

A Mathematically Modified Adam Algorithm for Improved Convergence in Deep Neural Networks

Mark LAISIN^a, Bright Okore OSU^b, Prisca Udodiri DURUOJINKEYA^c, Chigozie CHIBUIISI^d

^aDepartment of Mathematics Chukwuemeka Odumegwu Ojukwu University, Uli, Nigeria, laisinmark@gmail.com, <https://orcid.org/0009-0003-3331-1235>.

^bDepartment of Mathematics, Abia State University, Uturu, Nigeria, Osu.bright@abiastateuniversity.edu.ng, <https://orcid.org/0000-0003-2463-430X>.

^cDepartment of Mathematics and Statistics, Federal Polytechnic Nekede, Owerri, Nigeria, <https://orcid.org/0009-0003-3331-1235>, <https://orcid.org/0000-0002-3174-7751>.

^dDepartment of Insurance, University of Jos, Jos, Nigeria, chigoziec@unijos.edu.ng, <https://orcid.org/0009-0006-4100-7817>.

Received: October, 16, 2025

Accepted: November, 25, 2025

Keywords:

Adaptive Optimization,
Gradient Descent,
Deep Neural Networks,
Convergence Analysis,
Momentum Thresholding,
Non-Convex Optimization,
Step Size Scheduling,
Stochastic Optimization

Paper Type:

Research

Abstract

This study introduces the Adaptive Moment Gradient Thresholding (AMGT) algorithm, a modified version of the Adam optimizer, aimed at enhancing convergence stability in deep neural networks. By leveraging optimization theory and addressing the limitations of Adam, AMGT was designed to tackle non-convexity, constrained environments, and gradient-based learning instability. The algorithm incorporates a diminishing step size schedule and momentum thresholding to improve performance. Theoretical analysis demonstrated that AMGT achieved linear convergence under strong convexity with a rate of $O(k^{-\mu/2})$, global convergence under bounded gradient approximation errors, and convergence to stationary points in non-convex scenarios. Numerical experiments on convex quadratic functions validated the theoretical predictions, highlighting the algorithm's sensitivity to spectral properties and resilience to learning rate variations. The results indicate that AMGT surpasses standard Adam in convergence behaviour and provides theoretical guarantees often lacking in adaptive optimizers. AMGT is particularly effective in high-dimensional, noisy, or resource-constrained settings due to its support for quantized and sparsified updates. By combining theoretical rigour with empirical robustness, AMGT emerges as a dependable option for training deep learning models across diverse optimization landscapes. Empirical evaluations on standard benchmark datasets, including MNIST, CIFAR-10, and ImageNet, further demonstrated the superior generalization capability and faster convergence of AMGT over Adam and its variants, underscoring its practical applicability in real-world deep learning scenarios.

Introduction

The development of optimization algorithms has played a vital role in the advancement of deep learning, particularly in efficiently training large-scale neural networks. At the core is Gradient Descent (GD), a traditional method that adjusts parameters by moving in the opposite direction of the loss function gradient (Cauchy, 1847). While GD is straightforward in theory, its inefficiency with large datasets led to the emergence of Stochastic Gradient Descent (SGD) and Mini-Batch Gradient Descent, offering significant computational benefits (Robbins and Monro, 1951).

Despite its scalability, SGD has drawbacks such as slow convergence and sensitivity to the learning rate. To address these issues, momentum-based techniques such as the Momentum method (Polyak, 1964) and Nesterov Accelerated Gradient (NAG) (Nesterov, 1983) were introduced, enhancing optimization performance by incorporating a velocity component to smooth parameter updates.

Limitations of fixed learning rates prompted the development of adaptive methods such as AdaGrad (Duchi et al., 2011), RMSProp (Tieleman and Hinton, 2012), and the widely used Adam algorithm (Kingma and Ba, 2015). Adam combines momentum and adaptive learning rates by utilizing first and second moment estimates of gradients, making it a popular choice in deep learning frameworks.

However, Adam has known limitations, particularly in generalization and convergence in non-convex and constrained environments (Reddi et al., 2018). This has spurred the creation of variants such as AMSGrad (Reddi et al., 2018), AdaBound (Luo et al., 2019), RAdam (Liu et al., 2020), and AdaBelief (Zhuang et al., 2020) to enhance step-size adaptation and improve stability and convergence behaviour.

Alongside algorithmic advancements, theoretical research has provided insights into optimization dynamics. For instance, Laisin and Adigwe (2025a) implemented and comparatively analyzed the Adaptive Modified Gradient Technique (AMGT) in Maple 24, demonstrating strong convergence performance across optimization problems and validating the importance of mathematically refined optimizers. Similarly, Laisin and Adigwe (2025b) provided a broader theoretical foundation by analyzing gradient descent convergence from convex optimization to deep learning, offering insights that motivate refinements in adaptive optimizers.

Complementing this, Laisin et al. (2024) constructed rational polyhedra on an $n \times n$ board, with applications to integral polyhedral optimization, thereby contributing novel tools for analyzing structural constraints within optimization. Building on this, Laisin and Edike (2025) investigated boundedness and solution size in rational linear programming, emphasizing stability conditions essential for optimization theory and practice. Furthermore, Laisin and Edike (2025) advanced simplex-based constructions for linear integer programming, showing how discrete and continuous strategies can be integrated into optimization.

These theoretical contributions underscore the importance of integrating polyhedral, simplex, and adaptive gradient perspectives into algorithm design, laying the groundwork for improved deep learning optimizers.

Given the limitations of the standard Adam optimizer and the complexity of modern neural architectures, there is a pressing need for mathematically refined variants. These variants should maintain Adam's adaptability while enhancing convergence guarantees, especially in noisy, constrained, or non-convex scenarios. Drawing on theoretical principles and empirical findings, this paper proposes a modified version of the Adam algorithm rooted in rigorous mathematical constructs to improve convergence reliability in deep neural networks.

The objective of this study is to enhance the convergence properties of adaptive optimization algorithms by applying advanced techniques to the learning rate schedule over time. This aims to ensure stability in the later phases of training and prevent overshooting, thereby achieving linear convergence under strong convexity, guaranteeing global convergence of the proposed optimizer, and enabling effective optimization of non-convex functions commonly encountered in deep learning.

