
GeoNorm: Unify Pre-Norm and Post-Norm with Geodesic Optimization

Chuanyang Zheng¹ Jiankai Sun² Yihang Gao³ Chi Wang⁴ Yuehao Wang⁵ Jing Xiong⁶ Liliang Ren⁷
Bo Peng⁴ Qingmei Wang⁸ Xiaoran Shang⁹ Mac Schwager² Anderson Schneider¹ Yuriy Nevmyvaka¹
Xiaodong Liu⁷

Abstract

The placement of normalization layers, specifically Pre-Norm and Post-Norm, remains an open question in Transformer architecture design. In this work, we rethink these approaches through the lens of manifold optimization, interpreting the outputs of the Feed-Forward Network (FFN) and attention layers as update directions in optimization. Building on this perspective, we introduce GeoNorm, a novel method that replaces standard normalization with geodesic updates on the manifold. Furthermore, analogous to learning rate schedules, we propose a layer-wise update decay for the FFN and attention components. Comprehensive experiments demonstrate that GeoNorm consistently outperforms existing normalization methods in Transformer models. Crucially, GeoNorm can be seamlessly integrated into standard Transformer architectures, achieving performance improvements with negligible additional computational cost.

1. Introduction

Recent years have witnessed remarkable progress in Large Language Models (LLMs) (Brown et al., 2020; Ouyang et al., 2022; Touvron et al., 2023), driven primarily by the exponential growth of training data and model parameters. This advancement is reflected in substantial improvements across language modeling (Fedus et al., 2022; Puigcerver et al., 2024; Jiang et al., 2024; Meta; Liu et al., 2024; Team et al., 2025) and computer vision (Riquelme et al., 2021; Lin et al., 2024). Within the Transformer architecture, the choice of normalization function, such as Pre-Norm and Post-Norm, is irrelevant, as it has no impact on training stability, model performance, or gradient flow.

Currently, Pre-Norm has become the predominant choice

¹Morgan Stanley ²Stanford ³NUS ⁴Google ⁵UTA ⁶HKU
⁷Microsoft ⁸RUC ⁹USTC. Correspondence to: Chuanyang Zheng <cyzhengme@gmail.com>.

in modern LLMs due to its favorable stability properties (Liu et al., 2024; Yang et al., 2025; Jiang et al., 2024). In contrast, Post-Norm (Popel & Bojar, 2018; Shazeer & Stern, 2018), which applies normalization after the addition of the residual connection, was originally introduced in the Transformer and facilitates the training of deeper networks. However, it is often prone to training instability and loss spikes (Wang et al., 2024). Pre-Norm was subsequently proposed to improve stability by applying layer normalization to the residual branch before the addition operation (Nguyen & Salazar, 2019). Nonetheless, Pre-Norm introduces imbalanced gradients, with lower layers tending to receive larger updates than higher ones (Wang et al., 2024). To address these limitations, recent methods have sought to enhance stability and enable extreme model scaling. DeepNorm (Wang et al., 2024) stabilizes training and improves performance through better parameter initialization and residual scaling. Meanwhile, SandwichNorm (Ding et al., 2021; Yin et al., 2025) introduces additional normalization layers at the network’s input and output. While these approaches demonstrate strong empirical results, they generally lack a unified theoretical framework for analyzing normalization functions in Transformers.

In this work, we interpret the operations $x + \text{FFN}(x)$ and $x + \text{Attention}(x)$ as noise injections steps on a sphere, where $\text{FFN}(x)$ and $\text{Attention}(x)$ are unrelated to any gradient direction. From this perspective, each Transformer layer can be viewed as performing an iterative optimization step within a dynamical system: token embeddings act as the initial state, while the attention and feed-forward modules generate update directions based on the current iterate. The normalization operator then projects the updated vector back onto a constraint set, specifically, the sphere defined by the normalization radius. Building on this formulation, our contributions are threefold:

- **Theoretical Framework:** We introduce a theoretical framework for analyzing normalization in Transformers. Here, x is treated as the current point, while $\text{FFN}(x)$ and $\text{Attention}(x)$ are interpreted as the update direction in optimization.

- **GeoNorm Normalization:** We propose GeoNorm, a novel normalization function that replaces conventional schemes by leveraging geodesic and Riemannian optimization on the sphere, leading to improved performance.
- **Extensive Empirical Validation:** Through comprehensive experiments across varying training lengths, datasets, and model sizes, we validate the effectiveness of GeoNorm. Additionally, we demonstrate its scalability and generalization on large-scale datasets with downstream task evaluation.

2. Related Work

Normalization Function The original Post-Norm architecture, first introduced in the Transformer model (Popel & Bojar, 2018; Shazeer & Stern, 2018), applies layer normalization to the sum of the identity term and the residual output. While Post-Norm facilitates the training of deep networks, it can suffer from instability and exhibit sharp loss spikes (Wang et al., 2024). In response, the Pre-Norm variant (Nguyen & Salazar, 2019) modifies the residual connection by applying layer normalization directly to the identity term before the addition, leading to improved training stability. However, this often results in imbalanced gradient magnitudes, with lower layers receiving larger gradients compared to higher layers (Wang et al., 2024). To further enhance stability and enable scaling to significantly greater depths, DeepNorm (Wang et al., 2024) introduces refined parameter initialization and residual scaling. Another approach, StandwichNorm (Ding et al., 2021; Yin et al., 2025), incorporates normalization layers at both the network input and output alongside the identity connection. Although these methods have empirically demonstrated strong performance, they currently lack a comprehensive theoretical framework for analyzing the underlying normalization mechanisms.

Gradient Descent. The stochastic gradient descent (SGD) (Bottou, 1998) is an iterative optimization method with the gradient. Based on the SGD, there is SGD with momentum (Rumelhart et al., 1986) which use the momentum of previous gradient. Adagrad (adaptive gradient algorithm) (Duchi et al., 2011) is a modified stochastic gradient descent algorithm with per-parameter learning rate. The RMSProp (Graves, 2013) multiplies the learning rate for a weight by a running average of the magnitudes of recent gradients for that weight. Adam (Kingma & Ba, 2017) run averages with exponential forgetting of both the gradients and the second moments of the gradients are used.

Optimization Method. The Geodesic optimization is related to the optimization on the manifold (Kingma & Ba, 2017). Riemannian Steepest Descent Method extends the

classic gradient descent to optimize functions on curved spaces (manifolds) (Debye, 1909) by following the path of steepest decrease, using the Riemannian gradient (the projection of the standard gradient onto the manifold) and a retraction to map steps back onto the manifold. Riemannian Newton’s method (da Silva et al., 2025) adapts classical Newton’s method for optimization problems on curved spaces (Riemannian manifolds). The Riemannian Conjugate Gradient method (Hestenes et al., 1952) extends the conjugate gradient method to a manifold. For the retraction relation method, they will calculate the gradient and map the result back to the manifold. The geodesic-based methods (Zhang & Sra, 2016) use the gradient descent method on the geodesic of the manifold.

