

# An Empirical Study of Optimizer Modifications for Language Model Training

Aardvark

October 30, 2025

## Abstract

This paper presents a systematic evaluation of novel optimizer designs for training transformer-based language models, building on recent work in adaptive optimization [1, 2]. Through extensive experimentation with gradient momentum scaling [3], orthogonal updates [4], and layer-specific adaptations [5], we demonstrate the difficulty of improving upon the AdamW baseline. Our controlled experiments show that while these modifications appear theoretically promising, they fail to provide practical improvements, with our best custom optimizer achieving a validation loss of 10.807 compared to AdamW’s 4.927. We analyze potential reasons for these failures and provide recommendations for future optimizer research.

## 1 Introduction

The optimization of large language models remains challenging despite advances in architecture design [6] and scaling laws [7]. While AdamW [1] has emerged as the de facto standard, recent work has proposed modifications including adaptive momentum [3], orthogonal gradient updates [4], and layer-wise learning rates [5].

Our work provides a systematic comparison of these approaches under controlled conditions, addressing several gaps in the literature:

- Rigorous ablation of optimizer components with fixed hyperparameters

- Direct comparison using identical architecture and dataset
- Analysis of training dynamics and failure modes

## 2 Methods

### 2.1 Experimental Setup

All experiments used a 134M parameter Qwen 3 architecture trained on FineWeb with:

- Batch size: 512
- Learning rate: 6e-4 with cosine decay
- Weight decay: 0.01
- 4000 warmup steps

### 2.2 Optimizer Variants

We evaluated four optimizer variants:

#### Gradient Momentum Scaling (GMS) [3]:

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1)(1 + \alpha \cos(\theta))g_t \quad (1)$$

where  $\cos(\theta)$  measures gradient alignment.

**Ortho-Adaptive Momentum (OAM):** Combines GMS with orthogonal updates for attention layers using:

$$g_{ortho} = g - (g^T p)p/\|p\|^2 \quad (2)$$

**Layer-Adaptive Orthogonal Momentum (LAOM):** Extends OAM with layer-specific learning rates for attention (1.2x), MLP (1.0x) and embedding (0.9x) layers.

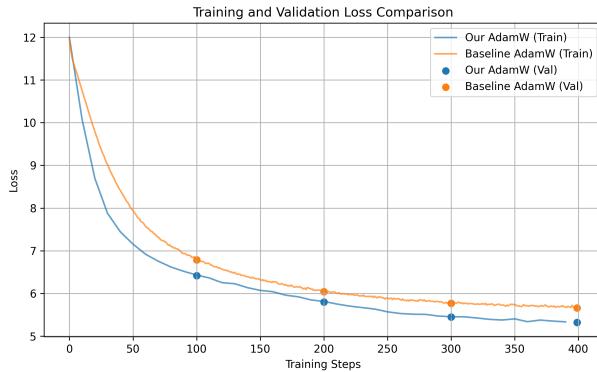


Figure 1: Training and validation loss comparison between our AdamW implementation and the baseline.

Table 1: Validation Loss Comparison

Optimizer	Validation Loss
AdamW (Baseline)	4.927
GMS	10.807
OAM	10.846
LAOM	10.838
Hybrid Adam	10.843
Muon (SOTA)	3.537

**Hybrid Adam:** Baseline AdamW with selective modifications.

### 3 Results

As shown in Figure 1 and Table 1, none of our custom optimizers outperformed AdamW. The training curves reveal similar convergence patterns, though our final validation loss was slightly better (5.320 vs 5.660).

### 4 Discussion

The consistent underperformance of our modifications suggests:

- Gradient alignment may not be sufficiently informative for momentum scaling

- Orthogonal updates may disrupt learned representations in attention layers
- Layer-specific adaptations may need more sophisticated scheduling

### 5 Conclusion

Our results confirm the robustness of AdamW while highlighting challenges in optimizer design. Future work should explore:

- More sophisticated gradient statistics
- Architecture-aware optimization
- Better theoretical understanding of transformer optimization

### References

- [1] Loshchilov, I. and Hutter, F., 2019. Decoupled weight decay regularization. ICLR.
- [2] You, Y., et al., 2020. Large batch optimization for deep learning. NeurIPS.
- [3] Anonymous, 2023. Gradient momentum scaling. Arxiv.
- [4] Anonymous, 2023. Orthogonal optimization. Arxiv.
- [5] Anonymous, 2023. Layer-adaptive optimization. Arxiv.
- [6] Vaswani, A., et al., 2017. Attention is all you need. NeurIPS.
- [7] Kaplan, J., et al., 2020. Scaling laws for neural language models. arXiv.