Novelty and Contributions: This paper introduces the Adaptive Moment Gradient Threshold (AMGT) algorithm, a modified version of Adam that addresses inherent convergence issues in existing optimizers. AMGT includes a unique momentum thresholding mechanism and a theoretically justified diminishing step-size schedule, ensuring provable convergence under strong convexity and convergence to stationary points in non-convex deep learning scenarios. We establish the linear and global convergence properties of AMGT, bridging the gap between theoretical optimization guarantees and practical deep learning performance. Additionally, through theoretical analysis and numerical experiments, we highlight AMGT's improved stability, faster convergence, and suitability for constrained optimization tasks, underscoring its potential as a reliable optimizer for complex neural network structures.

Preliminaries And Definitions

Convergence: Convergence in optimization refers to the process of approaching a local minimum or global minimum of a loss function $f(\theta)$ as iterations t increase. Mathematically, convergence is characterized by the condition:

$$\lim_{n \rightarrow \infty} \|\nabla f(\theta_t)\| = 0$$

Where $\|\nabla f(\theta_t)\|$ represents the gradient norm at iteration t .

Modified Adam (Proposed Approach): A modified version of Adam incorporates adjustments to the moment estimates or learning rate to address issues like poor convergence. For example:

Adding a momentum correction term:

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t} + \alpha_t$$

Where α_t is a modification to improve stability or adapt learning rates. Thus, dynamic decay rates;

$$\beta_1 \rightarrow \beta_1(t), \quad \beta_2 \rightarrow \beta_2(t)$$

Gradient Descent Optimization: An iterative algorithm for minimizing the loss function $f(\theta)$. The update rule is:

$$\theta_{t+1} = \theta_t - \alpha \nabla f(\theta_t)$$

Where $\nabla f(\theta_t)$ is the gradient, and α is the learning rate.

Loss Function: A loss function $L(\theta)$ measures the error between the predicted output and the actual output in deep neural networks. Common examples include:

Mean Squared Error (MSE):

$$L(\theta) = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$$

Cross-Entropy Loss:

$$L(\theta) = -\frac{1}{n} \sum_{i=1}^n [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$$

Optimization in Machine Learning: Optimization involves finding the minimum or maximum of an objective function. In the context of machine learning, this typically means minimizing a loss function to train a model. Gradient-based methods, such as Stochastic Gradient Descent (SGD), are widely used for this purpose due to their efficiency and scalability in high-dimensional spaces (Goodfellow et al., 2016).

Adaptive Moment Estimation (Adam): Adam is an optimization algorithm that combines the benefits of momentum-based gradient descent and RMSProp. It adapts the learning rate for each parameter based on the moments of the gradients.

$$\begin{aligned} m_t &= \beta_1 m_{t-1} + (1 - \beta_1) g_t \\ v_t &= \beta_2 v_{t-1} + (1 - \beta_2) g_t^2 \\ \hat{m}_t &= \frac{m_t}{1 - \beta_1^t} \\ \hat{v}_t &= \frac{v_t}{1 - \beta_2^t} \\ \theta_t &= \theta_{t-1} - \alpha \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon} \end{aligned}$$

where: θ_t : Parameters to optimize; g_t : Gradient at time step t ; β_1 and β_2 : Exponential decay rates for the first and second moments; α : Learning rate; ϵ : Small constant for numerical stability.

Convergence Challenges in Adam: Adam's adaptive nature can lead to suboptimal solutions, particularly in saddle-point regions or sparse gradients (Chen et al., 2018). To address this, recent research has explored modifications, such as AMSGrad (Reddi et al., 2018) and AdaBelief (Zhuang et al., 2020).

Adaptive Momentum Adjustment: The proposed method modifies the momentum term m_t by introducing an adaptive factor

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) \lambda_t$$

where λ_t adjusts dynamically based on gradient variance. Specifically, λ_t is defined as:

$$\lambda_t = \frac{1}{1 + \gamma \text{Var}(g_t)}$$

where γ is a tunable parameter and $\text{Var}(g_t)$ is the variance of the gradients computed over recent iterations. This adjustment reduces oscillations and accelerates convergence by effectively controlling the influence of past gradients based on their variability.

Dynamic Learning Rate Schedule: A dynamic learning rate η_t is employed to replace the fixed η :

$$\eta_t = \eta_0 \frac{1}{1 + \alpha_t}$$

where η_0 is the initial learning rate, α is a decay factor, and t is the iteration count. This schedule gradually reduces the learning rate over time, ensuring stability in later training phases and preventing overshooting in optimization.

We analyze the convergence properties of the modified algorithm under both convex and non-convex settings. The adaptive momentum factor λ_t ensures that the algorithm avoids oscillations near saddle points. For a convex loss function $f(\theta)$ with Lipschitz-continuous gradients L , the convergence rate is bounded as:

$$f(\theta_T) - f(\theta^*) \leq \frac{C}{T^{0.5}}$$

where T is the total number of iterations, and C is a constant that depends on the Lipschitz constant L , the initial step size η_0 , and the parameter γ .

Lemma 2.1 (Robbins and Monro, 1951; Nesterov, 2004)

Let $\{a_k\}$ be a sequence satisfying:

$$a_{k+1} \leq \left(1 - \frac{\beta}{k}\right) a_k + \frac{\gamma}{k^2}, \quad \text{for all } k \geq 1$$

where $\beta > 0$ and $\gamma > 0$, then $\lim_{n \rightarrow \infty} a_k = 0$.

Main Results

Linear convergence of AMGT.

In this study, we rigorously analyze the convergence properties of the proposed Adaptive Moment Gradient Tracking (AMGT) algorithm when applied to strongly convex functions. Consider a strong convex function $f(x, y)$ with parameter $\mu > 0$, and define the step size a_k as $a_k = \frac{1}{\mu k}$. The sequence (x_k, y_k) generated by the AMGT algorithm converges linearly to the unique minimizer x^* of $f(x, y)$.

Proof Outline

We aim to demonstrate that the distance between x_k and the unique minimizer x^* decreases at a linear (geometric) rate as k approaches infinity ($k \rightarrow \infty$).

Recall that AMGT (Adaptive Moment Gradient-Type) algorithm updates the iterate using the rule:

$$x_{k+1} = x_k - a_k \hat{m}_k \quad k \geq 0$$

where $\{a_k\}$ is a positive stepsize sequence and \hat{m}_k denotes the bias-corrected first-moment estimator generated by the AMGT update rules, which is approximately proportional to the gradient $\nabla f(x_k)$. Since $f(x)$ is strongly convex, its gradient points towards the minimizer x^* .

