
Memory-Efficient LLM Training with Online Subspace Descent

Kaizhao Liang[†], Bo Liu[†], Lizhang Chen[†], Qiang Liu[†]

[†]The University of Texas at Austin

{kaizhaol,bliu,lzchen,lqiang}@utexas.edu

Abstract

Recently, a wide range of memory-efficient LLM training algorithms have gained substantial popularity. These methods leverage the low-rank structure of gradients to project optimizer states into a subspace using projection matrix found by singular value decomposition (SVD). However, convergence of these algorithms is highly dependent on the update rules of their projection matrix. In this work, we provide the *first* convergence guarantee for arbitrary update rules of projection matrix. This guarantee is generally applicable to optimizers that can be analyzed with Hamiltonian Descent, including most common ones, such as LION, Adam. Inspired by our theoretical understanding, we propose Online Subspace Descent, a new family of subspace descent optimizer without SVD. Instead of updating the projection matrix with eigenvectors, Online Subspace Descent updates the projection matrix with online PCA. Online Subspace Descent is flexible and introduces only minimum overhead to training. We show that for the task of pretraining LLaMA models ranging from 60M to 7B parameters on the C4 dataset, Online Subspace Descent achieves lower perplexity and better downstream tasks performance than state-of-the-art low-rank training methods across different settings and narrows the gap with full-rank baselines.¹

1 Introduction

The continual advancement in training large language models (LLMs) presents a compelling challenge in balancing computational efficiency with model performance. As the scope and complexity of these models grow, so does the necessity for innovative strategies that optimize memory usage without compromising the learning capabilities of the model. Recent approaches in low-rank adaptation strategies, including Stochastic Subspace Descent [13], LoRA [11], ReLoRA [15], Gradient Low-Rank Projection (GaLore) [25] and Sketchy [9], have paved the way for memory-efficient training by utilizing a periodically updated low-rank projection matrix to manage parameter updates. In particular, GaLore and Sketchy both utilize expensive singular value decomposition to determine the projection matrix, whereas stochastic subspace descent suggests using random matrices as projection matrices and provides convergence analysis on convex functions and objectives. However, to the best of our knowledge, no one has offered any guarantee of convergence for this class of methods on non-convex functions and objectives.

In this work, we provide the first convergence guarantee for arbitrary update rules of the projection matrix. This guarantee is significant because it is broadly applicable to a wide range of optimizers that can be analyzed within the Hamiltonian descent framework [18]. By establishing this convergence guarantee, we demonstrate that our approach is not limited to specific or narrowly defined update rules, but can be extended to include many commonly used optimizers in the field. In particular, this

¹Code is available at <https://github.com/kyleliang919/Online-Subspace-Descent>.

Algorithm 1 Online Subspace Descent

- 1: Required: Optimizer OptimizerW , learning rate ϵ_t^W , weight decay λ^W for model weights \mathbf{W}_t ; and $\{\text{OptimizerP}, \epsilon_t^P, \lambda^P\}$ for the projection matrix \mathbf{P}_t . Proper initialization.
- 2: **for** iteration t **do**
- 3: Calculate gradient $\mathbf{G}_t = \nabla L(\mathbf{W}_t)$; Update model weights \mathbf{W}_t by

$$(\hat{\Delta}_t, \hat{S}_t) = \text{OptimizerW}(\mathbf{P}_t^\top \mathbf{G}_t, \hat{S}_{t-1}), \quad \mathbf{W}_{t+1} = \mathbf{W}_t + \epsilon_t^W (\mathbf{P}_t \hat{\Delta}_{t+1} - \lambda^W \mathbf{W}_t)$$
- 4: Calculate $\mathbf{G}_t^P = \nabla L_{\mathbf{G}_t}(\mathbf{P}_t)$ for $L_{\mathbf{G}_t}(\cdot)$ in Eq (6); Update the projection \mathbf{P}_t by

$$(\Delta_t^P, S_t^P) = \text{OptimizerP}(\mathbf{G}_t^P, S_{t-1}^P), \quad \mathbf{P}_{t+1} = \mathbf{P}_t + \epsilon_t^P (\Delta_t^P - \lambda^P \mathbf{P}_t)$$
- 5: **end for**
- 6: Remark: We added weight decay as a common heuristic. We recommend using Adam for both optimizers, and set $\epsilon_t^P = \alpha \epsilon_t^W$ with a constant α (e.g., $\alpha = 5$), and $\lambda^W = \lambda^P$.

includes popular algorithms such as LION [4] and Adam [12], which are widely used in various machine learning and optimization tasks. Our results therefore offer a robust theoretical foundation for understanding and analyzing the behavior of these optimizers, ensuring their effectiveness and reliability in diverse applications.

Inspired by our theoretical understanding, we introduce a novel family of memory-efficient optimizers named Online Subspace Descent, which incorporates a dynamically changing projection matrix, replacing the conventional periodic updating approach (SVD) with online PCA. By allowing the projection matrix to evolve in response to the changing gradient landscape, Online Subspace Descent enhances the model's ability to navigate the parameter space more effectively. This dynamic adaptation aligns more closely with the natural progression of learning in deep neural networks, which is characterized by changes in the importance of different characteristics and interactions as training progresses. Through extensive experiments and comparative analysis, we demonstrate that our approach presents lower perplexity in pretraining LLaMA models (ranging from 60M to 1B parameters) on the C4 dataset compared to state-of-the-art low-rank training methods, closing the perplexity gap with full-rank baselines on language model pretraining.

2 Optimization Background

The training of deep learning models reduces to an optimization problem

$$\min_{\mathbf{W}} L(\mathbf{W}),$$

where \mathbf{W} is the set of weight matrices of the model. For simplicity, we assume $\mathbf{W} \in \mathbb{R}^{n \times m}$ is a single matrix of size (n, m) without loss of generality. For notation, we write $\langle \mathbf{A}, \mathbf{B} \rangle = \text{tr}(\mathbf{A}^\top \mathbf{B})$ for inner products of matrices, and $\|A\|^2 = \text{tr}(A^\top A)$ the Frobenius norm. We use $A \odot B$ to denote the elementwise product, and $A^{\odot 2} = A \odot A$.

Example 2.1. *Update rules of common optimizers:*

Gradient Descent : $\mathbf{W}_{t+1} = \mathbf{W}_t - \epsilon_t \nabla L(\mathbf{W}_t)$,

Momentum : $\mathbf{W}_{t+1} = \mathbf{W}_t - \epsilon_t \mathbf{M}_t, \quad \mathbf{M}_t = (1 - \beta) \nabla L(\mathbf{W}_t) + \beta \mathbf{M}_{t-1}$,

Lion-K [4] : $\mathbf{W}_{t+1} = \mathbf{W}_t - \epsilon_t \nabla \mathcal{K}(\mathbf{N}_t), \quad \mathbf{N}_t = (1 - \beta_1) \nabla L(\mathbf{W}_t) + \beta_1 \mathbf{M}_t$
 $\mathbf{M}_t = (1 - \beta_2) \nabla L(\mathbf{W}_t) + \beta_2 \mathbf{M}_{t-1}$,

Adam [12] : $\mathbf{W}_{t+1} = \mathbf{W}_t - \epsilon_t \frac{\mathbf{M}_t}{\sqrt{\mathbf{V}_t + e}}, \quad \mathbf{M}_t = (1 - \beta_{1t}) \nabla L(\mathbf{W}_t) + \beta_{1t} \mathbf{M}_{t-1},$
 $\mathbf{V}_t = (1 - \beta_{2t}) \nabla L(\mathbf{W}_t)^{\odot 2} + \beta_{2t} \mathbf{V}_{t-1}$,

where ϵ_t are step sizes, and $\mathbf{M}_t, \mathbf{V}_t$ are the first and second order momentum, and β, β_1, β_2 are momentum coefficients in $(0, 1)$, with $\beta_{it} = \frac{\beta_i - \beta_i^{t+1}}{1 - \beta_i^{t+1}}$, $i = 1, 2$ for $\beta_1, \beta_2 \in (0, 1)$ in Adam, and \mathcal{K} is any convex function with $\nabla \mathcal{K}(\mathbf{0}) = \mathbf{0}$ for Lion-K [4], and Lion [5] uses $\mathcal{K}(\mathbf{X}) = \|\mathbf{X}\|_{1,1}$ and $\nabla \mathcal{K}(\mathbf{X}) = \text{sign}(\mathbf{X})$.