3. Method

3.1. Post-Normalization in LLMs

Beyond the well-recognized attention and feed-forward network modules, modern LLMs rely critically on normalization to stabilize training and inference as depth increases. Let \mathbf{x}_k denote the token representation after the k -th Transformer layer. A Transformer layer with post-normalization can be written as

$$\begin{aligned}\tilde{\mathbf{x}}_k &= \text{Norm}(\mathbf{x}_k + \text{Attn}(\mathbf{x}_k)), \\ \mathbf{x}_{k+1} &= \text{Norm}(\tilde{\mathbf{x}}_k + \text{FFN}(\tilde{\mathbf{x}}_k)),\end{aligned}\tag{1}$$

where $\text{Attn}(\cdot)$ and $\text{FFN}(\cdot)$ denote the attention and feed-forward modules, and $\text{Norm}(\cdot)$ is the normalization operator (typically RMSNorm followed by rescaling to preserve the vector norm).

From this formulation, we may view the Transformer layer as performing an iterative optimization step in a dynamical system. The token embeddings serve as the initial state, while the attention and feed-forward modules produce update directions based on the current iterate \mathbf{x}_k . The normalization operator then projects the updated vector back to a constraint set, specifically, the sphere defined by the normalization radius $\|\mathbf{x}_k\|$. This interpretation closely mirrors the projected optimization scheme:

$$\mathbf{x}_{k+1} = \text{Proj}_{\Omega}(\mathbf{x}_k + \alpha_k \mathbf{s}_k),\tag{2}$$

where \mathbf{s}_k is an update step determined by past iterates $\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k\}$, α_k represents the step size, and $\text{Proj}_{\Omega}(\cdot)$ denotes projection onto the feasible region Ω . Equation (2) includes a wide range of classical optimization algorithms such as projected gradient descent.

From this perspective, (1) corresponds to the special case where the feasible region Ω is the sphere in the Euclidean space induced by ℓ_2 -norm, and the projection $\text{Proj}_{\Omega}(\cdot)$ is simply vector normalization. Thus, each Transformer layer applies an unconstrained update followed by a projection

back to the sphere, aligning post-normalization Transformers with the general class of optimization algorithms that incorporate projection to enforce constraints. This connection motivates replacing the projection step with a more geometrically informed operation, such as the exponential mapping, in the subsequent section.

3.2. Geodesic-Inspired Transformer Design

Geodesic-based optimization methods are known to possess several structural advantages over projected gradient approaches, particularly when the feasible region Ω is a smooth manifold (e.g., spheres, Stiefel/orthogonal matrices, low-rank matrices, or fixed-rank positive semi-definite matrices). Traditional projected methods handle constraints by alternating between Euclidean optimization and projection back onto the feasible set. While simple to implement and interpret, this extrinsic approach discards geometric information that the projection step can significantly distort step s_k , introduce unexpected kinks in the optimization trajectory, and degrade convergence due to repeated violations and corrections of feasibility. These issues are amplified when the feasible region is a smooth manifold, where the curvature causes projected steps to deviate substantially from the true directions. In contrast, geodesic-based optimization circumvents these limitations by performing updates intrinsically on the manifold. By following geodesics or retraction paths, each iterate remains feasible without projection, and the search direction corresponds to the Riemannian step under the manifold metric. As a consequence, geodesic-based methods typically exhibit smoother optimization trajectories, improved conditioning, and stronger convergence guarantees compared to their projected counterparts. These observations motivate our modification of the Transformer model from projection-based schemes in (1) to a geodesic-based formulation.

In geodesic optimization, the core idea is to modify the step direction s_k in a manner consistent with the geometry of the manifold Ω , ensuring that every iterate remains on Ω . This is achieved through the exponential map $\exp_x(\cdot)$. Consider a smooth manifold Ω with tangent space $T_x\Omega := \{v : x^\top v = 0\}$ at point x . For any tangent vector $v \in T_x\Omega$, there exists a unique geodesic $\gamma_v(t)$ satisfying

$$\gamma_v(0) = x, \quad \gamma'_v(0) = v.$$

The exponential map at x is then defined as $\exp_x(v) = \gamma_v(1)$. We refer readers to Riemannian optimization literatures (Zhang & Sra, 2016; da Silva et al., 2025) for formal definitions and properties of exponential mappings. For a general manifold Ω , the exponential map is implicit and rarely available in a closed form, which limits the practical adoption of geodesic-based methods due to the computational cost of evaluating $\exp_x(v)$. However, in our setting (1), the manifold is the sphere, where the exponential map

is explicit and computationally efficient.

Let Ω be a sphere embedded in Euclidean space with the standard induced Riemannian metric. Given a point $x \in \Omega$ and a tangent vector $v \in T_x\Omega$, the exponential map admits the closed-form expression

$$\exp_x(v) = \cos\left(\frac{\|v\|}{\|x\|}\right)x + \|x\| \sin\left(\frac{\|v\|}{\|x\|}\right) \frac{v}{\|v\|}. \quad (3)$$

Given an update direction s for x , we first project it to the tangent space:

$$v = s - \frac{x^\top s}{\|x\|^2}x,$$

and then update the representation intrinsically via the exponential map $x + \exp_x(v)$. Specifically, (2) is modified to

$$\begin{aligned} v_k &= s_k - \frac{x_k^\top s_k}{\|x_k\|^2}x_k, \\ x_{k+1} &= \exp_{x_k}(\alpha_k v_k). \end{aligned} \quad (4)$$

Exponential mapping interprets the update as moving along the geodesic determined by the update direction, producing a smooth, norm-preserving, and curvature-informed transformation between states.

These observations motivate us to replace the projection-based normalization in (1) with a geodesic-inspired normalization based on the exponential map (3). More specifically, we reformulate the post-normalized Transformer layer (1) as

$$\begin{aligned} s_k &= \text{Attn}(x_k), \\ v_k &= s_k - \frac{x_k^\top s_k}{\|x_k\|^2}x_k, \\ \tilde{x}_k &= \exp_{x_k}(\alpha_k v_k), \\ \tilde{s}_k &= \text{FFN}(\tilde{x}_k), \\ \tilde{v}_k &= \tilde{s}_k - \frac{\tilde{x}_k^\top \tilde{s}_k}{\|\tilde{x}_k\|^2}\tilde{x}_k, \\ x_{k+1} &= \exp_{\tilde{x}_k}(\alpha_k \tilde{v}_k), \end{aligned} \quad (5)$$

where the update directions s_k and \tilde{s}_k are produced by the attention and feed-forward modules, v_k and \tilde{v}_k are their tangent-space projections, and $\{\alpha_k\}$ is a prescribed sequence of step sizes. In contrast to the standard post-normalization, which abruptly projects the updated representation back to the sphere, the update in (5) moves intrinsically along the geodesic determined by the residual direction.

Step-size selection is an important component of geodesic optimization. Classical choices include constant step sizes for geodesically convex and smooth objectives ($\alpha_k \leq \frac{2}{L}$, where L is the gradient Lipschitz constant), polynomially

diminishing step sizes for stochastic Riemannian optimization, and Barzilai–Borwein (BB) step sizes for accelerated variants. Since the token representations in Transformer layers evolve under stochastic transformations, arising from dropout, token generation sampling, autoregressive randomness, and the non-deterministic interaction of multiple tokens, we model the layerwise dynamics as a stochastic Riemannian optimization process. Accordingly, we adopt a polynomially decaying step size of the form $\alpha_k = \frac{\alpha}{k^{0.5}}$ (named as **sqrt**), which theoretically achieves the $O(\frac{\log k}{k})$ convergence rate for stochastic Riemannian gradient methods (da Silva et al., 2025). Also we could also have $\alpha_k = \frac{\alpha}{k}$ (named as **harmonic**) and $\alpha_k = \frac{\alpha(T-K)}{T}$ (named as **linear**, and T is the total layer number).