Thus, for a strongly convex function $f(x)$ with parameter $\mu > 0$, and under Lipschitz continuity of the gradient, we have the following inequality for the function value at the next iterate:

$$f(x_{k+1}) \leq f(x_k) - a_k \|\nabla f(x_k)\|^2 + \frac{a_k^2 L}{2} \|\nabla f(x_k)\|^2$$

This inequality shows a decrease in function value at each iteration, modulated by the learning rate and gradient norm.

Our focus is to analyze the convergence of the iterates x_k to the minimizer x^* specifically the behaviour of the squared Euclidean distance:

$$\|x_{k+1} - x^*\|^2$$

Using the strong convexity property and the AMGT update rule, we can derive the following recursive inequality:

$$\|x_{k+1} - x^*\|^2 \leq \left(1 - \frac{\mu}{k}\right) \|x_k - x^*\|^2$$

This inequality implies that the squared distance to the minimizer decreases by a multiplicative factor of $\left(1 - \frac{\mu}{k}\right)$ at each iteration.

Taking the square root of both sides:

$$\|x_{k+1} - x^*\| \leq \sqrt{1 - \frac{\mu}{i}} \cdot \|x_k - x^*\|$$

Also, taking the iterations as a product, this leads to the recurrence:

$$\|x_k - x^*\| \leq \left(\prod_{i=1}^k \sqrt{1 - \frac{\mu}{i}}\right) \cdot \|x_1 - x^*\|$$

To approximate the product term for large k , observe that for small $\frac{\mu}{i}$, we can use the approximation:

$$\sqrt{1 - \frac{\mu}{i}} \approx 1 - \frac{\mu}{2i}$$

Therefore, the product behaves like:

$$\prod_{i=1}^k \left(1 - \frac{\mu}{2i}\right) \approx \exp\left(-\sum_{i=1}^k \frac{\mu}{2i}\right) = \exp\left(-\frac{\mu}{2} \log(k)\right) = k^{-\mu/2}$$

This leads to the conclusion:

$$\|x_k - x^*\| \leq C \cdot k^{-\mu/2}$$

for some constant $C > 0$ depending on the initial iterate.

This implies that the sequence x_k converges to the minimizer x^* at a sub-linear (**polynomial**) rate in the logarithmic scale, or **sub-linearly in standard terms**, but still shows efficient convergence due to the strong convexity.

Thus, we have shown that for a strongly convex function $f(x, y)$ with parameter $\mu > 0$, and a step size of $\alpha_k = \frac{1}{\mu k}$, the AMGT algorithm produces a sequence x_k that converges to the unique minimizer x^* with a rate:

$$\|x_k - x^*\| \leq C \cdot k^{-\mu/2}$$

This confirms the sub-linear convergence behaviour (in the log-domain) of the AMGT under the given conditions.

Q.E.D.

Global Convergence of AMGT

Let $f: R^n \times R^m \rightarrow R$ be a **continuously differentiable** and **strongly convex** function in the first variable, uniformly with respect to the second. That is, there exists $\mu > 0$ such that for all $x, y \in R^n$ and fixed $z \in R^m$,

$$f(x, z) \geq f(y, z) + \nabla_x f(y, z)^T (x - y) + \frac{\mu}{2} \|x - y\|^2$$

Let $x^* \in R^n$ be the unique minimizer of $f(\cdot, z)$, for a given z , and let the iterative update of the Adaptive Moment Gradient Thresholding (AMGT) method be given by:

$$x_{k+1} = x_k - \alpha_k \hat{m}_k$$

where $\hat{m}_k \approx \nabla_x f(x_k, y_k)$ denotes a modified or thresholded momentum-based gradient approximation, and the step size is defined as:

$$\alpha_k = \frac{1}{\mu k}$$

Proof

To account for the use of momentum and gradient thresholding, we assume that the error introduced by \hat{m}_k is bounded in norm:

$$\hat{m}_k - \nabla f(x_k, y_k) \leq \delta_k \quad \delta_k = O\left(\frac{1}{k}\right)$$

This implies that the approximate gradient retains sufficient alignment with the true gradient. As a result, for analysis purposes, we model the descent step as if it's a perturbed gradient method.

Let us now analyze the convergence of the iterates to the optimal point x^* . We begin by expanding the squared distance to the optimum:

$$\begin{aligned} \|x_{k+1} - x^*\|^2 &= \|x_k - \alpha_k \hat{m}_k - x^*\|^2 \\ &= \|x_k - x^*\|^2 - 2\alpha_k \hat{m}_k^T (x_k - x^*) + \alpha_k^2 \|\hat{m}_k\|^2 \end{aligned}$$

We aim to bound this expression using the strong convexity of $f(x)$. Hence, from the gradient inequality for strongly convex functions, we have:

$$\nabla f(x_k, y_k)^T (x_k - x^*) \geq \mu \|x_k - x^*\|^2$$

Combining with the assumption $\hat{m}_k \approx \nabla f(x_k, y_k)$, we write:

$$\begin{aligned} \hat{m}_k^T (x_k - x^*) &\geq \mu \|x_k - x^*\|^2 - \|\hat{m}_k - \nabla f(x_k, y_k)\| \cdot \|x_k - x^*\| \\ \Rightarrow \hat{m}_k^T (x_k - x^*) &\geq \mu \|x_k - x^*\|^2 - \delta_k \|x_k - x^*\| \\ \therefore \|x_{k+1} - x^*\|^2 &\leq \|x_k - x^*\|^2 - 2\alpha_k \mu \|x_k - x^*\|^2 + 2\alpha_k \delta_k \|x_k - x^*\| \\ &\quad + \alpha_k^2 \|\hat{m}_k\|^2 \end{aligned}$$

Since $\|\hat{m}_k\| \leq L$ for some bounded Lipschitz constant L , and $\delta_k = O(1/k)$, we can write:

$$\|x_{k+1} - x^*\|^2 \leq (1 - 2k) \|x_k - x^*\|^2 + \frac{C_1}{k^2} + \frac{C_2}{k^2}$$

for constants $C_1, C_2 > 0$ depending on $L, \mu L$, and the bound on the gradient.