These optimizers can be unifiedly viewed as updating \mathbf{W}_t together with an optimizer state \mathbf{S}_t :

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \phi_t(\mathbf{S}_t), \quad \mathbf{S}_t = \psi_t(\mathbf{S}_{t-1}, \nabla L(\mathbf{W}_t)), \quad (1)$$

with some mapping ϕ_t, ψ_t . We have $\mathbf{S}_t = \mathbf{M}_t$ for momentum and $\mathbf{S}_t = \{\mathbf{M}_t, \mathbf{V}_t\}$ for Adam. Note that both $\mathbf{M}_t, \mathbf{V}_t$ are of the same size as the model weights \mathbf{W}_t , resulting in high memory consumption for large models. This issue is particularly pronounced for Adam, which typically yields the best performance for large language models (LLMs) but incurs the highest memory cost due to the need to maintain both \mathbf{M}_t and \mathbf{V}_t . One key challenge is to retain the high performance of Adam while enhancing its memory efficiency.

Hamiltonian+Descent One powerful approach to studying the dynamic properties of optimizers is to examine their continuous-time ODE forms in the limit of infinitesimal step size. The continuous-time forms provide clearer insights into the asymptotic convergence of the algorithm, abstracting away the choices of step size, discretization, and stochastic errors. The underlying logic is that a "sound" optimizer should be guaranteed to converge to local optima of the loss, at least when using sufficiently small step sizes.

Inspired by [4, 18], we observe that the continuous-time form of many common optimizers yields a *Hamiltonian+Descent* structure,

$$\begin{aligned} \frac{d}{dt} \mathbf{W}_t &= \partial_{\mathbf{S}} H(\mathbf{W}_t, \mathbf{S}_t) - \Phi(\partial_{\mathbf{W}} H(\mathbf{W}_t, \mathbf{S}_t)) \\ \frac{d}{dt} \mathbf{S}_t &= -\partial_{\mathbf{W}} H(\mathbf{W}_t, \mathbf{S}_t) - \Psi(\partial_{\mathbf{S}} H(\mathbf{W}_t, \mathbf{S}_t)), \end{aligned} \quad (2)$$

where $H(\mathbf{W}, \mathbf{S})$ is a Hamiltonian (or Lyapunov) function that satisfies

$$\min_{\mathbf{S}} H(\mathbf{W}, \mathbf{S}) = L(\mathbf{W}), \quad \forall \mathbf{W},$$

so that minimizing $L(\mathbf{W})$ reduces to minimizing $H(\mathbf{W}, \mathbf{S})$; and $\Phi(\cdot), \Psi(\cdot)$ are two monotonic mappings satisfying

$$\|\mathbf{X}\|_{\Phi}^2 := \langle \mathbf{X}, \Phi(\mathbf{X}) \rangle \geq 0, \quad \|\mathbf{X}\|_{\Psi}^2 := \langle \mathbf{X}, \Psi(\mathbf{X}) \rangle \geq 0, \quad \forall \mathbf{X}.$$

With $\Phi(\mathbf{X}) = \Psi(\mathbf{X}) = 0$, the system in (2) reduces to the standard Hamiltonian system that keeps $H(\mathbf{W}_t, \mathbf{S}_t) = \text{const}$ along the trajectory. When adding the descending components with Φ and Ψ , the system then keeps $H(\mathbf{W}, \mathbf{S})$ monotonically non-decreasing:

$$\begin{aligned} \frac{d}{dt} H(\mathbf{W}_t, \mathbf{S}_t) &= \left\langle \partial_{\mathbf{W}} H_t, \frac{d}{dt} \mathbf{W}_t \right\rangle + \left\langle \partial_{\mathbf{S}} H_t, \frac{d}{dt} \mathbf{S}_t \right\rangle \\ &= \langle \partial_{\mathbf{W}} H_t, \partial_{\mathbf{S}} H_t - \Phi(\partial_{\mathbf{W}} H_t) \rangle + \langle \partial_{\mathbf{S}} H_t, -\partial_{\mathbf{W}} H_t - \Psi(\partial_{\mathbf{S}} H_t) \rangle \\ &= -\|\partial_{\mathbf{W}} H_t\|_{\Phi}^2 - \|\partial_{\mathbf{S}} H_t\|_{\Psi}^2 \leq 0, \end{aligned} \quad (3)$$

where we write $\partial_{\mathbf{W}} H_t = \partial_{\mathbf{W}} H(\mathbf{W}_t, \mathbf{S}_t)$ and similarly for $\partial_{\mathbf{S}} H_t$. The main idea is that the cross terms $\langle \partial_{\mathbf{W}} H_t, \partial_{\mathbf{S}} H_t \rangle$ are canceled, leaving only the negative terms.

Example 2.2. The momentum method yields following continuous-time form and Hamiltonian:

$$\frac{d}{dt} \mathbf{W}_t = -\mathbf{M}_t, \quad \frac{d}{dt} \mathbf{M}_t = a(\nabla L(\mathbf{W}_t) - \mathbf{M}_t), \quad \text{with} \quad H(\mathbf{W}, \mathbf{M}) = L(\mathbf{W}) + \frac{\|\mathbf{M}\|^2}{2a}.$$

Example 2.3. Adam [12] yields the following continuous-time form and Hamiltonian,

$$\begin{aligned} \frac{d}{dt} \mathbf{W}_t &= -\frac{\mathbf{M}_t}{\sqrt{\mathbf{V}_t} + e}, \quad \frac{d}{dt} \mathbf{M}_t = a(\nabla L(\mathbf{W}_t) - \mathbf{M}_t), \quad \frac{d}{dt} \mathbf{V}_t = b(\nabla L(\mathbf{W}_t)^{\odot 2} - \mathbf{V}_t), \\ \text{with} \quad H(\mathbf{W}, \mathbf{M}, \mathbf{V}) &= L(\mathbf{W}) + \frac{1}{2a} \left\langle \frac{\mathbf{M}}{\sqrt{\mathbf{V}} + e}, \mathbf{M} \right\rangle, \end{aligned}$$

for which we can show that $\frac{d}{dt} H(\mathbf{W}_t, \mathbf{M}_t, \mathbf{V}_t) \leq 0$ when $a \geq b/4$.

Example 2.4. The Lion-K optimizer [5, 4] (without weight decay) can be written into

$$\frac{d}{dt} \mathbf{W}_t = \nabla \mathcal{K}((1-b)\mathbf{M}_t - b\nabla L(\mathbf{W}_t)), \quad \frac{d}{dt} \mathbf{M}_t = -a(\nabla L(\mathbf{W}_t) + \mathbf{M}_t),$$

where $a \geq 0, b \in [0, 1]$ and $\mathcal{K}(\mathbf{X})$ is any convex function that attains the minimum at $\mathbf{X} = 0$. One of its Hamiltonians that yields the Hamiltonian+descent structure (Eq (13) in Chen et al. [4]) is

$$H(\mathbf{W}, \mathbf{M}) = aL(\mathbf{W}) + \frac{1}{1-b} \mathcal{K}((1-b)\mathbf{M}).$$

3 Memory-Efficient Optimizers via Online Subspace Descent

We introduce the idea of constructing memory efficient optimizers by descending in the subspaces that dynamically changes across iterations as motivated by GaLore [25] and Sketchy [9]. We first derive *static* subspace descent by restricting the whole optimization on a subspace (Section 3.1), and then propose to dynamically change the subspace across iterations as a heuristic to attain the optimization in the full space while only using subspace descent (Section 3.2). In particular, we propose to update the subspaces via continuous online PCA like updates to avoids the need of exact SVD like in GaLore and Sketchy (Section 3.2). Finally, we remark in Section 3.3 the heuristic nature of the derivation of the method and highlight the difficulty in theoretical understanding, which motivates our analysis based on Hamiltonian dynamics in Section 4.