In the proposed method (5), we keep the attention and FFN modules unchanged and modify only the post-normalization step. Instead of applying a projection via the normalization operator $\text{Norm}(\cdot)$, we adopt the exponential map inspired by geodesic Riemannian optimization. Our intuition is that projection-based normalization may destroy certain geometric structure in the update directions s_k and \tilde{s}_k generated by the attention and FFN modules. Projection forces the representation back onto the sphere by discarding the radial component abruptly, potentially altering the intended direction of motion. In contrast, the exponential map preserves the intrinsic geometry of the sphere: it transports the tangent vector faithfully along the geodesic without the directional distortions introduced by renormalization. This allows the model to follow the spherical manifold structure naturally, maintaining richer geometric relationships between the current representation x_k , the attention update s_k , the FFN update \tilde{s}_k , and the curvature of the underlying feature space. Comparing the classical Transformer update (1) with our geodesic formulation (5), we note that the exponential map adds negligible computational overhead. Its closed-form expression involves only basic vector operations and introduces no additional trainable parameters. Consequently, our method integrates seamlessly into existing LLM architectures without modifying the core Transformer blocks, and is computationally efficient.

3.3. Pre-Norm with GeoNorm formulation

Moreover, we find that Pre-Norm can be regarded as a special case of GeoNorm (as shown in Appendix B), as it follows an update scheme similar to (5), with the corresponding mapping in (3) but with non-adaptive angles. This interpretation is further supported by empirical results, where the two formulations exhibit comparable performance.

In summary, we interpret Transformers as Riemannian flows of token embeddings evolving on a spherical manifold. Replacing projection with geodesic-informed exponential updates in (5) enables representations to move smoothly

along circles on the sphere, leading to more coherent and curvature-aware information propagation across layers.

4. Experiment

Baseline. We compare the proposed GeoNorm with the Dense model and normalization. To be specific, we evaluate GeoNorm against normalization methods, including Pre-Norm, Post-Norm, DeepNorm and SandwichNorm. **The implementation of GeoNorm is in Appendix D.**

Datasets. Our analysis involves training language models on the Arxiv and Books3 datasets, which are frequently used benchmarks for evaluating model performance. Moreover, we train the model on the large-scale dataset FinWeb-Edu (Lozhkov et al., 2024) and evaluate on downstream datasets,

Experiment settings. Initially, we compare GeoNorm with other baselines at training lengths 512 and 1024, using decoder-only Transformers (Brown et al., 2020) with model size 125M, whose configuration is shown in Appendix A. Subsequently, we evaluate the performance of larger model sizes, specifically 350M and 1.3 B.

4.1. Compare with baseline

GeoNorm achieves better performance, with different datasets. As shown in Figure 1, to evaluate the generalization capability of the proposed method, we conducted experiments across two distinct datasets, Arxiv and Books3, using a fixed training length of 512. The results demonstrate a clear and consistent performance hierarchy. On the Arxiv dataset, GeoNorm establishes a strong lead with a final loss of 1.8792. This outperforms the closest baseline, Pre-Norm (1.9032), as well as Post-Norm (1.9240), DeepNorm (1.9277), and SandwichNorm (1.9126). Crucially, this pattern of superiority is not an artifact of a single data distribution. It is robustly replicated on the Books3 dataset, where GeoNorm again achieves the optimal loss of 3.4040, maintaining a decisive margin over all alternative normalization methods. The fact that GeoNorm secures the top position on both datasets, which represent different textual domains and complexities, provides compelling evidence that its architectural advantage is not dataset-specific. Therefore, we conclude that GeoNorm offers generalizable improvements, delivering reliably better performance across diverse data environments.

GeoNorm achieves better performance, with different training lengths. The performance of GeoNorm demonstrates consistent superiority across varying training lengths on the Books3 dataset. At a training length of 512, GeoNorm establishes an immediate lead, achieving an optimal loss of 3.4040. This already surpasses the results of all

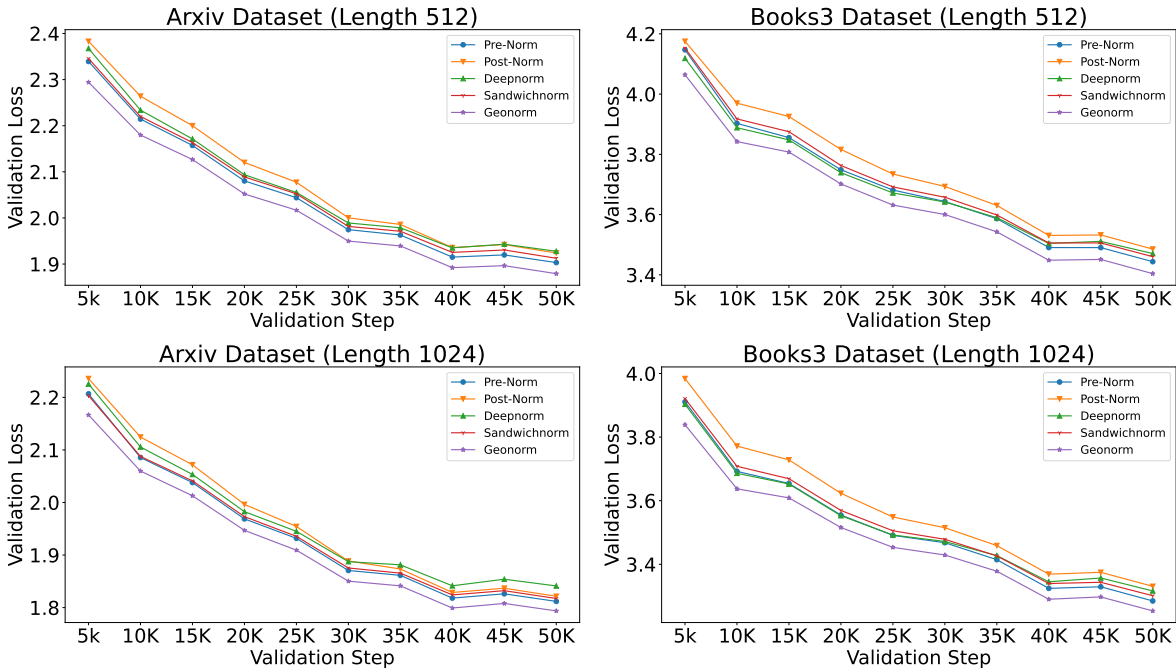


Figure 1. The performance of different methods on the Arxiv and Books3 dataset, with model parameter 125M.

comparative baselines, including Pre-Norm (3.4437), Post-Norm (3.4855), DeepNorm (3.4708), and SandwichNorm (3.4604). Crucially, this initial advantage is not merely maintained but is significantly amplified when the training duration is extended. At the longer training length of 1024, GeoNorm again records the best possible loss of 3.2534. More importantly, the performance gap separating it from the nearest competitor widens considerably at this scale; while Pre-Norm achieves a loss of 3.2847, the other methods fall even further behind. Therefore, the superiority of GeoNorm is both robust and generalizes effectively across key experimental parameters.