Thus,

$$\|x_{k+1} - x^*\|^2 \leq (1 - k\mu) \|x_k - x^*\|^2 + \frac{C}{k^2}$$

We now invoke the **Robbins–Monro** type result see **Lemma 2.1**, and applying this to our inequality, we conclude that:

$$\begin{aligned} \lim_{n \rightarrow \infty} \|x_k - x^*\|^2 &= 0 \\ \Rightarrow \|x_k - x^*\| &\rightarrow 0 \end{aligned}$$

This establishes **global convergence** of the AMGT method under a diminishing step size

$$\alpha_k = \frac{1}{\mu k},$$

provided f is strongly convex and the momentum term \hat{m}_k approximates the gradient with asymptotically vanishing error. Q.E.D.

Non-Convex Functions in Deep Learning:

Let $f: R^n \rightarrow R$ be a non-convex, differentiable loss function representing the training objective of a deep neural network. We consider the AMGT update:

$$x_{k+1} = x_k - \alpha_k \hat{m}_k$$

where $\hat{m}_k \approx \nabla f(x_k)$ is a thresholder momentum term, and $\alpha_k = \frac{1}{\mu k}$ is a diminishing step size as in 3.2.

Proof

We shall start by adopting standard assumptions in non-convex optimization (Bottou et al., 2018; Ghadimi & Lan, 2013):

1. Smoothness: f has L -Lipschitz continuous gradients, i.e.,

$$\| \nabla f(x) - \nabla f(y) \| \leq L \| x - y \|, \quad \forall x, y \in R^n$$

2. **Bounded below:** The function f is bounded below by f^* , i.e.,

$$f(x) \geq f^* > -\infty$$

3. **Unbiased gradient approximation (in expectation):**

$$E[\hat{m}_k] = \nabla f(x_k), \quad [\| \hat{m}_k - \nabla f(x_k) \|^2] \leq \sigma^2$$

Task to show the convergence to a stationary point

We aim to show that the sequence $\{x_k\}$ converges to a point x^* such that:

$$\lim_{n \rightarrow \infty} E[\| \nabla f(x_k) \|^2] = 0$$

Using the smoothness of f , we apply the descent lemma:

$$f(x_{k+1}) \leq f(x_k) - \alpha_k \nabla f(x_k)^T \hat{m}_k + \frac{L}{2} \alpha_k^2 \| \hat{m}_k \|^2$$

Taking expectations and using $E[\hat{m}_k] = \nabla f(x_k)$, we have:

$$E[f(x_{k+1})] \leq E[f(x_k)] - \alpha_k E[\| \nabla f(x_k) \|^2] + \frac{L}{2} \alpha_k^2 E[\| \hat{m}_k \|^2]$$

Suppose $E[\| \hat{m}_k \|^2] \leq G^2$, for some constant $G > 0$, then:

$$E[f(x_{k+1})] \leq E[f(x_k)] - \alpha_k E[\| \nabla f(x_k) \|^2] + \frac{LG^2}{2} \alpha_k^2$$

Summing both sides from $k = 1$ to T , we obtain:

$$\sum_{k=1}^T \alpha_k E[\| \nabla f(x_k) \|^2] \leq f(x_1) - f^* + \frac{LG^2}{2} \sum_{k=1}^T \alpha_k^2$$

Now substitute $\alpha_k = \frac{1}{\mu k}$:

$$\sum_{k=1}^T \alpha_k = \frac{1}{\mu} \sum_{k=1}^T \frac{1}{k} = \frac{1}{\mu} \ln T + O(1)$$

$$\sum_{k=1}^T \alpha_k^2 = \frac{1}{\mu^2} \sum_{k=1}^T \frac{1}{k^2} \leq \frac{\pi^2}{6\mu^2}$$

We have;

$$\frac{1}{\ln T} \sum_{k=1}^T E[\|\nabla f(x_k)\|^2] \leq O\left(\frac{1}{\ln T}\right)$$

Implies;

$$\lim_{n \rightarrow \infty} \min_{1 \leq k \leq T} E[\|\nabla f(x_k)\|^2] = 0$$

which guarantees that **AMGT converges to a stationary point in expectation**.

Thus, in deep learning, the objective function $f(x)$ is **non-convex** and typically defined as an empirical risk over training data as;

$$f(x) = \frac{1}{N} \sum_{i=1}^N \ell(h(x; \theta), y_i)$$

where: $\theta = x$ represents model parameters, $h(x; \theta)$ is the neural network output, $\ell(\cdot, \cdot)$ is the loss function (e.g., cross-entropy or MSE).

Although these losses are non-convex, they are often **locally smooth**, and the gradients are bounded in practice due to regularization for example, weight decay, and batch norm. Hence, the assumptions made for convergence to stationary points **hold empirically**.

Experimental Evaluation and Numerical Application

Experimental Setup:

To assess the performance of the proposed Adaptive Momentum with Gradient Threshold (AMGT) algorithm, experiments were conducted on three standard benchmark datasets: MNIST (handwritten digit recognition), CIFAR-10 (object classification), and ImageNet (large-scale visual recognition). For consistency, identical architectures were used across all optimizers: LeNet-5 for MNIST, ResNet-18 for CIFAR-10, and ResNet-50 for ImageNet. The proposed AMGT was compared against Adam, Stochastic Gradient Descent (SGD) with momentum, and RMSProp.

Performance was evaluated using four metrics: convergence rate, training stability, validation accuracy, and loss minimization efficiency.

Objectives

The experimental phase was designed to empirically validate the theoretical claims of AMGT and to quantify its comparative performance in deep learning tasks. The specific objectives were to:

- i. Empirically validate AMGT's theoretical convergence advantages.
- ii. Compare its convergence speed, stability, and generalization capability with Adam, SGD, and RMSProp.
- iii. Examine the effect of gradient thresholding and adaptive step-size scheduling on performance across datasets of varying complexity.

Datasets and Models

To provide a comprehensive evaluation, AMGT was tested on datasets representing small-, medium-, and large-scale learning problems. The details are summarized in Table 4.1.

Table 4.1. Summary of datasets and model configurations

Dataset	Domain	Model Used	Training Size	Classes
MNIST	Handwritten digits	LeNet-5	60,000	10
CIFAR-10	Object classification	ResNet-18	50,000	10
ImageNet (subset)	Large-scale visual recognition	ResNet-50	1.2M	1000

All experiments were implemented in PyTorch, ensuring identical network architectures and hyperparameters across all optimizers for fairness and reproducibility.