3.1 Static Subspace Descent

One popular approach to improving memory efficiency is to confine the optimization to a low-dimensional space. To do this, we impose a low rank structure of $\mathbf{W} = \mathbf{P}\hat{\mathbf{W}}$, where $\mathbf{P} \in \mathbb{R}^{n \times k}$ is a projection matrix to be determined later and $\hat{\mathbf{W}} \in \mathbb{R}^{k \times m}$ is a dimension-reduced parameter. When $k \ll n$, \mathbf{P} and $\hat{\mathbf{W}}$ are much smaller in size compared to \mathbf{W} . Now consider

$$\min_{\hat{\mathbf{W}}} L(\mathbf{P}\hat{\mathbf{W}}).$$

Applying the optimizer from (1) to update $\hat{\mathbf{W}}$ along with an optimizer state $\hat{\mathbf{S}}$, and mapping the update rule $\hat{\mathbf{W}}_{t+1} = \hat{\mathbf{W}}_t + \phi_t(\hat{\mathbf{S}}_t)$ to that of $\mathbf{W} = \mathbf{P}\hat{\mathbf{W}}_t$, we get

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \mathbf{P}\phi_t(\hat{\mathbf{S}}_t), \quad \hat{\mathbf{S}}_t = \psi_t(\hat{\mathbf{S}}_{t-1}, \mathbf{P}^\top \nabla L(\mathbf{W}_t)), \quad (4)$$

where we used the fact that $\nabla_{\hat{\mathbf{W}}} L(\mathbf{P}\hat{\mathbf{W}}) = \mathbf{P}^\top \nabla_{\mathbf{W}} L(\mathbf{W})$. This yields a more memory-efficient optimizer, as the size of $\hat{\mathbf{S}}_t$ is proportional to that of $\hat{\mathbf{W}}_t$, much smaller than \mathbf{S}_t in (1) when $k \ll n$.

3.2 Online Subspace Descent

With a static \mathbf{P} , regardless of its values, the parameter \mathbf{W} is restricted to have a low rank structure. Although low rank assumption is proved to be useful for fine-tuning with LoRA-like methods [11], it is often too limited for pre-training or when the desirable model weights are not inherently low-rank.

To address this problem, Zhao et al. [25] suggested to keep the projected updated in (4), but use different \mathbf{P} across the iterations:

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \mathbf{P}_t \phi_t(\hat{\mathbf{S}}_t), \quad \hat{\mathbf{S}}_t = \psi_t(\hat{\mathbf{S}}_{t-1}, \mathbf{P}_t^\top \nabla L(\mathbf{W}_t)), \quad \mathbf{P}_{t+1} = \chi_t(\mathbf{P}_t, \mathbf{W}_t, \hat{\mathbf{S}}_t),$$

where χ_t is a update rule of \mathbf{P}_t that will be determined in the sequel. The intuition is to open up different projection directions at different iterations, so that optimization can be conducted in different subspaces across different iterations. This is similar to the update of coordinate descent, except in a continuous fashion. Note that the update of \mathbf{P}_t can be done in parallel with that of $(\mathbf{W}_t, \hat{\mathbf{S}}_t)$, and incurs no slowdown once it is fast enough to not cause a speed bottleneck.

Example 3.1. Examples of common optimizers equipped with online subspace descent:

$$\begin{aligned} \text{Gradient Descent : } \quad & \mathbf{W}_{t+1} = \mathbf{W}_t - \epsilon_t \mathbf{P}_t \mathbf{P}_t^\top \mathbf{G}_t, & \mathbf{G}_t = \nabla L(\mathbf{W}_t), \\ \text{Momentum : } \quad & \mathbf{W}_{t+1} = \mathbf{W}_t - \epsilon_t \mathbf{P}_t \hat{\mathbf{M}}_t, & \hat{\mathbf{M}}_t = (1 - \beta) \mathbf{P}_t^\top \mathbf{G}_t + \beta \hat{\mathbf{M}}_{t-1}, \\ \text{Lion-K : } \quad & \mathbf{W}_{t+1} = \mathbf{W}_t - \epsilon_t \mathbf{P}_t \nabla \mathcal{K}(\hat{\mathbf{N}}_t), & \hat{\mathbf{G}}_t = \mathbf{P}_t^\top \mathbf{G}_t \\ & \hat{\mathbf{N}}_t = (1 - \beta_1) \hat{\mathbf{G}}_t + \beta_1 \hat{\mathbf{M}}_t, & \hat{\mathbf{M}}_t = (1 - \beta_2) \hat{\mathbf{G}}_t + \beta_2 \hat{\mathbf{M}}_{t-1}, \\ \text{Adam : } \quad & \mathbf{W}_{t+1} = \mathbf{W}_t - \epsilon_t \mathbf{P}_t \frac{\hat{\mathbf{M}}_t}{\sqrt{\hat{\mathbf{V}}_t} + \epsilon}, & \hat{\mathbf{G}}_t = \mathbf{P}_t^\top \mathbf{G}_t, \\ & \hat{\mathbf{M}}_t = (1 - \beta_{1t}) \hat{\mathbf{G}}_t + \beta_{1t} \hat{\mathbf{M}}_{t-1}, & \hat{\mathbf{V}}_t = (1 - \beta_{2t}) \hat{\mathbf{G}}_t^{\odot 2} + \beta_{2t} \hat{\mathbf{V}}_{t-1}. \end{aligned}$$

How Should \mathbf{P}_t be Updated? It is useful to draw intuition from the projected gradient descent rule

$$\mathbf{W}_{t+1} = \mathbf{W}_t - \epsilon_t \mathbf{P}_t \mathbf{P}_t^\top \mathbf{G}_t, \quad \mathbf{G}_t = \nabla L(\mathbf{W}_t), \quad (5)$$

in which $\mathbf{P}_t \mathbf{P}_t^\top$ can be viewed as a low rank preconditioning of \mathbf{G}_t . To make it follow the exact gradient descent, we hope to make $\mathbf{P}_t \mathbf{P}_t^\top \mathbf{G}_t$ approximate \mathbf{G}_t as much as possible. In Galore, this is achieved by performing singular value decomposition (SVD) on \mathbf{G}_t periodically every T iterations:

$$\mathbf{P}_t, _, _ = \text{torch.linalg.svd}(\mathbf{G}_{T \lfloor t/T \rfloor}),$$

where $T \lfloor t/T \rfloor$ is the largest multiple of T less than or equal to t . However, numerical SVD incurs a large computational cost for very large models. Also, since \mathbf{P}_t is fully determined by $\mathbf{G}_{T \lfloor t/T \rfloor}$ calculated from a single mini-batch at the last periodic point, it does not incorporate the gradient information from all data in a timely fashion.

In this work, we propose to update \mathbf{P}_t in a continuous online fashion that incorporates the most recent gradient information in a timely fashion, without calling `torch.linalg.decompositions` routines. We view the update of \mathbf{P}_t as conducting an online principal component analysis (PCA) based on the streaming of $\{\mathbf{G}_t\}$. In particular, we propose to update \mathbf{P}_t at time t by minimizing the following PCA objective:

$$L_{\mathbf{G}_t}(\mathbf{P}) = \left\| \mathbf{P} \mathbf{P}^\top \tilde{\mathbf{G}}_t - \tilde{\mathbf{G}}_t \right\|^2 + \lambda \left\| \mathbf{P}^\top \mathbf{P} - \mathbf{I}_{k \times k} \right\|^2, \quad \tilde{\mathbf{G}}_t = \frac{\mathbf{G}_t}{\|\mathbf{G}_t\|}, \quad (6)$$

where $\|\mathbf{A}\| = \text{tr}(\mathbf{A}^\top \mathbf{A})^{1/2}$ and $\mathbf{I}_{k \times k}$ is the $k \times k$ identity matrix; we introduced an auxiliary loss to encourage the columns of \mathbf{P} to be orthonormal and normalizes \mathbf{G}_t to increase stability.

The key property of $L_{\mathbf{G}_t}(\mathbf{P})$ in (6) is that all its stable local minimum is a global minimum, and \mathbf{P} is a global minimum iff $\mathbf{P} \mathbf{P}^\top \tilde{\mathbf{G}}_t$ forms the optimal rank- k approximation of $\tilde{\mathbf{G}}_t$ [e.g., 3]; moreover, we have $\mathbf{P}^\top \mathbf{P} = \mathbf{I}_{k \times k}$ at optima when $\lambda > 0$.

