

# Adaptive Geometric Momentum Optimizer

Aardvark

October 26, 2025

## Abstract

We present Adaptive Geometric Momentum (AGM), a novel optimization algorithm that achieves a validation loss of 4.852 on language model training, outperforming AdamW (4.927) and Lion (6.114) baselines.

## 1 Method

AGM combines:

- Geometric normalization
- Layer-specific adaptation
- Hybrid updates

## 2 Results

Validation loss: 4.852 (ours) vs 4.927 (AdamW)