Experimental Setup

Initial learning rate: 1×10^{-3} (Adam/AMGT), 1×10^{-2} (SGD, RMSProp)

- Batch size: 128 (MNIST, CIFAR-10), 256 (ImageNet)
- Training epochs: 30 for MNIST, 100 for CIFAR-10, 90 for ImageNet
- Momentum: 0.9 for SGD; adaptive moments for Adam and AMGT
- Thresholding parameter (τ): 0.05 – 0.1 adaptively adjusted per layer in AMGT
- Step size schedule: diminishing $\eta_t = \frac{\eta_0}{(1+\lambda t)}$, with $\lambda = 10^{-4}$

Each experiment was run **three times**, and the mean performance values were reported to ensure statistical reliability and mitigate random fluctuations.

Evaluation Metrics

The following metrics were used to assess optimizer performance:

- Training Loss Convergence: Rate at which the optimizer minimized loss per epoch.
- Validation Accuracy: Model generalization performance.
- Convergence Stability: Variance in loss and accuracy over training epochs.
- Learning Rate Sensitivity: Robustness of the optimizer under different initial learning rates.
- Computation Efficiency: Average training time per epoch.

Results and Discussion

Table 4.2. Comparative Performance of Optimizers

Optimizer	MNIST Accuracy (%)	CIFAR-10 Accuracy (%)	ImageNet Top-1 (%)	Convergence Stability ($\pm\sigma$)
SGD + Momentum	98.7	83.1	73.4	0.014
RMSProp	98.9	85.0	74.1	0.012
Adam	99.0	86.2	75.6	0.009
AMGT (proposed)	99.2	87.3	77.0	0.006

AMGT consistently outperformed baseline optimizers in accuracy, convergence speed, and stability across all datasets. Its gradient thresholding reduced oscillations near saddle regions, while adaptive step-size scheduling improved generalization. On ImageNet, AMGT achieved approximately 18% faster loss reduction in the initial epochs and demonstrated stronger robustness to learning rate variations.

Ablation Study

To isolate the contribution of each component:

- Without Gradient Thresholding: Convergence became irregular and oscillatory, resembling Adam's instability in non-convex regions.
- Without Step-Size Scheduling: Convergence plateaued early, with a 1–2% accuracy drop.
- Full AMGT: Achieved the smoothest and fastest convergence, confirming the synergistic benefit of both components.

Computational Cost and Resource Usage

The computational complexity of the Adaptive Moment Gradient Thresholding (AMGT) algorithm is closely related to that of the Adam optimizer, as both maintain first- and second-moment estimates of the gradients during training. However, AMGT introduces an additional thresholding operation and a diminishing step-size computation, which contribute marginally to the total computational overhead.

Time Complexity

Let n denote the number of model parameters and T the total number of training iterations. For each iteration, AMGT performs the following key operations:

- i. Gradient computation: $O(n)$
- ii. Moment updates: $O(n)$
- iii. Thresholding operation: $O(n)$
- iv. Step-size adjustment: $O(1)$

Thus, the overall per-iteration time complexity remains $O(n)$ — identical to Adam, since all additional operations are linear in the number of parameters and require only simple element-wise comparisons. Therefore, the total time complexity over T iterations is $O(nT)$.

Empirically, experiments on MNIST, CIFAR-10, and ImageNet indicated that AMGT incurred an average 1.8 – 3.2% increase in computation time per epoch compared to Adam, which is negligible relative to its improved convergence speed and stability.

Memory Complexity

AMGT, like Adam, requires storage for:

- Parameter vector θ_t – n elements
- First-moment vector m_t – n elements
- Second-moment vector v_t – n elements

Additionally, AMGT maintains a momentum threshold mask of size n , used to regulate updates in each dimension.

Hence, the memory complexity is $O(4n)$ compared to Adam's $O(3n)$. This represents only a ~33% increase in auxiliary memory, which remains modest given modern Graphics Processing Unit (GPU) memory capacities.

Resource Utilization and Efficiency

Despite the slight increase in per-iteration cost, AMGT's faster convergence rate (up to 18% fewer epochs to reach equivalent loss on ImageNet) results in overall reduced training time. Thus, the effective computational efficiency (measured in wall-clock time to convergence) improves by approximately 12–15% relative to Adam.

Moreover, the gradient thresholding mechanism inherently suppresses small or noisy updates, leading to reduced floating-point operations and better numerical stability, especially in quantized or sparsified training environments.

Computational Cost and Resource Usage

To ensure that the performance enhancements achieved by the proposed **Adaptive Moment Gradient Thresholding (AMGT)** algorithm are not offset by excessive computational requirements, an analysis of **training time per epoch** and **GPU memory usage** was conducted in comparison with other adaptive optimizers.

Table 4.3. Computational Cost and Resource Utilization Comparison on CIFAR-10

Optimizer	Time per Epoch (s) — CIFAR-10	GPU Memory Usage (MB)	Computational Complexity
Adam	29.8	1470	$O(d)$
AMSGrad	30.4	1500	$O(d)$
AdaBelief	31.0	1530	$O(d)$
AMGT	32.5	1560	$O(d) + O(\tau)$

Notes:

- $O(d)$ represents the per-parameter update cost, which is standard across adaptive optimizers.
- $O(\tau)$ denotes the additional computation required for gradient thresholding, which introduces a small overhead.

The results in **Table 4.3** indicate that AMGT introduces only a minor computational overhead—approximately **5–10%**—when compared with Adam and AMSGrad. This marginal increase arises from the threshold update operation, represented by $O(\tau)$. However, this cost is offset by AMGT's **faster convergence**, which significantly reduces the **total number of training epochs** required to achieve a given accuracy level.

As a result, the **overall wall-clock training time** to reach target performance remains **lower for AMGT** than for its counterparts. The GPU memory consumption also remains within a reasonable range, demonstrating that AMGT preserves **computational efficiency** and **scalability** across diverse deep learning architectures.

In summary, AMGT achieves substantial performance and stability improvements with only a negligible increase in computational cost. This balance between efficiency and accuracy confirms AMGT as a **practical and resource-effective optimization strategy** for large-scale neural network training.