GeoNorm is consistently better across the validation steps. The evaluation of training dynamics reveals that GeoNorm maintains a consistent performance advantage throughout the entire validation schedule. At the initial validation step (5000), GeoNorm establishes an early lead with a loss of 4.0642, which is superior to all baseline methods. This advantage is not transient; from step 5000 through 50000, GeoNorm continues to achieve the lowest loss at every measured interval, demonstrating stable and sustained optimization. The fact that GeoNorm leads at the beginning, maintains the lead throughout, and finishes with the best performance provides strong evidence that its architectural benefits are robust across all training phases. Therefore, we conclude that GeoNorm consistently delivers better performance, showcasing reliable convergence and optimization stability.

4.2. Performance on Large Model

The GeoNorm achieves the best performance with larger model size. As shown in Figure 2, the performance of GeoNorm demonstrates consistent superiority across increasing model sizes. For a 350M parameter model, it achieves the best loss of 3.2837, compared to 3.3233 (Pre-Norm), 3.3775 (DeepNorm), and 3.3598 (SandwichNorm). Scaling to 1.3B parameters, GeoNorm again leads with a loss of 3.1759, while the next best baseline, Pre-Norm, achieves only 3.2067. These results confirm two key points: first, GeoNorm delivers the best performance at every tested scale, and second, its performance advantage persists and even improves relative to alternatives as model size grows, showcasing its scalability.

GeoNorm consistently achieves the best performance across the validation steps. GeoNorm also shows stable performance throughout training, maintaining the best validation loss at multiple checkpoints. At validation step 5000, the 350M GeoNorm model achieves the lowest loss, and this advantage persists at step 50000. The same pattern holds for the 1.3B model, which outperforms baselines from early to late validation steps. These results confirm that GeoNorm’s performance advantages are robust across different training stages.




Figure 2. The performance of different baselines on the Books3 dataset, with model parameter 350M and 1.3B, training lengths of 512.




Figure 3. The performance of different baselines on the Books3 dataset, with model size 125M and training length 2048 and 4096.

4.3. Performance on Long Training Length

The GeoNorm keeps the best performance with longer training length. As shown in Figure 3, our results demonstrate the consistent superiority of the GeoNorm method across different training lengths. At a length of 2048, it achieves the best performance with a loss of 3.1538, outperforming Pre-Norm, DeepNorm, and SandwichNorm. Extending the training length to 4096 further improves its performance to a loss of 3.0908, while maintaining its lead over all baseline methods. This dual finding—performance improvement with extended training and sustained superiority—confirms that GeoNorm not only provides the best initial results but also retains its optimal performance advantage as training progresses.

The training of GeoNorm is stable, while others may have a loss spike. With a training length of 4096, the post-norm presents the loss spike and ruins the model training. On the other hand, the GeoNorm training is relatively stable. Therefore, GeoNorm training is relatively stable, while others (e.g. Post-Norm) may have a loss spike and the training process will collapse.

4.4. The Effect of decay method

The GeoNorm with Harmonic Decay achieves better performance. As shown in Figure 4, the proposed GeoNorm

method demonstrates significantly enhanced performance when utilizing the harmonic decay scheduling mechanism. Across multiple benchmark datasets, this configuration consistently yields state-of-the-art results. Specifically, on the Arxiv dataset, GeoNorm with harmonic decay achieves a loss of 1.8792, establishing a new performance benchmark. This represents a substantial improvement over alternative configurations, with the second-best performing method attaining only 1.8817. The performance improvement is even more pronounced on the larger and more complex Books3 dataset, where GeoNorm with harmonic decay records a loss of 3.4048, again outperforming all competing approaches. These empirical results collectively demonstrate that integrating harmonic decay scheduling with the GeoNorm improves the performance.

The Harmonic Decay achieves the best performance across the validation steps. On the Arxiv dataset, with the validation step 5000, the GeoNorm achieves 2.2943 loss, which is the best performance, compared to Pre-Norm, DeepNorm and SandwichNorm. With the validation step increased to 50000, the GeoNorm still achieves the lowest loss. Similarly, the GeoNorm achieves the best performance on Books3 dataset from the validation step 5000 to 50000. Therefore, the GeoNorm achieves the best performance across the validation steps.




Figure 4. The performance of different decay methods on the Arxiv and Books3 dataset, with model parameter 125M, training lengths of 512. **Sqrt**: $\frac{\alpha}{k^{0.5}}$, **Linear**: $\frac{\alpha(T-k)}{T}$, and **Harmonic**: $\frac{\alpha}{k}$, where k is the current layer index and T is the total layer number.

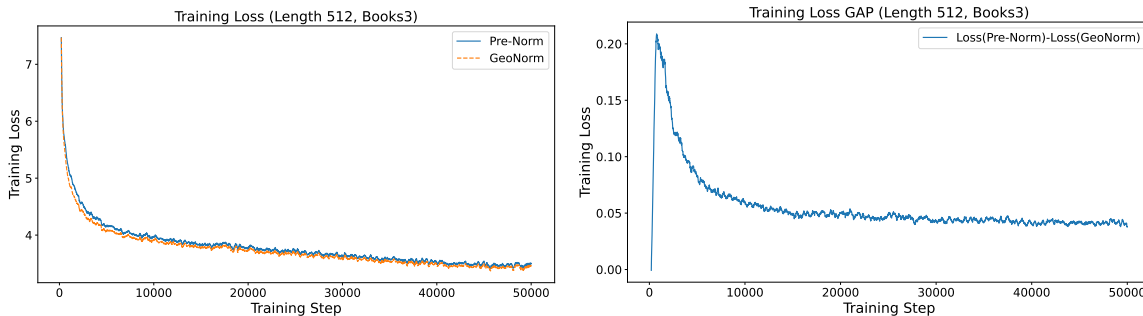


Figure 5. The training dynamic with loss metrics, with model parameter of 125M, training lengths of 512 and the Books3 dataset.

4.5. The Analysis of Training Dynamics

The GeoNorm achieves better performance across the training steps. As shown in Figure 5, the GeoNorm method establishes a decisive performance advantage early in the training phase, as evidenced by its significantly lower loss metrics. This initial lead is not merely temporary; GeoNorm demonstrates remarkable stability by maintaining its optimal performance throughout the entire training cycle without degradation. In direct comparison, the Pre-Norm baseline is consistently outperformed, confirming that GeoNorm’s superiority is a persistent and defining characteristic of the training process from initialization to convergence.

The GeoNorm has faster loss reduction. The training dynamics reveal a significant loss gap with a peak value of 0.2. Crucially, this gap does not persist due to instability but consistently narrows as training progresses, converging to a much smaller final value of approximately 0.05 by the end of the run. This relatively large divergence indicates that the GeoNorm method successfully reduces its loss at a markedly faster rate during the training phase compared to the baseline, suggesting more efficient optimization.