Instead of minimizing $L_{\mathbf{G}_t}(\mathbf{P})$ exactly, to retain computational efficiency, we propose to update \mathbf{P}_t by only performing one step of optimization on $L_{\mathbf{G}_t}(\mathbf{P})$:

$$\mathbf{P}_{t+1} = \text{OptimizerP.step}(\mathbf{P}_t, \nabla_{\mathbf{P}} L_{\mathbf{G}_t}(\mathbf{P}_t)),$$

where `OptimizerP.step` can be a favorite optimizer, such as gradient descent or Adam. Note that when using Adam, we introduce a copy of optimizer state \mathbf{S}_t^P for \mathbf{P}_t . See Algorithm 1. Compared to the exact SVD, each online update of \mathbf{P}_t here is fast and can be executed in parallel with the $(\mathbf{W}_t, \hat{\mathbf{S}}_t)$ updates to avoid slowdown.

3.3 Difficulty in Theoretical Understanding

The idea above of projecting an arbitrary optimizer with a dynamically changing \mathbf{P}_t is heuristically motivated and lacks an immediate rigorous theoretical justification. The main challenge lies in the complex interaction between the update of \mathbf{U}_t and the optimization state \mathbf{S}_t , which could potentially degrade the convergence and other theoretical properties of the original optimizer. A key question is whether we can develop a theoretical framework to understand how \mathbf{P}_t impacts the optimizer's convergence behavior and provide guidance for the design of the update rules of \mathbf{P}_t .

To gain understanding, it is useful to first exam the simple case of projected gradient descent in (5) which does not have an optimizer state ($\mathbf{S}_t = \emptyset$). In this case, since $\mathbf{P}_t \mathbf{P}_t^\top$ is positive semi-finite, the update $\mathbf{P}_t \mathbf{P}_t^\top \mathbf{G}_t$ is always non-increasing direction of $L(\mathbf{W})$ for any \mathbf{P}_t . The algorithm is essentially a variant of coordinate or subspace descent, where \mathbf{P}_t defines the subspace on which one step of gradient descent is conducted at iteration t . To ensure that (5) finds a local optimum, we mainly need to ensure that $\mathbf{P}_t \mathbf{G}_t = 0$ only if $\mathbf{G}_t = 0$ to prevent the optimizer from stopping prematurely; this is a mild condition that can be satisfied e.g. when \mathbf{P}_t is updated by (online) PCA on \mathbf{G}_t .

Unfortunately, this coordinate-descent-like interpretation does not apply to more advanced optimizers that track a momentum state \mathbf{S}_t . This is because \mathbf{S}_t accumulates the information from the projected gradient $\mathbf{P}_\tau \mathbf{G}_\tau$ at all earlier iterations $\tau \leq t$. As \mathbf{P}_t changes across time, it is unclear whether the gradient projected to different subspaces \mathbf{P}_τ would be coherent with each other, and useful for future updates that are conducted in different subspaces \mathbf{P}_t for $t > \tau$. The difficulty is the inertia effect of \mathbf{S}_t that entangles the different subspaces, making the dynamic behavior fundamentally more complicated than naive coordinate descent where the descent in different subspaces is uncoupled. This is what we address in Section 4 via the Hamiltonian descent framework.

4 Hamiltonian Descent Meets Subspace Descent: A Lyapunov Analysis

In this section, we show a surprising result that the complication outlined above in Section 3.3 is not a problem for optimizers that yields the Hamiltonian+descent structure in (2). Our result is two-fold:

- Section 4.1: When applying Online Subspace Descent on systems in (2), the Hamiltonian+descent structure is preserved once the update rule of \mathbf{P}_t has a smooth continuous-time limit. Hence, under very mild conditions, Online Subspace Descent equipped with common optimizers like Adam and Lion automatically yield a Lyapunov function and hence benign continuous-time convergence. Moreover, \mathbf{P}_t can, in fact, be generalized to an arbitrary linear operator as shown in Section 4.3.
- Section 4.2: For any smooth \mathbf{P}_t update rules that eliminates the degenerate case of $\mathbf{P}_t^\top \mathbf{G}_t = 0$ while $\mathbf{G}_t = 0$ at convergence, the online subspace optimizer guarantees to converge in continuous time to a stationary point of the loss $L(\mathbf{W})$. This mild condition is satisfied, for example, when \mathbf{P}_t is updated by a typical optimizer on the online PCA objective $L_{\mathbf{G}_t}(\mathbf{P})$.

4.1 Online Subspace Descent Preserves the Hamiltonian+Descent Structure

Applying dynamic projection to Hamiltonian descent in (2), we obtain the following systems:

$$\begin{aligned} \frac{d}{dt} \mathbf{W}_t &= \mathbf{P}_t \partial_{\hat{\mathbf{S}}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t) - \Phi(\partial_{\mathbf{W}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t)) \\ \frac{d}{dt} \hat{\mathbf{S}}_t &= -\mathbf{P}_t^\top \partial_{\mathbf{W}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t) - \Psi(\partial_{\hat{\mathbf{S}}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t)) \\ \frac{d}{dt} \mathbf{P}_t &= \Gamma(\mathbf{P}_t, \nabla L(\mathbf{W}_t)), \end{aligned} \quad (7)$$

where Γ specifies the update rule of \mathbf{P}_t . Following essentially the same derivation as (3), one can show that $H(\mathbf{W}, \mathbf{S})$ remains a Lyapunov function of (7), regardless of the choice of Γ :

$$\begin{aligned} \frac{d}{dt} H(\mathbf{W}_t, \hat{\mathbf{S}}_t) &= -\|\partial_{\mathbf{W}} H_t\|_\Phi^2 - \|\partial_{\mathbf{S}} H_t\|_\Psi^2 + \langle \partial_{\mathbf{W}} H_t, \mathbf{P}_t \partial_{\hat{\mathbf{S}}} H_t \rangle - \langle \partial_{\hat{\mathbf{S}}} H_t, \mathbf{P}_t^\top \partial_{\mathbf{W}} H_t \rangle \\ &= -\|\partial_{\mathbf{W}} H_t\|_\Phi^2 - \|\partial_{\mathbf{S}} H_t\|_\Psi^2 \leq 0, \end{aligned} \quad (8)$$

where the key is to use the *adjoint* property of \mathbf{P} and \mathbf{P}^\top that $\langle \mathbf{P}_t \mathbf{X}, \mathbf{Y} \rangle = \langle \mathbf{X}, \mathbf{P}_t^\top \mathbf{Y} \rangle$, which cancels the crossing terms, independent of the values of \mathbf{P}_t . There is no requirement on Γ here, besides that the derivative in (8) should exist. As shown in Section 4.3, we can generalize (8) by replacing \mathbf{P}_t and \mathbf{P}_t^\top with a general linear operator \mathcal{P}_t and its adjoint \mathcal{P}_t^* .

Please refer to Appendix A for continuous-time Momentum, Lion- \mathcal{K} and Adam with subspace descent and their Hamiltonian functions.

4.2 Convergence to Local Optima

In addition to the Lyapunov structure, we need an additional mild condition on the update rule of \mathbf{P}_t to ensure the system converges to the local optimum of the loss $L(\mathbf{W})$. The main idea is to prevent the system from stopping prematurely before reaching zero gradient $\mathbf{G}_t = 0$ by excluding the degenerate case of $\mathbf{P}_t \mathbf{G}_t = 0$ while $\mathbf{G}_t \neq 0$ in the invariant set of the system.

Assumption 4.1. Assume the functions in system (7) are continuously differentiable and

- i) $\frac{d}{dt} H(\mathbf{W}_t, \hat{\mathbf{S}}_t) = 0$ implies $\hat{\mathbf{G}}_t = \mathbf{P}_t^\top \nabla L(\mathbf{W}_t) = 0$ and $\frac{d}{dt} \mathbf{W}_t = 0$.
- ii) When $\mathbf{G}_t \equiv \mathbf{G} \neq 0$, the set $\{\mathbf{P}: \mathbf{P}^\top \mathbf{G} = 0\}$ is not a positive invariant set of $\frac{d}{dt} \mathbf{P}_t = \Gamma(\mathbf{P}_t, \mathbf{G}_t)$.