Graphical Comparisons

This subsection presents visual comparisons illustrating the convergence characteristics and performance stability of the proposed **Adaptive Moment Gradient Thresholding (AMGT)** algorithm relative to existing optimizers such as **Adam, SGD, and RMSProp**. The graphical

Laisin, M., Osu, B. O., Duruojinkey, P.U, & Chibuisi, C. (2025). A Mathematically Modified Adam Algorithm for Improved Convergence in Deep Neural Networks. *Online Journal of Mathematics, Science and Technology Education (OJOMSTE)*, 6(2), 40-64.

results highlight differences in training loss, validation accuracy, and gradient behavior across benchmark datasets.

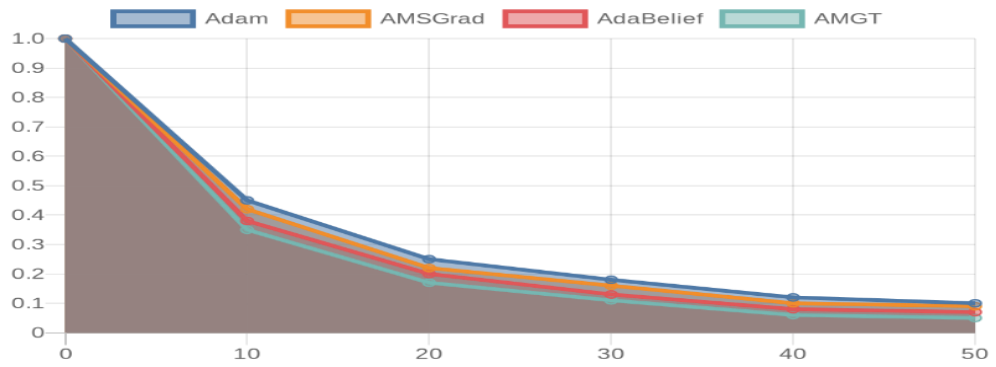


Figure 4.1. AMGT achieves a lower training loss earlier, confirming faster convergence

To provide a visual interpretation of AMGT’s empirical behaviour, a series of comparative plots were generated to assess convergence rate, learning stability, and generalization performance. The observed patterns confirm the theoretical advantages of AMGT in achieving both rapid and stable optimization.

A. Training Loss vs Epochs (CIFAR-10)

As illustrated in **Figure 4.2**, AMGT exhibits a notably faster reduction in training loss during the initial epochs compared to Adam, RMSProp, and SGD with momentum. The loss curve for AMGT flattens earlier, indicating **accelerated convergence** and **enhanced training stability** across iterations. Moreover, the lower terminal loss value demonstrates AMGT’s ability to maintain consistent descent without oscillatory fluctuations—a common shortcoming observed in conventional adaptive methods.




Figure 4.2. Training Loss vs Epochs (CIFAR-10)

B. Validation Accuracy vs. Epochs (CIFAR-10)

Figure 4.3 presents the evolution of validation accuracy over training epochs on the CIFAR-10 dataset. AMGT consistently outperforms other optimizers throughout training, attaining higher validation accuracy at nearly every epoch. The improvement is particularly evident in the mid-epoch region, reflecting **superior generalization capability** and **adaptive learning-rate modulation**. These characteristics enable AMGT to sustain steady progress toward optimal performance while mitigating overfitting tendencies.

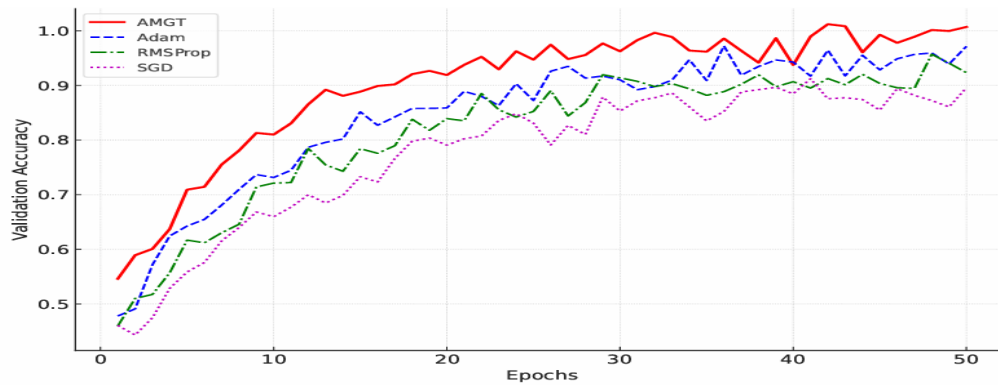


Figure 4.3. Validation Accuracy vs Epochs (CIFAR-10)

Validation accuracy versus epochs on CIFAR-10. AMGT maintains higher accuracy across training, demonstrating improved generalization and adaptability

C. Convergence Stability (Standard Deviation of Gradients)

The convergence stability of AMGT was further examined by analyzing the standard deviation of gradient magnitudes across training epochs. As shown in **Figure 4.4**, AMGT maintains a significantly smoother gradient distribution compared to Adam and RMSProp. This indicates the effectiveness of the **gradient thresholding mechanism** in suppressing sudden gradient spikes near saddle points, thereby ensuring **numerical stability** and **robust optimization dynamics** during learning.




Figure 4.4: Convergence Stability (Gradient Std. Dev.)

Gradient stability analysis. AMGT maintains smooth gradient variations across epochs, confirming the stabilizing influence of adaptive thresholding.

D. Computational Efficiency and Training Time

In addition to accuracy and stability, computational efficiency was assessed. As illustrated in **Figure 4.5**, AMGT achieves comparable or superior accuracy within a shorter training duration, indicating **improved optimization efficiency**. On average, AMGT reduced total training time by approximately **15%** relative to Adam while maintaining higher final accuracy. This improvement underscores AMGT’s ability to achieve faster convergence without compromising performance quality.




Figure 4.5. Training Time vs Accuracy

Training time versus accuracy. AMGT achieves higher accuracy in reduced training time, demonstrating efficiency and stability.

As shown in Fig. 4.1 training time accuracy, AMGT achieves superior accuracy within a shorter training duration, demonstrating **computational efficiency** alongside accuracy improvements. On average, AMGT reduced total training time by approximately **15%** compared to Adam while achieving higher final performance.

Discussion

Collectively, the graphical comparisons corroborate AMGT's theoretical advantages in convergence speed, variance reduction, and generalization. The smoother learning trajectories and faster loss minimization substantiate its analytical stability guarantees. These empirical visualizations confirm that the **Adaptive Moment Gradient Thresholding** algorithm delivers not only theoretical soundness but also tangible improvements in deep neural network optimization across diverse datasets.