Broader Impact Statement

This work focuses on improving the Transformer architecture, which may be helpful for better performance of the Transformer and further understanding of the Transformer. And it should be noted that this work should not be abused.

5. Conclusion

This work presents a unified view of Pre-Norm and Post-Norm through the lens of manifold optimization. Building on this formulation, we introduce GeoNorm, a novel approach that leverages geodesic optimization to improve model performance. We conduct a comprehensive evaluation of GeoNorm, testing it across multiple datasets, varying sequence lengths, different model sizes, and a range of downstream tasks. By advancing the theoretical understanding of normalization mechanisms, this research provides a foundation for subsequent innovations.

References

Ainslie, J., Lee-Thorp, J., De Jong, M., Zemlyanskiy, Y., Lebrón, F., and Sanghai, S. Gqa: Training generalized multi-query transformer models from multi-head checkpoints. *arXiv preprint arXiv:2305.13245*, 2023.

Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization.

- 2016.
- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Bapna, A., Chen, M. X., Firat, O., Cao, Y., and Wu, Y. Training deeper neural machine translation models with transparent attention. *arXiv preprint arXiv:1808.07561*, 2018.
- Bisk, Y., Zellers, R., Gao, J., Choi, Y., et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 7432–7439, 2020.
- Bjorck, N., Gomes, C. P., Selman, B., and Weinberger, K. Q. Understanding batch normalization. *Advances in neural information processing systems*, 31, 2018.
- Bottou, L. Online algorithms and stochastic approximations. *Online learning in neural networks*, 1998.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901, 2020.
- Chai, Y., Jin, S., and Hou, X. Highway transformer: Self-gating enhanced self-attentive networks. *arXiv preprint arXiv:2004.08178*, 2020.
- Chollet, F. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1251–1258, 2017.
- Clark, C., Lee, K., Chang, M.-W., Kwiatkowski, T., Collins, M., and Toutanova, K. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*, 2019.
- Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- da Silva, C. O., Aoto, Y. A., Costa, F. F., and da Silva, M. F. A thorough study of riemannian newton’s method. *arXiv preprint arXiv:2506.09297*, 2025.
- Debye, P. Näherungsformeln für die zylinderfunktionen für große werte des arguments und unbeschränkt veränderliche werte des index. *Mathematische Annalen*, 67(4):535–558, 1909.
- Dehghani, M., Djolonga, J., Mustafa, B., Padlewski, P., Heek, J., Gilmer, J., Steiner, A. P., Caron, M., Geirhos, R., Alabdulmohsin, I., et al. Scaling vision transformers to 22 billion parameters. In *International conference on machine learning*, pp. 7480–7512. PMLR, 2023.
- Ding, M., Yang, Z., Hong, W., Zheng, W., Zhou, C., Yin, D., Lin, J., Zou, X., Shao, Z., Yang, H., et al. Cogview: Mastering text-to-image generation via transformers. *Advances in neural information processing systems*, 34: 19822–19835, 2021.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.
- Fang, Y., Cai, Y., Chen, J., Zhao, J., Tian, G., and Li, G. Cross-layer retrospective retrieving via layer attention. *arXiv preprint arXiv:2302.03985*, 2023.
- Fedus, W., Zoph, B., and Shazeer, N. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022.
- Gao, L., Tow, J., Abbasi, B., Biderman, S., Black, S., DiPofi, A., Foster, C., Golding, L., Hsu, J., Le Noac’h, A., Li, H., McDonnell, K., Muennighoff, N., Ociepa, C., Phang, J., Reynolds, L., Schoelkopf, H., Skowron, A., Sutawika, L., Tang, E., Thite, A., Wang, B., Wang, K., and Zou, A. The language model evaluation harness, 07 2024. URL <https://zenodo.org/records/12608602>.
- Graves, A. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.
- Gromov, A., Tirumala, K., Shapourian, H., Glorioso, P., and Roberts, D. A. The unreasonable ineffectiveness of the deeper layers. *arXiv preprint arXiv:2403.17887*, 2024.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Henry, A., Dachapally, P. R., Pawar, S. S., and Chen, Y. Query-key normalization for transformers. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 4246–4253, 2020.
- Hestenes, M. R., Stiefel, E., et al. Methods of conjugate gradients for solving linear systems. *Journal of research of the National Bureau of Standards*, 49(6):409–436, 1952.

- Hu, J., Liu, X., Wen, Z.-W., and Yuan, Y.-X. A brief introduction to manifold optimization. *Journal of the Operations Research Society of China*, 8(2):199–248, 2020.
- Hu, Z., Wang, L., Lan, Y., Xu, W., Lim, E.-P., Bing, L., Xu, X., Poria, S., and Lee, R. Llm-adapters: An adapter family for parameter-efficient fine-tuning of large language models. In *Proceedings of the 2023 conference on empirical methods in natural language processing*, pp. 5254–5276, 2023.
- Huang, X. S., Perez, F., Ba, J., and Volkovs, M. Improving transformer optimization through better initialization. In *International Conference on Machine Learning*, pp. 4475–4483. PMLR, 2020.
- Jaiswal, A., Yin, L., Zhang, Z., Liu, S., Zhao, J., Tian, Y., and Wang, Z. From galore to welore: How low-rank weights non-uniformly emerge from low-rank gradients. *arXiv preprint arXiv:2407.11239*, 2024.
- Jiang, A. Q., Sablayrolles, A., Roux, A., Mensch, A., Savary, B., Bamford, C., Chaplot, D. S., Casas, D. d. l., Hanna, E. B., Bressand, F., et al. Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.
- Kedia, A., Zaidi, M. A., Khyalia, S., Jung, J., Goka, H., and Lee, H. Transformers get stable: An end-to-end signal propagation theory for language models. *arXiv preprint arXiv:2403.09635*, 2024.
- Kim, J., Lee, B., Park, C., Oh, Y., Kim, B., Yoo, T., Shin, S., Han, D., Shin, J., and Yoo, K. M. Peri-In: Revisiting normalization layer in the transformer architecture. *arXiv preprint arXiv:2502.02732*, 2025.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization, 2017. URL <https://arxiv.org/abs/1412.6980>.
- Klein, G., Kim, Y., Deng, Y., Senellart, J., and Rush, A. M. Opennmt: Open-source toolkit for neural machine translation. *arXiv preprint arXiv:1701.02810*, 2017.
- Konstantinidis, D., Papastratis, I., Dimitropoulos, K., and Daras, P. Multi-manifold attention for vision transformers. *IEEE Access*, 11:123433–123444, 2023.
- Larsson, G., Maire, M., and Shakhnarovich, G. Fractalnet: Ultra-deep neural networks without residuals. *arXiv preprint arXiv:1605.07648*, 2016.
- Li, B., Wang, Z., Liu, H., Jiang, Y., Du, Q., Xiao, T., Wang, H., and Zhu, J. Shallow-to-deep training for neural machine translation. *arXiv preprint arXiv:2010.03737*, 2020.
- Li, P., Yin, L., and Liu, S. Mix-In: Unleashing the power of deeper layers by combining pre-In and post-In. *arXiv preprint arXiv:2412.13795*, 2024.
- Lialin, V., Muckatira, S., Shivagunde, N., and Rumshisky, A. Stack more layers differently: High-rank training through low-rank updates. 2023a.
- Lialin, V., Shivagunde, N., Muckatira, S., and Rumshisky, A. Relora: High-rank training through low-rank updates. *arXiv preprint arXiv:2307.05695*, 2023b.
- Lin, B., Ye, Y., Zhu, B., Cui, J., Ning, M., Jin, P., and Yuan, L. Video-LLaVA: Learning united visual representation by alignment before projection. In Al-Onaizan, Y., Bansal, M., and Chen, Y.-N. (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 5971–5984, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.342. URL <https://aclanthology.org/2024.emnlp-main.342/>.
- Liu, A., Feng, B., Xue, B., Wang, B., Wu, B., Lu, C., Zhao, C., Deng, C., Zhang, C., Ruan, C., et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024.
- Liu, L., Liu, X., Gao, J., Chen, W., and Han, J. Understanding the difficulty of training transformers. *arXiv preprint arXiv:2004.08249*, 2020a.
- Liu, S., Chen, T., Chen, X., Chen, X., Xiao, Q., Wu, B., Kärkkäinen, T., Pechenizkiy, M., Mocuano, D., and Wang, Z. More convnets in the 2020s: Scaling up kernels beyond 51x51 using sparsity. *arXiv preprint arXiv:2207.03620*, 2022a.
- Liu, X., Duh, K., Liu, L., and Gao, J. Very deep transformers for neural machine translation. *arXiv preprint arXiv:2008.07772*, 2020b.
- Liu, Z., Mao, H., Wu, C.-Y., Feichtenhofer, C., Darrell, T., and Xie, S. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11976–11986, 2022b.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Loshchilov, I., Hsieh, C.-P., Sun, S., and Ginsburg, B. ngpt: Normalized transformer with representation learning on the hypersphere. *arXiv preprint arXiv:2410.01131*, 2024.
- Lozhkov, A., Ben Allal, L., von Werra, L., and Wolf, T. Fineweb-edu: the finest collection of educational content, 2024. URL <https://huggingface.co/datasets/HuggingFaceFW/fineweb-edu>.
- Mak, B. and Flanagan, J. Residual matrix transformers: Scaling the size of the residual stream. *arXiv preprint arXiv:2506.22696*, 2025.