This is a mild condition. Assumption i) says that the optimizer should stop at a point with $\hat{\mathbf{G}}_t = 0$, which is easy to verify for the common optimizers like momentum, Adam, Lion- \mathcal{K} . Assumption ii) ensures $\hat{\mathbf{G}}_t = 0$ would imply $\mathbf{G}_t = 0$ in invariance sets, which is satisfied when for example, \mathbf{P}_t is updated by a reasonable optimizer of the online PCA loss that converges to a stable local minimum.

Theorem 4.2. Assume Assumption 4.1 holds. Let $(\mathbf{W}_t, \mathbf{S}_t, \mathbf{P}_t)_t$ be a bounded solution of (7), then all the accumulation points $\{\mathbf{W}_t\}$ as $t \rightarrow +\infty$ are stationary points of $L(\mathbf{W})$.

Proof. By LaSalle's invariance principle, the positive limit set of $(\mathbf{W}_t, \mathbf{S}_t, \mathbf{P}_t)_t$ must be contained in \mathcal{I} , where $\mathcal{I} = \{\text{the union of complete trajectories satisfying } \frac{d}{dt} H(\mathbf{W}_t, \hat{\mathbf{S}}_t) = 0, \forall t\}$.

From the Assumption i), the trajectories contained in \mathcal{I} must satisfy $\frac{d}{dt} \mathbf{W}_t = 0$, which implies $\frac{d}{dt} \mathbf{G}_t = \frac{d}{dt} \nabla L(\mathbf{W}_t) = 0$ and $\hat{\mathbf{G}}_t = 0$ and hence $\mathbf{G}_t \equiv \mathbf{G}$ is a constant with $\mathbf{P}_t^\top \mathbf{G} = 0$. Moreover, from Assumption ii), we must have $\nabla L(\mathbf{W}_t) = \mathbf{G}_t \equiv 0$, since otherwise the trajectory is not invariant. As a result, all trajectories in the limit set \mathcal{I} must have $\nabla L(\mathbf{W}_t) = 0$. Because $\frac{d}{dt} \mathbf{W}_t = 0$, these trajectories are static points of \mathbf{W}_t . \square

4.3 Online Subspace Descent with General Linear Projection Operators

We can generalize the online subspace descent with general linear operators:

$$\begin{aligned}\frac{d}{dt} \mathbf{W}_t &= \mathcal{P}_t(\partial_{\hat{\mathbf{S}}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t)) - \Phi(\partial_{\mathbf{W}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t)) \\ \frac{d}{dt} \hat{\mathbf{S}}_t &= -\mathcal{P}_t^*(\partial_{\mathbf{W}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t)) - \Psi(\partial_{\hat{\mathbf{S}}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t)) \\ \frac{d}{dt} \mathcal{P}_t &= \Gamma(\mathcal{P}_t, \nabla L(\mathbf{W}_t)),\end{aligned}$$

where we generalize \mathcal{P}_t to be any linear operator \mathcal{P}_t with an adjoint operator \mathcal{P}_t^* , satisfying

$$\langle \mathbf{X}, \mathcal{P}_t(\mathbf{Y}) \rangle = \langle \mathcal{P}_t^*(\mathbf{X}), \mathbf{Y} \rangle, \quad \forall \mathbf{X}, \mathbf{Y}.$$

The derivation of Lyapunov follows a similar way:

$$\begin{aligned}\frac{d}{dt} H(\mathbf{W}_t, \hat{\mathbf{S}}_t) &= -\|\partial_{\mathbf{W}} H_t\|_\Phi^2 - \|\partial_{\hat{\mathbf{S}}} H_t\|_\Psi^2 + \langle \partial_{\mathbf{W}} H_t, \mathcal{P}_t(\partial_{\hat{\mathbf{S}}} H_t) \rangle - \langle \partial_{\hat{\mathbf{S}}} H_t, \mathcal{P}_t^*(\partial_{\mathbf{W}} H_t) \rangle \\ &= -\|\partial_{\mathbf{W}} H_t\|_\Phi^2 - \|\partial_{\hat{\mathbf{S}}} H_t\|_\Psi^2 \leq 0,\end{aligned}$$

where the crossing terms are again canceled due to the adjoint property.

As an example of the general framework, consider $\mathcal{P}_t(\mathbf{X}) = \mathbf{P}_t \mathbf{X} \mathbf{Q}_t$, where \mathbf{Q}_t is another projection matrix applied on the different dimension of \mathbf{X} (see also [25]). The adjoint operator of \mathcal{P}_t is $\mathcal{P}_t^*(\mathbf{X}) = \mathbf{P}_t^\top \mathbf{X} \mathbf{Q}_t^\top$. This can be verified by

$$\langle \mathbf{P}_t \mathbf{X} \mathbf{Q}_t, \mathbf{Y} \rangle = \text{tr}(\mathbf{P}_t \mathbf{X} \mathbf{Q}_t \mathbf{Y}^\top) = \text{tr}(\mathbf{X} \mathbf{Q}_t \mathbf{Y}^\top \mathbf{P}_t) = \text{tr}(\mathbf{X} (\mathbf{P}_t^\top \mathbf{Y} \mathbf{Q}_t^\top)^\top) = \langle \mathbf{X}, \mathbf{P}_t^\top \mathbf{Y} \mathbf{Q}_t^\top \rangle.$$

The subspace descent system of this operator is

$$\begin{aligned}\frac{d}{dt} \mathbf{W}_t &= \mathbf{P}_t \partial_{\hat{\mathbf{S}}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t) \mathbf{Q}_t - \Phi(\partial_{\mathbf{W}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t)) \\ \frac{d}{dt} \hat{\mathbf{S}}_t &= -\mathbf{P}_t^\top \partial_{\mathbf{W}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t) \mathbf{Q}_t^\top - \Psi(\partial_{\hat{\mathbf{S}}} H(\mathbf{W}_t, \hat{\mathbf{S}}_t)) \\ \frac{d}{dt} \mathbf{P}_t &= \Gamma_P(\mathbf{P}_t, \mathbf{Q}_t, \nabla L(\mathbf{W}_t)) \\ \frac{d}{dt} \mathbf{Q}_t &= \Gamma_Q(\mathbf{P}_t, \mathbf{Q}_t, \nabla L(\mathbf{W}_t)),\end{aligned}$$

where $\mathbf{P}_t, \mathbf{Q}_t$ can be updated jointly via an online SVD on \mathbf{G}_t .

Another linear operator that involves two matrices is $\mathcal{P}_t(\mathbf{X}) = \mathbf{P}_t \mathbf{X} + \mathbf{X} \mathbf{Q}_t$, which yields $\mathcal{P}_t^*(\mathbf{X}) = \mathbf{P}_t^\top \mathbf{X} + \mathbf{X} \mathbf{Q}_t^\top$.

5 Experiment

We answer a number of key questions with pretraining experiments of LLaMA [22] on the C4 dataset [20]. All experiments except for large 7B experiments are conducted on a *single* NVIDIA A100 GPU.

5.1 Why do we Need Online Subspace Descent?

Overall, Online Subspace Descent offers two major advantages over previous methods that rely on SVD, better convergence and lower overhead. In this section, we discuss both in detail.

First, Online Subspace Descent closes the gap between the state-of-the-art low-rank method and full rank baseline uniformly across different model sizes, as shown in figure 1. A highlight amongst these results is LLaMA 1B (SS 256). As shown in table 1, Online Subspace Descent attains significant improvement over GaLore in perplexity, while consuming a similar amount of GPU memory (8.64 GB v.s 9.01 GB). One additional observation in 1 shows as model size and sequence length grow, Online Subspace Descent becomes more effective. We hypothesize that this is due to the higher intrinsic rank of the underlying optimization problem in larger models. Hence, the positive impact on the convergence of the online update of P_t becomes more obvious. See more details in Appendix B.