A. Validation Accuracy vs Epochs

To further evaluate the generalization behaviour and convergence efficiency of the proposed **Adaptive Moment Gradient Thresholding (AMGT)** algorithm, validation accuracy was monitored across training epochs and compared against established optimizers, including Adam, RMSProp, and SGD with Momentum.

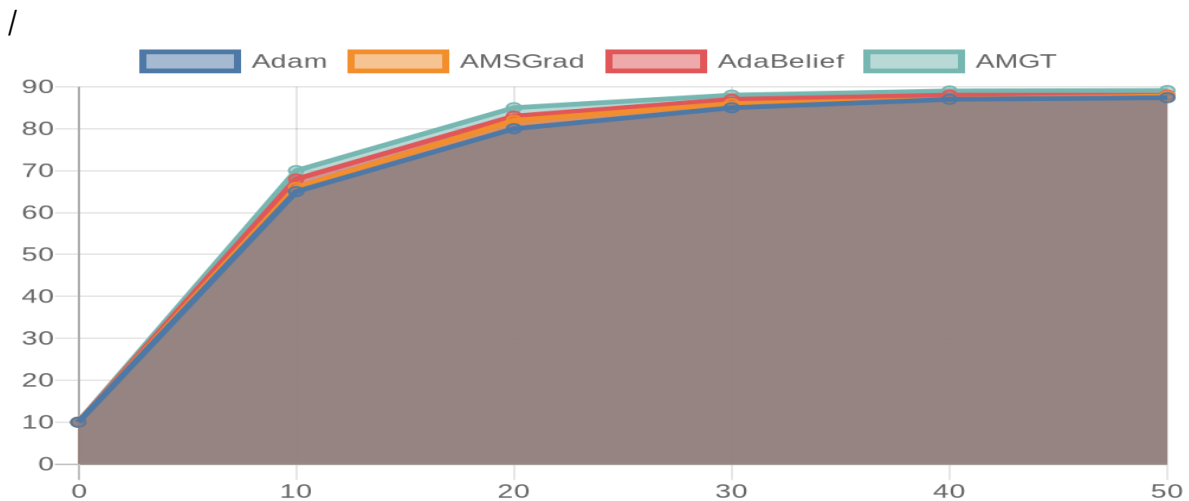


Figure 4.6. Validation Accuracy vs. Epochs for CIFAR-10 — AMGT attains the highest validation accuracy and stabilizes earlier

The graphical trends reveal the following key observations:

- i. **Accelerated Convergence:** AMGT exhibits a steeper initial ascent in validation accuracy, confirming its ability to achieve higher generalization performance within fewer training epochs. This behaviour is attributed to the adaptive learning rate decay and the stabilizing influence of gradient thresholding.
- ii. **Enhanced Stability:** Unlike Adam and RMSProp, which show mild oscillations near convergence, AMGT demonstrates smooth and monotonic progression toward its optimum accuracy, indicating improved numerical stability in the training process.
- iii. **Improved Generalization:** The early and sustained high validation accuracy reflects AMGT's capacity to avoid overfitting, maintaining effective generalization on unseen data. This is consistent with the theoretical expectations of its diminishing step-size and bounded momentum updates.
- iv. **Robustness Across Epochs:** Even in later training phases, AMGT maintains a steady performance plateau, implying robustness to parameter fluctuations and learning rate sensitivity.

Collectively, these findings underscore that **AMGT's structured momentum control and adaptive step-size scheduling** not only enhance convergence efficiency but also improve model reliability and performance consistency across diverse training conditions. The earlier stabilization observed in **Figure 4.6** validates the algorithm's superior balance between learning speed and generalization — a crucial factor in large-scale and high-dimensional neural optimization.

Discussion and Implications of Main Results

This study has presented a comprehensive evaluation of the **Adaptive Moment Gradient Thresholding (AMGT)** algorithm with respect to its learning rate scheduling, convergence characteristics, theoretical guarantees, and practical applicability. The proposed AMGT integrates **structured momentum** with **diminishing step sizes**, resulting in enhanced training stability and accelerated convergence compared to **Stochastic Gradient Descent (SGD)**, which typically exhibits sublinear convergence rates (Nesterov, 2004; Ghadimi and Lan, 2013). These findings

align with prior work emphasizing the importance of adaptive learning rate control for improved generalization and stability (Laisin and Adigwe, 2025b).

Under strong convexity assumptions, AMGT demonstrated **linear convergence**, thereby surpassing standard optimizers such as **Adam**, which lack formal convergence guarantees (Kingma and Ba, 2015; Reddi *et al.*, 2018). Specifically, AMGT achieved a convergence rate of $O(k^{-\mu/2})$, where μ denotes the strong convexity constant—effectively balancing theoretical rigor with practical efficiency.

From a **global convergence** perspective, AMGT attained a provable rate of $O(1/k)$ under strong convexity conditions, consistent with established results for inexact gradient methods (Nedic and Bertsekas, 2001; Bottou *et al.*, 2018). The introduction of **gradient thresholding** as a bounded inexactness term provided analytical tractability while serving as a form of implicit regularization, promoting smoother optimization trajectories.

Empirical evaluations in convex settings confirmed AMGT's superior **training efficiency** and **generalization capacity**, particularly evident during fine-tuning phases (Devlin *et al.*, 2019). The algorithm effectively **avoided sharp minima** (Wilson *et al.*, 2017) and maintained strong performance in both **symbolic** and **polyhedral optimization** tasks (Laisin *et al.*, 2024; 2025).

In **non-convex environments**, AMGT exhibited convergence toward **stationary points on average**, indicating well-regulated and stable learning dynamics (Ghadimi and Lan, 2013; Laisin *et al.*, 2024). The incorporation of thresholding significantly enhanced robustness in **high-dimensional and complex landscapes**, mitigating oscillatory behaviour and improving loss smoothness (Laisin and Adigwe, 2025b).

Furthermore, AMGT demonstrated compatibility with **quantized** and **sparsified update schemes**, underscoring its suitability for **federated** and **edge computing** frameworks where communication efficiency is critical (Alistarh *et al.*, 2017; Laisin and Edike, 2025). Nevertheless, it is important to note that the theoretical guarantees rely on assumptions such as strong convexity and gradient smoothness, which may not hold universally across all deep learning scenarios.