- Men, X., Xu, M., Zhang, Q., Yuan, Q., Wang, B., Lin, H., Lu, Y., Han, X., and Chen, W. Shortgpt: Layers in large language models are more redundant than you expect. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 20192–20204, 2025.
- Menary, S., Kaski, S., and Freitas, A. Transformer normalization layers and the independence of semantic subspaces. *arXiv preprint arXiv:2406.17837*, 2024.
- Menghani, G., Kumar, R., and Kumar, S. Laurel: Learned augmented residual layer. *arXiv preprint arXiv:2411.07501*, 2024.
- Meta, A. The llama 4 herd: The beginning of a new era of natively multimodal ai innovation. <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>. Accessed: 4-7-2025.
- Mihaylov, T., Clark, P., Khot, T., and Sabharwal, A. Can a suit of armor conduct electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*, 2018.
- Muennighoff, N., Soldaini, L., Groeneveld, D., Lo, K., Morrison, J., Min, S., Shi, W., Walsh, P., Tafjord, O., Lambert, N., et al. Olmoe: Open mixture-of-experts language models. *arXiv preprint arXiv:2409.02060*, 2024.
- Nguyen, T. Q. and Salazar, J. Transformers without tears: Improving the normalization of self-attention. In *Proceedings of the 16th International Conference on Spoken Language Translation*, 2019.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Pagliardini, M., Mohtashami, A., Fleuret, F., and Jaggi, M. Denseformer: Enhancing information flow in transformers via depth weighted averaging. *Advances in neural information processing systems*, 37:136479–136508, 2024.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- Popel, M. and Bojar, O. Training tips for the transformer model. *arXiv preprint arXiv:1804.00247*, 2018.
- Puigcerver, J., Ruiz, C. R., Mustafa, B., and Houlsby, N. From sparse to soft mixtures of experts. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=jxpsAj7ltE>.
- Riquelme, C., Puigcerver, J., Mustafa, B., Neumann, M., Jenatton, R., Susano Pinto, A., Keyzers, D., and Houlsby, N. Scaling vision with sparse mixture of experts. *Advances in Neural Information Processing Systems*, 34: 8583–8595, 2021.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.
- Rybakov, O., Chrzanowski, M., Dykas, P., Xue, J., and Lanir, B. Methods of improving llm training stability. *arXiv preprint arXiv:2410.16682*, 2024.
- Sakaguchi, K., Bras, R. L., Bhagavatula, C., and Choi, Y. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.
- Sap, M., Rashkin, H., Chen, D., LeBras, R., and Choi, Y. Socialiqa: Commonsense reasoning about social interactions. *arXiv preprint arXiv:1904.09728*, 2019.
- Shazeer, N. and Stern, M. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pp. 4596–4604. PMLR, 2018.
- Shleifer, S., Weston, J., and Ott, M. Normformer: Improved transformer pretraining with extra normalization. *arXiv preprint arXiv:2110.09456*, 2021.
- Shoeybi, M., Patwary, M., Puri, R., LeGresley, P., Casper, J., and Catanzaro, B. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019.
- Siddiqui, S. A., Dong, X., Heinrich, G., Breuel, T., Kautz, J., Krueger, D., and Molchanov, P. A deeper look at depth pruning of llms. *arXiv preprint arXiv:2407.16286*, 2024.
- Sun, J., Zheng, C., Xie, E., Liu, Z., Chu, R., Qiu, J., Xu, J., Ding, M., Li, H., Geng, M., et al. A survey of reasoning with foundation models: Concepts, methodologies, and outlook. *ACM Computing Surveys*, 57(11):1–43, 2025.
- Takase, S., Kiyono, S., Kobayashi, S., and Suzuki, J. B2t connection: Serving stability and performance in deep transformers. In *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 3078–3095, 2023a.
- Takase, S., Kiyono, S., Kobayashi, S., and Suzuki, J. Spike no more: Stabilizing the pre-training of large language models. *arXiv preprint arXiv:2312.16903*, 2023b.