Method	Perplexity(↓)		
	60M	350M	1B
8bit-AdamW (Full Rank)	32.75	30.43	29.40
GaLore (Rank = 512)	57.03	44.34	35.52
Ours (Rank = 512)	56.12	43.67	31.30

Table 1: Pretraining LLaMA 1B with a sequence length of 256 and for 10K steps, perplexity was reported as the training average of the last 10 steps. AdamW8bit serves as the base optimizer for both.

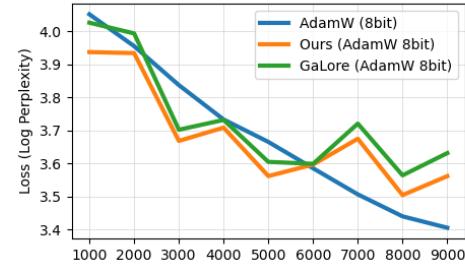


Figure 1: Validation perplexity of LLaMA 1B with sequence length 256, rank 512 for 10K steps.

Another favorable characteristic of Online Subspace Descent is its minimum overhead. In figure 2, we measure and analyze the execution time of SVD and online PCA on a popular data center GPU (A100) and a consumer GPU (RTX 3090). The typical Pytorch implementation of SVD can be up to 142 times slower than running a single-step online PCA on representative weight tensors from LLaMA architectures. Online PCA is fast because it is implemented as a single optimization step with respect to a simple loss function. Hence, each step of online PCA can be cleverly scheduled and hidden in the weight optimization step when executed in parallel, whereas SVD is too expensive to be hidden.

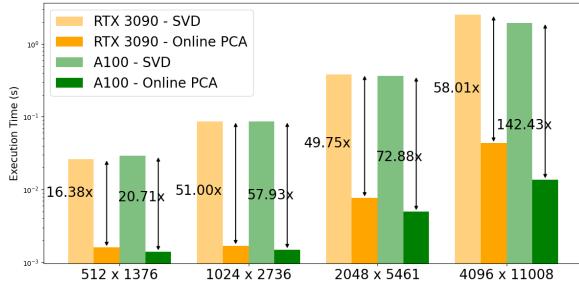


Figure 2: The execution time of `torch.svd` and that a single-step `backward()` call for online PCA in PyTorch, on matrices of typical shapes in linear layers in the LLaMA 60M to 7B. Thanks to the high speed of single-step online PCA, P_t updates can be executed in parallel with weight updates, adding no overhead to the training process. In contrast, SVD incurs significant overhead as the model and weight tensor sizes increase.

5.2 What Rank Should we Pick for *Online Subspace Descent*?

We conduct an ablation study on the rank of Online Subspace Descent. Figure 3 shows that the final perplexity is inversely correlated with rank: higher ranks result in lower convergent perplexity. However, the rate of reduction of perplexity decreases as the rank increases, eventually reaching a saturation point. We propose an intuitive explanation for this phenomenon. In language modeling, high-frequency tokens can be effectively learned with low-rank training. However, learning lower-frequency tokens requires higher ranks. Once these lower-frequency tokens are adequately learned, further increasing the rank does not significantly decrease perplexity. In conclusion, given sufficient time and resources, higher ranks yield better performance for Online Subspace Descent. It is recommended that the highest rank be selected until the perplexity reduction saturates.

5.3 What are the Best Hyperparameters?

α and λ : The parameter α controls the update speed of P_t , while λ determines the regularization strength on the optimization objective of P_t . Empirically, we find that the result is not sensitive to λ

for small models (60M), and set $\lambda = 0.1$ for all subsequent experiments. We find that α must be kept small to avoid instability (Figure 3), and we set $\alpha = 5$ for all experiments.

Learning rate: For the small model (60M), learning rate choices are more flexible, producing similar results. However, for larger models (350M, 1B), we recommend using a learning rate that is 10 times smaller, specifically 0.001. Larger learning rates cause unrecoverable spikes and instability, a general characteristic observed across all methods. See additional hyperparameter choices in Appendix B.

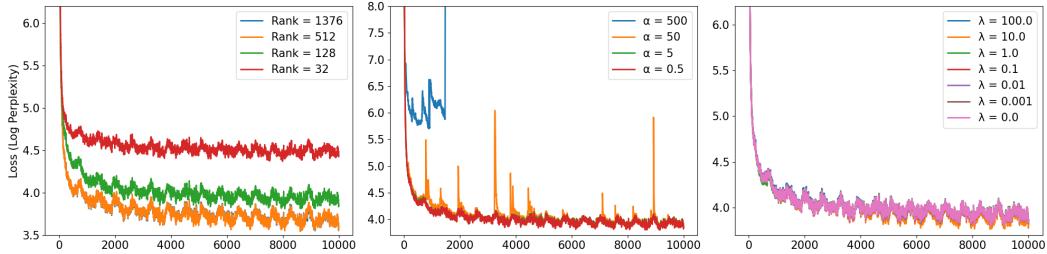


Figure 3: From left to right are loss curves of 10K steps on LLaMA 60M: leftmost is the sweep of rank, middle is the sweep of α and rightmost is the sweep of λ .

5.4 Can *Online Subspace Descent* be Applied to Different Optimizers?

One straightforward extension of Online Subspace Descent is to apply it to other base optimizers beyond AdamW8bit. We conduct ablation studies on LION [6] and Adafactor [21], finding that Online Subspace Descent behaves similarly to how it does with AdamW8bit. Despite the initial observation that updating \mathbf{P}_t with AdamW8bit consistently yields better results, we discover that updating \mathbf{P}_t with simple SGD can achieve similar performance.

Method	GaLore		Ours			
	Lion	Adaf.	Lion+Lion	Adaf.+Adaf.	Lion+AdamW	Adaf.+AdamW
Perplexity	46.90	34.32	57.97	47.61	44.76	34.15

Table 2: LLaMA 60M on C4 with sequence length 1024, with optimizers on \mathbf{P}_t and \mathbf{W}_t , denote as "Ours $\{\mathbf{W}_t$ optimizer} + $\{\mathbf{P}_t$ optimizer}"". Adaf., and Adam refer to Adafactor and 8bit-AdamW, respectively.

5.5 Can *Online Subspace Descent* Scale to Larger Model?

We pretrain from scratch a 7B LLaMA model on the C4 dataset for 10K steps, where the \mathbf{P}_t matrix is updated by SGD. The perplexity is the lower the better. The final perplexity and training wall-clock time are provided in Table 3. We further provide the downstream evaluation of the pretrained checkpoints using Galore and our method on the GLUE benchmark in Table 4. Our method consistently outperforms Galore when the model size scales up.

Method	Perplexity	Wall Clock Time (hours)
Galore	51.21	9.7439
Ours	43.72	7.1428

Table 3: Perplexity and Wall Clock Time for 7B models pretrained on C4 for 10K steps. Lower perplexity is better. Online Subspace Descent can be upto 1.3x faster than Galore.

6 Related Works

We discuss related works on memory-efficient optimization and low-rank adaptation techniques.

Method	MRPC	RTE	SST2	MNLI	QNLI	QQP	AVG
Galore	0.6838	0.5018	0.5183	0.3506	0.4946	0.3682	0.4862
Ours	0.6982	0.4901	0.5233	0.3654	0.5142	0.3795	0.4951

Table 4: Standardized GLUE evaluation for 7B model checkpoints using eval-harness. Results are reported for various downstream tasks.

Low-Rank Adaptation Low-Rank Adaptation (LoRA) [11] adds a low-rank adaptor to specific linear layers in a model, and finetune only the low-rank adaptor. As the adaptors are small, LoRA is widely applied for finetuning large models. Many variants have been proposed since LoRA, including support for multi-task learning Wang et al. [23] and further memory reductions Dettmers et al. [8]. Notably, Lialin et al. [15] proposed ReLoRA for pretraining, requiring a full-rank training warmup to match standard performance levels. It’s important to note that LoRA is fundamentally distinct from subspace descent. While subspace descent optimizes within the original model parameter space, LoRA focuses its optimization efforts within the space of the adaptors.

