[\|\hat{\mathbf{M}}_{t-1} - \mathbf{M}_{t-1}\|_*] + \eta L t + \frac{2t\sigma}{\sqrt{B}} \quad (49)$$

Similar to Equation 40, we can unroll the above equation to get:

$$\sum_{t=1}^T \mathbb{E}[\|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_*] \leq \frac{\eta L T}{1 - \beta} + \frac{2\sigma T}{\sqrt{B}(1 - \beta)} \quad (50)$$

where we also used the fact that $\hat{\mathbf{M}}_0 = \mathbf{M}_0$. \square

Proof of Theorem 4.5. Taking the expectation and unrolling the descent lemma, Lemma D.3, we can write:

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^T \mathbb{E}[\|\bar{\mathbf{G}}_t\|_*] &\leq \frac{\mathcal{L}(\mathbf{W}_0) - \mathbb{E}[\mathcal{L}(\mathbf{W}_{T+1})]}{\eta T} + \frac{2}{T} \sum_{t=0}^T \mathbb{E}[\|\hat{\mathbf{M}}_t - \bar{\mathbf{G}}_t\|_*] + \frac{\eta^2 L}{2} \\ &\leq \frac{\mathcal{L}(\mathbf{W}_0) - \mathbb{E}[\mathcal{L}(\mathbf{W}_{T+1})]}{\eta T} + \frac{\eta^2 L}{2} + \frac{2}{T} \sum_{t=0}^T \mathbb{E}[\|\hat{\mathbf{M}}_t - \mathbf{M}_t\|_*] + \frac{2}{T} \sum_{t=0}^T \mathbb{E}[\|\mathbf{M}_t - \bar{\mathbf{G}}_t\|_*] \end{aligned} \quad (51)$$

Plugging Equation 41 (Lemma D.4) and Equation 50 (Lemma D.5) into Equation 51 yields

$$\left(1 - \frac{2\beta}{1 - \beta}\right) \frac{1}{T} \sum_{t=0}^T \mathbb{E}[\|\bar{\mathbf{G}}_t\|_*] \leq \frac{\mathcal{L}(\mathbf{W}_0) - \mathbb{E}[\mathcal{L}(\mathbf{W}_{T+1})]}{\eta T} + \frac{\eta^2 L}{2} + \frac{4\sigma}{(1 - \beta)\sqrt{B}} + \frac{2\eta L}{1 - \beta} \quad (52)$$

Let $\beta \in (0, \frac{1}{3})$, $\eta < 1$, and $B = T$; then we can rewrite the above equation as follows:

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E}[\|\bar{\mathbf{G}}_t\|_*] \leq \mathcal{O} \left(\frac{\mathcal{L}(\mathbf{W}_0) - \mathbb{E}[\mathcal{L}(\mathbf{W}_{T+1})]}{\eta T} + \eta L + \frac{\sigma}{\sqrt{T}} \right) \quad (53)$$

Finally, by letting $\eta = \Theta \left(\sqrt{\frac{\mathcal{L}(\mathbf{W}_0) - \mathbb{E}[\mathcal{L}(\mathbf{W}_{T+1})]}{TL}} \right)$, we have the final upper bound as follows:

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E}[\|\bar{\mathbf{G}}_t\|_*] \leq \mathcal{O} \left(\frac{(\mathcal{L}(\mathbf{W}_0) - \mathbb{E}[\mathcal{L}(\mathbf{W}_{T+1})])^{\frac{1}{2}} \sqrt{L} + \sigma}{\sqrt{T}} \right). \quad (54)$$

\square