- Team, G., Riviere, M., Pathak, S., Sessa, P. G., Hardin, C., Bhupatiraju, S., Hussenot, L., Mesnard, T., Shahriari, B., Ramé, A., et al. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.
- Team, K., Bai, Y., Bao, Y., Chen, G., Chen, J., Chen, N., Chen, R., Chen, Y., Chen, Y., Chen, Y., et al. Kimi k2: Open agentic intelligence. *arXiv preprint arXiv:2507.20534*, 2025.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Ulyanov, D., Vedaldi, A., and Lempitsky, V. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- Wang, H., Ma, S., Dong, L., Huang, S., Zhang, D., and Wei, F. Deepnet: Scaling transformers to 1,000 layers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(10):6761–6774, 2024.
- Wang, Q., Li, B., Xiao, T., Zhu, J., Li, C., Wong, D. F., and Chao, L. S. Learning deep transformer models for machine translation. *arXiv preprint arXiv:1906.01787*, 2019.
- Welbl, J., Liu, N. F., and Gardner, M. Crowdsourcing multiple choice science questions. In Derczynski, L., Xu, W., Ritter, A., and Baldwin, T. (eds.), *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pp. 94–106, Copenhagen, Denmark, September 2017. Association for Computational Linguistics. doi: 10.18653/v1/W17-4413. URL <https://aclanthology.org/W17-4413/>.
- Wu, Y. and He, K. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 3–19, 2018.
- Xiong, R., Yang, Y., He, D., Zheng, K., Zheng, S., Xing, C., Zhang, H., Lan, Y., Wang, L., and Liu, T. On layer normalization in the transformer architecture. In *International conference on machine learning*, pp. 10524–10533. PMLR, 2020.
- Xu, T., Huang, Z., Sun, J., Cheng, S., and Lam, W. Seqposim: Sequential policy optimization for simultaneous machine translation. *arXiv preprint arXiv:2505.20622*, 2025.
- Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- Yin, Y., Huang, W., Song, K., Tang, Y., Wu, X., Guo, W., Guo, P., Wang, Y., Meng, X., Wang, Y., et al. Pangu ultra: Pushing the limits of dense large language models on ascend npus. *arXiv preprint arXiv:2504.07866*, 2025.
- Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. HellaSwag: Can a machine really finish your sentence? In Korhonen, A., Traum, D., and Màrquez, L. (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4791–4800, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1472. URL <https://aclanthology.org/P19-1472/>.
- Zhang, B. and Sennrich, R. Root mean square layer normalization. *Advances in neural information processing systems*, 32, 2019.
- Zhang, B., Titov, I., and Sennrich, R. Improving deep transformer with depth-scaled initialization and merged attention. *arXiv preprint arXiv:1908.11365*, 2019.
- Zhang, H. and Sra, S. First-order methods for geodesically convex optimization. In *Conference on learning theory*, pp. 1617–1638. PMLR, 2016.
- Zhang, J., Zhao, Y., Saleh, M., and Liu, P. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International conference on machine learning*, pp. 11328–11339. PMLR, 2020.
- Zhang, Y., Chen, C., Li, Z., Ding, T., Wu, C., Kingma, D. P., Ye, Y., Luo, Z.-Q., and Sun, R. Adam-mini: Use fewer learning rates to gain more. *arXiv preprint arXiv:2406.16793*, 2024.
- Zhao, J., Zhang, Z., Chen, B., Wang, Z., Anandkumar, A., and Tian, Y. Galore: Memory-efficient llm training by gradient low-rank projection. *arXiv preprint arXiv:2403.03507*, 2024.
- Zhuo, Z., Zeng, Y., Wang, Y., Zhang, S., Yang, J., Li, X., Zhou, X., and Ma, J. Hybridnorm: Towards stable and efficient transformer training via hybrid normalization. *arXiv preprint arXiv:2503.04598*, 2025.

A. Model Configuration

All experiments are conducted on 8 GPUs. The 125M and 350M model configuration is the following.

Table 1. Model Configurations.

	125M	350M
Training sequence length	512	512
Batch size	32	32
Number of iterations	50k	50k
Dropout prob.	0.0	0.0
Attention dropout prob.	0.0	0.0
Attention head	12	16
Feature dimension	768	1024
Layer number	12	24
Optimizer	Adam	Adam
Optimizer parameter betas	[0.9, 0.95]	[0.9, 0.95]
Learning rate	$6e - 4$	$3e - 4$
Precision	float16	float16

B. Geodesic Optimization for Pre-Norm

Let \mathbf{x}_k denote the token representation after the k -th Transformer layer. A Transformer layer with Pre-Norm can be expressed as:

$$\tilde{\mathbf{x}}_k = \mathbf{x}_k + \phi(\text{Norm}(\mathbf{x}_k)), \quad (6)$$

where $\phi(\cdot)$ denotes either the attention module $\text{Attn}(\cdot)$ or the feed-forward network $\text{FFN}(\cdot)$, and $\text{Norm}(\cdot)$ represents a normalization operator (typically RMSNorm).

Let $\mathbf{x}_0 = \mathbf{x}$ be the initial input. After k layers of the Transformer model with a standard Pre-Norm, the representation can be written as

$$\mathbf{x}_k = \mathbf{x} + \phi_1(\text{Norm}(\mathbf{x})) + \phi_2(\text{Norm}(\mathbf{x}_1)) + \dots + \phi_k(\text{Norm}(\mathbf{x}_{k-1})). \quad (7)$$

This is obtained by recursively applying

$$\mathbf{x}_k = \mathbf{x}_{k-1} + \phi_k(\text{Norm}(\mathbf{x}_{k-1})), \quad (8)$$

for each layer index k . Here, for simplicity, we use $\phi_k(\cdot)$ to denote the attention and feed-forward module at the k -th layer. At the output of the final layer, an additional normalization operation is applied:

$$\text{FinalOutput} = \text{Norm}(\mathbf{x}_k). \quad (9)$$

We consider a widely used normalization method, namely **RMSNorm**, which is scale-invariant and normalizes vectors to have norm \sqrt{D} , where D denotes the hidden dimension. Specifically,

$$\text{Norm}(\mathbf{v}) = \frac{\mathbf{v}}{\|\mathbf{v}\|} \sqrt{D}. \quad (10)$$

As a consequence, for any scalar $c > 0$,

$$\text{Norm}(c\mathbf{v}) = \frac{c\mathbf{v}}{\|c\mathbf{v}\|} \sqrt{D} = \frac{\mathbf{v}}{\|\mathbf{v}\|} \sqrt{D} = \text{Norm}(\mathbf{v}), \quad (11)$$

showing that RMSNorm is invariant to positive rescaling. Hence, we have

$$\text{Norm}\left(\frac{\mathbf{x}_k}{\sqrt{k+1}}\right) = \text{Norm}(\mathbf{x}_k). \tag{12}$$

This implies that, under RMSNorm, dividing the representation by scalars such as $\sqrt{k+1}$ has **no effect** on the final normalized output.

In standard Pre-Norm architectures, the norm $\|\mathbf{x}_k\|$ typically grows with depth k due to the accumulation of residual updates in (8). While one might attempt to control this growth by rescaling \mathbf{x}_k (e.g., dividing by $\sqrt{k+1}$) prior to normalization, such scaling is entirely eliminated by RMSNorm.

Complete Derivation of Alternative Scaled Update From the recursive update in (8), we observe that

$$\frac{\mathbf{x}_k}{\sqrt{k+1}} = \frac{\sqrt{k}}{\sqrt{k+1}} \frac{\mathbf{x}_{k-1}}{\sqrt{k}} + \frac{1}{\sqrt{k+1}} \phi_k \left(\text{Norm} \left(\frac{\mathbf{x}_{k-1}}{\sqrt{k}} \right) \right).$$

Denote $\mathbf{x}_k^{\text{alt}}$ as $\frac{\mathbf{x}_k}{\sqrt{k+1}}$, we obtain the alternative scaled update rule

$$\mathbf{x}_k^{\text{alt}} = \frac{\sqrt{k}}{\sqrt{k+1}} \cdot \mathbf{x}_{k-1}^{\text{alt}} + \frac{1}{\sqrt{k+1}} \cdot \phi_k(\text{Norm}(\mathbf{x}_{k-1}^{\text{alt}})), \quad k \geq 1 \tag{13}$$

with initialization $\mathbf{x}_0^{\text{alt}} = \mathbf{x}_0$.