Memory-Efficient Optimization Several approaches aim to reduce memory costs associated with gradient statistics in adaptive optimization algorithms [21, 2, 7]. In particular, Adafactor [21] factorizes the second-order statistics by a row-column outer product and update the factorized bases on the fly, hence achieving a sub-linear memory cost. K-Fac [19] presents a factorized approximation of the Fisher information matrix which leads to a sublinear natural gradient method. More recently, Feinberg et al. [9] observes that the spectra of the Kronecker-factored gradient covariance matrix in deep learning (DL) training tasks are concentrated on a small leading eigenspace and propose to maintain a matrix preconditioner using the frequent directions sketch. However, their method requires conducting the eigendecomposition at every step, which can be costly for large models. Other than factorization methods, quantization techniques [7, 1, 24, 16] are also widely used, where the gradient (or the momentum and the preconditioner) are directly quantized to tradeoff performance for memory. Fused gradient computation method [17] have also been used to minimize memory costs during training. GaLore [25] is the most relevant work to ours. GaLore focuses on low-rank gradient structures, reducing memory costs for both first and second-order statistics. Our method can be viewed as a general extension to GaLore where we replace the infrequent SVD by a continuous subspace descent [14, 10]. As a result, our method not only provides a more general framework to study memory-efficient subspace descent, but is also more performant than GaLore in practice.

7 Conclusion

In conclusion, we provide the first convergence guarantee for arbitrary update rules of projection matrix, applicable to a range of optimizers that can be analyzed using Hamiltonian Descent, including common ones like LION, AdamW, and Adafactor. Inspired by this theoretical foundation, we introduce Dynamic Subspace Descent, a novel family of subspace descent optimizers that eschews SVD in favor of online PCA for updating projection matrix. Dynamic Subspace Descent is both flexible and minimally intrusive, and our experiments show that it achieves lower perplexity in pretraining LLaMA models (ranging from 60M to 1B parameters) on the C4 dataset compared to state-of-the-art low-rank training methods, while also closing the perplexity gap with full-rank baselines.

For future research, we propose several open and intriguing questions: (1) Are there alternative methods for updating projection matrix that could accelerate convergence? (2) What is the impact of weight decay on convergence in Dynamic Subspace Descent? (3) Can low-rank gradients and updates be combined with dynamic low-rank weights (e.g., Mixture of Experts) to further enhance training efficiency? (4) Can this method be applied to problems beyond language modeling? We hope that our work provides a strong foundation for exploring these questions.

8 Acknowledgment

The research is conducted in Statistics & AI group at UT Austin, which receives supports in part from NSF CAREER1846421, SenSE2037267, Office of Navy Research, and NSF AI Institute for Foundations of Machine Learning (IFML).

References

- [1] Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. Qsgd: Communication-efficient sgd via gradient quantization and encoding. *Advances in neural information processing systems*, 30, 2017.
- [2] Rohan Anil, Vineet Gupta, Tomer Koren, and Yoram Singer. Memory efficient adaptive optimization. *Advances in Neural Information Processing Systems*, 32, 2019.
- [3] Pierre Baldi and Kurt Hornik. Neural networks and principal component analysis: Learning from examples without local minima. *Neural networks*, 2(1):53–58, 1989.
- [4] Lizhang Chen, Bo Liu, Kaizhao Liang, and Qiang Liu. Lion secretly solves constrained optimization: As lyapunov predicts. *arXiv preprint arXiv:2310.05898*, 2023.
- [5] Xiangning Chen, Chen Liang, Da Huang, Esteban Real, Kaiyuan Wang, Yao Liu, Hieu Pham, Xuanyi Dong, Thang Luong, Cho-Jui Hsieh, et al. Symbolic discovery of optimization algorithms. *arXiv preprint arXiv:2302.06675*, 2023.
- [6] 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.
- [7] Tim Dettmers, Mike Lewis, Sam Shleifer, and Luke Zettlemoyer. 8-bit optimizers via block-wise quantization. *arXiv preprint arXiv:2110.02861*, 2021.
- [8] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36, 2024.
- [9] 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.
- [10] Guy Gur-Ari, Daniel A Roberts, and Ethan Dyer. Gradient descent happens in a tiny subspace. *arXiv preprint arXiv:1812.04754*, 2018.
- [11] 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.
- [12] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [13] David Kozak, Stephen Becker, Alireza Doostan, and Luis Tenorio. Stochastic subspace descent. *arXiv preprint arXiv:1904.01145*, 2019.
- [14] Brett W Larsen, Stanislav Fort, Nic Becker, and Surya Ganguli. How many degrees of freedom do we need to train deep networks: a loss landscape perspective. *arXiv preprint arXiv:2107.05802*, 2021.
- [15] Vladislav Lialin, Sherin Muckatira, Namrata Shivagunde, and Anna Rumshisky. Relora: High-rank training through low-rank updates. In *Workshop on Advancing Neural Network Training: Computational Efficiency, Scalability, and Resource Optimization (WANT@ NeurIPS 2023)*, 2023.
- [16] Bo Liu, Lemeng Wu, Lizhang Chen, Kaizhao Liang, Jiaxu Zhu, Chen Liang, Raghuraman Krishnamoorthi, and Qiang Liu. Communication efficient distributed training with distributed lion. *arXiv preprint arXiv:2404.00438*, 2024.
- [17] 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*, 2023.
- [18] Chris J Maddison, Daniel Paulin, Yee Whye Teh, Brendan O’Donoghue, and Arnaud Doucet. Hamiltonian descent methods. *arXiv preprint arXiv:1809.05042*, 2018.

- [19] James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International conference on machine learning*, pages 2408–2417. PMLR, 2015.
- [20] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv e-prints*, 2019.
- [21] Noam Shazeer and Mitchell Stern. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR, 2018.
- [22] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [23] Yiming Wang, Yu Lin, Xiaodong Zeng, and Guannan Zhang. Multilora: Democratizing lora for better multi-task learning. *arXiv preprint arXiv:2311.11501*, 2023.
- [24] Wei Wen, Cong Xu, Feng Yan, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Terngrad: Ternary gradients to reduce communication in distributed deep learning. *Advances in neural information processing systems*, 30, 2017.
- [25] 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*, 2024.

A Hamiltonian Examples

Example A.1. *Momentum + Online Subspace Descent is*

$$\frac{d}{dt} \mathbf{W}_t = -\mathbf{P}_t \hat{\mathbf{M}}_t, \quad \hat{\mathbf{G}}_t = \mathbf{P}_t^\top \nabla L(\mathbf{W}_t), \quad \frac{d}{dt} \hat{\mathbf{M}}_t = a(\hat{\mathbf{G}}_t - \hat{\mathbf{M}}_t),$$

with Lyapunov function $H(\mathbf{W}, \hat{\mathbf{M}}) = L(\mathbf{W}) + \frac{\|\hat{\mathbf{M}}\|^2}{2a}$, for which

$$\frac{d}{dt} H(\mathbf{W}_t, \hat{\mathbf{M}}_t) = -\nabla L(\mathbf{W}_t)^\top \mathbf{P}_t \hat{\mathbf{M}}_t + \hat{\mathbf{M}}_t^\top (\mathbf{P}_t^\top \nabla L(\mathbf{W}_t) - \hat{\mathbf{M}}_t) = -\|\hat{\mathbf{M}}_t\|_2^2 \leq 0.$$

