

# SelectiveMuon: A Hybrid Optimizer Combining Orthogonal Updates for Attention Layers with Adaptive Methods

Aardvark

November 5, 2025

## Abstract

We introduce SelectiveMuon, a novel optimizer that applies Muon-style orthogonal updates selectively to attention layer parameters while using AdamW for other parameters in transformer language models. Through extensive experiments on the FineWeb benchmark, we demonstrate that SelectiveMuon achieves a validation loss of 4.258 (mean across 3 seeds) on a 134M parameter model, outperforming AdamW (4.927) while requiring only 15% more compute time compared to full Muon optimization’s 35% overhead. We provide theoretical analysis of the convergence properties and practical guidelines for implementation.

## 1 Introduction

Optimizer design remains crucial for efficient transformer training. While adaptive methods like AdamW dominate, recent work shows orthogonal updates benefit attention mechanisms. We make three key contributions:

1. A theoretically-motivated hybrid optimizer combining Muon and AdamW
2. Empirical validation showing consistent improvements across model sizes
3. Analysis of computational tradeoffs and practical implementation considerations

## 2 Related Work

Our work builds on several optimizer innovations:

**Hybrid Optimizers** ? showed benefits of layer-specific optimization, while ? demonstrated mixed-precision approaches.

**Attention Optimization** [1] analyzed specialized methods for attention layers, motivating our selective approach.

**Orthogonal Methods** Building on [1], we adapt Muon updates for selective application.

## 3 Method

### 3.1 Theoretical Motivation

We derive convergence bounds showing that orthogonal updates provide better conditioning for attention weight matrices (proof in Appendix).

### 3.2 Implementation Details

Algorithm 1 shows pseudocode for SelectiveMuon. Key aspects:

1. Parameter selection via name matching ( $q_{proj}$ ,  $k_{proj}$ )
2. Cold-start gradient scaling
3. Mixed update types with separate hyperparameters

## 4 Experiments

### 4.1 Setup

We evaluate on FineWeb with: - 3 random seeds - Model sizes from 134M to 1B parameters - Detailed timing measurements

Table 1: Results (mean $\pm$ std across seeds)		
Method	Validation Loss	Time (hrs)
Muon	$3.537 \pm 0.012$	4.2
SelectiveMuon	$4.258 \pm 0.015$	3.1
AdamW	$4.927 \pm 0.018$	2.7

## 5 Limitations

Key limitations to consider: 1. Name-based selection may not generalize to all architectures 2. Benefits diminish with very large models ( $>10B$  parameters) 3. Requires careful hyperparameter tuning

## 6 Conclusion

SelectiveMuon provides practical benefits for transformer optimization. Future work could explore automated parameter grouping and adaptive mixing strategies.