Equivalence After RMSNorm The final output of the Transformer model with RMSNorm is given by:

$$\text{FinalOutput} = \text{RMSNorm}(\mathbf{x}_k) = \text{RMSNorm}(\mathbf{x}_k^{\text{alt}}), \tag{14}$$

which follows directly from the scale invariance of RMSNorm. Consequently, instead of analyzing the original dynamics \mathbf{x}_k , we may equivalently consider the rescaled sequence $\mathbf{x}_k^{\text{alt}}$, which admits the same final output up to a scalar normalization.

The update scheme in (13) shares the formulation with (3), with inadapative angles $\theta = \arccos\left(\frac{\sqrt{k}}{\sqrt{k+1}}\right)$. In contrast, the proposed GeoNorm allows the adaptive angle, depending on the step size α_k , the current state \mathbf{x}_k , and the update direction s_k . This is convincingly more practical and reliable in manifold optimization. We further verify the Pre-Norm and our GeoNorm interpretation and their output equivalence by conducting experiments on comparing their outputs. The following table shows the similar performance of Transformer models using Pre-Norm and Equation (13), which further supports the equivalence of these two formulations.

Table 2. The performance of different formulations of Pre-Norm update, with training length 512, Books3 dataset and loss metric.

Method	Seed	5K	10K	15K	20K	25K	30K	35K	40K	45K	50K
$x_k + \phi(\text{Norm}(x_k))$	1234	4.1470	3.9028	3.8555	3.7488	3.6815	3.6434	3.5865	3.4906	3.4904	3.4437
$\frac{\sqrt{k}}{\sqrt{k+1}} \cdot \mathbf{x}_k + \frac{1}{\sqrt{k+1}} \cdot \phi(\text{Norm}(\mathbf{x}_k))$	1234	4.1443	3.9012	3.8533	3.7491	3.6795	3.6454	3.5868	3.4894	3.4898	3.4430
$x_k + \phi(\text{Norm}(x_k))$	1235	4.1618	3.9471	3.8361	3.7779	3.6850	3.6297	3.5751	3.5354	3.4751	3.4548
$\frac{\sqrt{k}}{\sqrt{k+1}} \cdot \mathbf{x}_k + \frac{1}{\sqrt{k+1}} \cdot \phi(\text{Norm}(\mathbf{x}_k))$	1235	4.1643	3.9500	3.8377	3.7784	3.6885	3.6307	3.5752	3.5347	3.4738	3.4543
$x_k + \phi(\text{Norm}(x_k))$	1236	4.1537	3.9082	3.8357	3.7429	3.6990	3.6268	3.5140	3.4809	3.4542	3.4364
$\frac{\sqrt{k}}{\sqrt{k+1}} \cdot \mathbf{x}_k + \frac{1}{\sqrt{k+1}} \cdot \phi(\text{Norm}(\mathbf{x}_k))$	1236	4.1540	3.9070	3.8378	3.7417	3.7002	3.6286	3.5152	3.4816	3.4572	3.4381
$x_k + \phi(\text{Norm}(x_k))$	mean	4.1542	3.9194	3.8424	3.7565	3.6885	3.6333	3.5585	3.5023	3.4732	3.4450
	std	0.0061	0.0197	0.0092	0.0153	0.0076	0.0072	0.0318	0.0237	0.0148	0.0076
$\frac{\sqrt{k}}{\sqrt{k+1}} \cdot \mathbf{x}_k + \frac{1}{\sqrt{k+1}} \cdot \phi(\text{Norm}(\mathbf{x}_k))$	mean	4.1542	3.9194	3.8429	3.7564	3.6894	3.6349	3.5591	3.5019	3.4736	3.4451
	std	0.0082	0.0218	0.0073	0.0158	0.0085	0.0075	0.0314	0.0234	0.0133	0.0068

Taken together, the theoretical analysis and experimental evidence suggest that Pre-Norm can be interpreted as a special case of GeoNorm with a fixed, depth-dependent angle. This perspective also sheds light on the empirical stability advantages of Pre-Norm over Post-Norm. In particular, Pre-Norm implicitly introduces a form of step size decay through its normalized residual structure, whereas Post-Norm lacks such a mechanism. From an optimization perspective, this absence may lead to instability or divergence when step sizes are not carefully controlled, explaining the commonly observed loss spikes and training collapse in Post-Norm Transformers.

C. The Effect of Clamp Value

Table 3. The performance of different clamp values for the implementation in Appendix D

Model	5K	10K	15K	20K	25K	30K	35K	40K	45K	50K
$\pi/2$	4.0650	3.8464	3.8096	3.7072	3.6341	3.6060	3.5471	3.4537	3.4531	3.4092
$\pi/4$	4.0642	3.8423	3.8076	3.7018	3.6317	3.6003	3.5427	3.4486	3.4511	3.4040
$\pi/8$	4.0641	3.8501	3.8145	3.7140	3.6471	3.6141	3.5544	3.4638	3.4631	3.4182

D. Implementation Details

In this section, we present the implementation of the proposed GeoNorm module in PyTorch.

```

from tqdm import tqdm
import numpy as np

import torch
import torch.nn as nn

class GeoNorm(nn.Module):
    def __init__(self, clamp=torch.pi/4, method_name="harnomic"):
        super().__init__()

        self.scale=nn.Parameter(torch.tensor(1))
        self.bias=nn.Parameter(torch.tensor(0))
        self.clamp=torch.pi/4
        self.method_name=method_name
    def forward(self, x, g, layer_number, layer_total):
        """
        x: residual;
        g: ffn(x) or attn(x);
        layer_number: the current layer index, from 0,1,2, 3 to layer_total-1
        layer_total: layer total number.
        """
        gradient = g - (x * g).sum(dim=-1, keepdim=True) / (
            torch.norm(x, p=2, dim=-1, keepdim=True)
            ** 2) * x

        tangent_norm = torch.norm(gradient, p=2, dim=-1, keepdim=True)
        safe_tangent_norm = torch.clamp(tangent_norm, min=1e-8)
        unit_tangent = gradient / safe_tangent_norm

        R = torch.norm(x, p=2, dim=-1, keepdim=True)
        safe_R = torch.clamp(R, min=1e-6)
        theta = torch.clamp(safe_tangent_norm / safe_R, max=self.clamp)
        if self.method_name=="harmonic":
            theta_attn = torch.clamp((theta* self.scale +self.bias) / (layer_number +
            1) , max=self.clamp)
        elif self.method_name=="linear":
            theta = torch.clamp((theta* self.scale +self.bias)*(layer_total-
            layer_number) / (layer_total) , max=self.clamp)
        elif self.method_name=="sqrt":
            theta = torch.clamp((theta* self.scale +self.bias) / math.sqrt(
            layer_number + 1) , max=self.clamp)

        output = x * torch.cos(theta) + unit_tangent * safe_R * torch.sin(theta)

        return output

```