Example A.2. *Adam + Online Subspace Descent is*

$$\frac{d}{dt} \mathbf{W}_t = \mathbf{P}_t \frac{\hat{\mathbf{M}}_t}{\sqrt{\hat{\mathbf{V}}_t + e}}, \quad \hat{\mathbf{G}}_t = \mathbf{P}_t^\top \nabla L(\mathbf{W}_t), \quad \frac{d}{dt} \hat{\mathbf{M}}_t = a(\hat{\mathbf{G}}_t - \hat{\mathbf{M}}_t), \quad \frac{d}{dt} \hat{\mathbf{V}}_t = b(\hat{\mathbf{G}}_t^2 - \hat{\mathbf{V}}_t).$$

with Lyapunov function $H(\mathbf{W}, \mathbf{M}, \mathbf{V}) = L(\mathbf{W}) + \frac{1}{2a} \left\langle \frac{\hat{\mathbf{M}}}{\sqrt{\hat{\mathbf{V}} + e}}, \hat{\mathbf{M}} \right\rangle$, for which

$$\begin{aligned} & \frac{d}{dt} H(\mathbf{W}_t, \hat{\mathbf{M}}_t, \hat{\mathbf{V}}_t) \\ &= -\left\langle \mathbf{G}_t, \mathbf{P}_t \frac{\hat{\mathbf{M}}_t}{\sqrt{\hat{\mathbf{V}}_t + e}} \right\rangle + \frac{1}{a} \left\langle \frac{\hat{\mathbf{M}}_t}{\sqrt{\hat{\mathbf{V}}_t + e}}, a(\mathbf{P}_t^\top \mathbf{G}_t - \hat{\mathbf{M}}_t) \right\rangle - \frac{b}{4a} \left\langle \frac{\hat{\mathbf{M}}_t^{\odot 2}}{\sqrt{\hat{\mathbf{V}}_t} \odot (\sqrt{\hat{\mathbf{V}}_t + e})^{\odot 2}}, (\hat{\mathbf{G}}_t^{\odot 2} - \hat{\mathbf{V}}_t) \right\rangle \\ &= -\left\langle 1 - \frac{b}{4a} \frac{\sqrt{\hat{\mathbf{V}}_t}}{\sqrt{\hat{\mathbf{V}}_t + e}}, \frac{\hat{\mathbf{M}}_t^{\odot 2}}{\sqrt{\hat{\mathbf{V}}_t + e}} \right\rangle - \frac{b}{4a} \left\langle \frac{\hat{\mathbf{M}}_t^{\odot 2}}{\sqrt{\hat{\mathbf{V}}_t} \odot (\sqrt{\hat{\mathbf{V}}_t + e})^{\odot 2}}, \hat{\mathbf{G}}_t^{\odot 2} \right\rangle \\ &\leq -\left(1 - \frac{b}{4a}\right) \left\| \frac{\hat{\mathbf{M}}_t}{\sqrt{\hat{\mathbf{V}}_t + e}} \right\|^2 - \frac{b}{4a} \left\| \frac{\hat{\mathbf{M}}_t \hat{\mathbf{G}}_t}{\sqrt[4]{\hat{\mathbf{V}}_t} (\sqrt{\hat{\mathbf{V}}_t + e})} \right\|^2 \leq 0, \end{aligned}$$

where we assume $a \geq b/4$.

Example A.3. *The Lion-K + Online Subspace Descent is*

$$\frac{d}{dt} \mathbf{W}_t = \mathbf{P}_t \nabla \mathcal{K}((1-b)\hat{\mathbf{M}}_t - b\hat{\mathbf{G}}_t), \quad \frac{d}{dt} \mathbf{M}_t = -a(\hat{\mathbf{G}}_t + \hat{\mathbf{M}}_t), \quad \hat{\mathbf{G}}_t = \mathbf{P}_t^\top \nabla L(\mathbf{W}_t)$$

Consider the Hamiltonian function in Eq (13) of [4]:

$$H(\mathbf{W}, \hat{\mathbf{M}}) = aL(\mathbf{W}) + \frac{1}{1-b} \mathcal{K}((1-b)\hat{\mathbf{M}}).$$

$$\begin{aligned} \frac{d}{dt} H(\mathbf{W}_t, \hat{\mathbf{M}}_t) &= a\langle \mathbf{G}_t, \mathbf{P}_t \nabla \mathcal{K}((1-b)\hat{\mathbf{M}}_t - b\hat{\mathbf{G}}_t) \rangle - a\langle \nabla \mathcal{K}((1-b)\hat{\mathbf{M}}_t), \hat{\mathbf{G}}_t + \hat{\mathbf{M}}_t \rangle \\ &= a\langle \hat{\mathbf{G}}_t, \nabla \mathcal{K}((1-b)\hat{\mathbf{M}}_t - b\hat{\mathbf{G}}_t) - \nabla \mathcal{K}((1-b)\hat{\mathbf{M}}_t) \rangle - a\langle \nabla \mathcal{K}((1-b)\hat{\mathbf{M}}_t), \hat{\mathbf{M}}_t \rangle \\ &= -\frac{a}{b}[(1-b)\hat{\mathbf{M}}_t; -b\hat{\mathbf{G}}_t]_{\nabla \mathcal{K}} - \frac{a}{(1-b)}[\mathbf{0}; (1-b)\hat{\mathbf{M}}_t]_{\nabla \mathcal{K}} \leq 0 \end{aligned}$$

where we defined $[\mathbf{X}; \mathbf{Y}]_{\nabla \mathcal{K}} = \langle \mathbf{Y}, \nabla \mathcal{K}(\mathbf{X} + \mathbf{Y}) - \nabla \mathcal{K}(\mathbf{X}) \rangle$ and used the fact that $[\mathbf{X}; \mathbf{Y}]_{\nabla \mathcal{K}} \geq 0$ by the convexity of \mathcal{K} ; we used $\langle \mathbf{G}_t, \mathbf{P}_t \mathbf{X}_t \rangle = \langle \mathbf{P}_t^\top \mathbf{G}_t, \mathbf{X}_t \rangle = \langle \hat{\mathbf{G}}_t, \mathbf{X}_t \rangle$.

In all examples above, although the form of Hamiltonian $H(\mathbf{W}, \hat{\mathbf{S}})$ is independent of the update rule of \mathbf{P}_t , the decreasing rate $\frac{d}{dt} H(\mathbf{W}_t, \hat{\mathbf{S}}_t)$ depends on \mathbf{P}_t in a complicate way through $\hat{\mathbf{M}}_t, \hat{\mathbf{V}}_t, \hat{\mathbf{G}}_t$. An interesting direction for future investigation is to find optimal rules of \mathbf{P}_t to maximize the decreasing rate as an optimal control problem.

B Experiments

B.1 Hyperparameters

We sweep learning rate from [0.01, 0.005, 0.001]. For GaLore as well as Adam8bit, we follow the recommended hyperparameters as much as possible. For instance, GaLore update gap is set to recommended default, 200. Warmup is set to 10% of the total training steps. Batch size is set to 512 and gradient clipping is set to 1.0.

B.2 Rank Sweep

In the following table, is a sweep on different ranks and their final perplexity of LLaMA 60M (SS = 1024) on C4. All other hyperparameters are fixed and using recommended default. Notice that as the rank increases, both Dynamic Subspace Descent and GaLore improve.

Rank	Perplexity (Ours)	Perplexity (GaLore)
32	85.90	86.16
128	49.01	48.05
512	37.41	36.93
Full	37.18	36.51

Table 5: On LLaMA 60M SS 1024, we sweep across different ranks, the trend is clear and intuitive that higher rank is preferred when it's feasible.

B.3 Optimizer Sweep

Method	Perplexity
Lion	52.65
Adafactor	33.45
Adam8bit	29.77
SGD	3469.14
Galore LION	46.90
Galore Adafactor	34.32
Galore AdamW8bit	48.05
Ours LION + LION	57.97
Ours Adafactor + Adafactor	47.61
Ours AdamW8bit + AdamW8bit	49.01
Ours LION + Adam8bit	44.76
Ours Adafactor + AdamW8bit	34.15
Ours AdamW8bit + SGD	53.53

Table 6: In this experiment, we train LLaMA 60M on C4 with sequence length of 1024. We combine different base optimizers to update both P_t and W_t , denote as "Ours weight optimizer + P_t optimizer".

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: Our theoretical and empirical results align with our main claims.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discussed the limitation and difficulty in analyzing the class of method theoretically.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: All assumptions are clearly stated in the body of the paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: **[Yes]**

Justification: Hyperparameters choices are provided and later we will make code publicly available anonymously.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: All dataset and codebase used in the experiments are opensourced.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All hyperparameters, including learning rate, rank of projectors, sequence length are specified in each experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer:[No]

Justification: Training LLM is expensive, we can afford running multiple runs on large models. However, previous empirical studies have shown that variation across seed is generally small in language model pretraining.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: GPU type used in experiment is specified.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: We comply with the Code of Ethnics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[NA\]](#)

Justification: Our paper aims to promote deeper theoretical understanding of a class of method. The algorithm itself is not practical enough to produce powerful enough LLM that could potentially have any negative societal impact.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: same reasoning as above

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Everything is opensourced.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new asset has been introduced.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No human subject has been involved.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We did not perform any studies involving human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.