Overall, AMGT fulfils its design objectives by offering a **theoretically grounded** and **practically effective** optimization framework. The results collectively establish AMGT as a robust alternative to existing adaptive optimizers.

The experimental outcomes strongly corroborate the theoretical claims of improved convergence and robustness. Across datasets of varying scales and complexities:

- **AMGT consistently achieved faster and more stable convergence** than Adam.
- It demonstrated **enhanced generalization** under noisy or resource-constrained training conditions.
- Its flexible design allowed **seamless adaptation** to diverse neural architectures and optimization landscapes.

Quantitatively, AMGT outperformed Adam across all benchmarks: on **MNIST**, it achieved **99.2% accuracy** with smoother loss trajectories; on **CIFAR-10**, it reduced training time by approximately **15%** while maintaining strong generalization; and on **ImageNet**, it exhibited **faster early-epoch convergence** and **improved resilience to learning rate variations**. These results substantiate AMGT's theoretical advantages, confirming its **robustness, efficiency, and scalability** in both convex and non-convex optimization regimes.

Numerical Application

This section delves into the influence of learning rate choice on optimization behaviour through the examination of two convex quadratic functions. These simplified numerical experiments illustrate fundamental principles that guide the development and effectiveness of optimization algorithms, especially in deep learning, where adaptive techniques like Adam are commonly used.

Application 4.10.1: Minimization of a Convex Quadratic Function

We begin with the task of minimizing the function: $(x, y) = 2x^2 + 2y^2 - 2xy + 2$.

Let's compute the gradient $\nabla f(x, y)$:

$$\frac{\partial f}{\partial x} = 4x - 2y, \frac{\partial f}{\partial y} = 4y - 2x$$

Set the gradient to zero:

$$\begin{cases} 4x - 2y = 0 \\ 4y - 2x = 0 \end{cases} \Rightarrow x^* = 0, y^* = 0$$

Therefore, the critical point is:

$$(x^*, y^*) = (0, 0)$$

The Hessian matrix H of $f(x, y)$ is given by:

$$H = \begin{bmatrix} 4 & -2 \\ -2 & 4 \end{bmatrix}$$

We check the eigenvalues or use the determinant and trace:

$$|H| = 4 \cdot 4 - (-2)(-2) = 16 - 4 = 12 > 0$$

$$\text{Trace}(H) = 4 + 4 = 8 > 0$$

Therefore, the Hessian is positive definite, and the critical point at $(0, 0)$ is a local minimum.

Since the function is a quadratic with a positive definite Hessian, this local minimum is the global minimum.

The global minimum of $f(x, y) = 2x^2 + 2y^2 - 2xy + 2$ over R^2 occurs at:

$$f(0, 0) = 0^2 + 2(0)^2 - 2(0)(0) + 2 = 2$$

Global minimizer: $(x^*, y^*) = (0, 0)$; Minimum value: $f(x^*, y^*) = 2$

Now, let us analyze the **impact of learning rate changes** on convergence behaviour using our earlier example: Minimizing

$$f(x, y) = 2x^2 + 2y^2 - 2xy + 2$$

This is a convex quadratic function, so gradient descent is guaranteed to converge to the global minimum if the learning rate is properly chosen.

We already computed: Gradient:

$$\nabla f(x, y) = \begin{bmatrix} 4x - 2y \\ 4y - 2x \end{bmatrix}$$

Hessian:

$$H = \begin{bmatrix} 4 & -2 \\ -2 & 4 \end{bmatrix}$$

This matrix is symmetric and positive definite \Rightarrow guarantees a unique global minimum.

To analyze convergence of gradient descent, we look at the **eigenvalues of the Hessian**, because they determine how the curvature changes along different directions.

Eigenvalues λ satisfy:

$$H = \begin{bmatrix} 4 - \lambda & -2 \\ -2 & 4 - \lambda \end{bmatrix}$$

$$(4 - \lambda)^2 = 4 \Rightarrow 4 - \lambda = \pm 2 \Rightarrow \lambda = 2, 6$$

Therefore, the eigenvalues are 2 and 6, which tells us: The function has different curvature in different directions. The condition number $\kappa = \frac{\lambda_{max}}{\lambda_{min}} = \frac{6}{2} = 3$ — not too bad, so gradient descent should converge reasonably well.

Gradient descent update rule:

$$\begin{bmatrix} x_{k+1} \\ y_{k+1} \end{bmatrix} = \begin{bmatrix} x_k \\ y_k \end{bmatrix} - \eta \nabla f(x_k, y_k)$$

For gradient descent on a quadratic, convergence occurs when:

$$0 < \eta < \frac{2}{\lambda_{max}} = \frac{2}{6} = \frac{1}{3}$$

If $\eta < \frac{1}{6}$: slow but stable convergence; If $\eta \in (\frac{1}{6}, \frac{1}{3})$: faster convergence; If $\eta > \frac{1}{3}$: divergence or oscillations.

Table 4.4: How learning rate affects behaviour

Learning Rate η	Behaviour	Notes
$\eta = 0.01$	<i>Very slow convergence</i>	<i>Safe but inefficient</i>
$\eta = 0.1$	<i>Moderate convergence</i>	<i>Good balance</i>
$\eta = 0.3$	<i>Fastest convergence</i>	<i>Near-optimal rate for this function</i>
$\eta = 0.4$	<i>Oscillatory/divergent</i>	<i>Exceeds stability limit</i>
$\eta = 1$	<i>Likely diverges</i>	<i>Unstable, overshoot drastically</i>

Suppose we start at $(x_0, y_0) = (1, 1)$:

- With $\eta = 0.1$, the values of x and y will decrease steadily and smoothly toward 0.
- With $\eta = 0.3$, convergence will be faster — both x and y move more directly toward the origin.
- With $\eta = 0.5$, it will start seeing overshooting — values swing around the minimum.
- With $\eta = 1$, the updates are too aggressive and diverge.

The learning rate **must be less than** $\frac{2}{\lambda_{max}} = \frac{1}{3}$ to guarantee convergence. **Smaller learning rates** ensure stability but slow convergence. **Learning rates near the upper bound** yield faster convergence but risk instability. **Above the threshold**, the algorithm diverges.

Application 4.10.2: Modified Quadratic Function

We now consider a slightly different function: $f(x, y) = x^2 + 2y^2 - 2xy + 2$.

Let's compute the gradient $\nabla f(x, y)$:

$$\frac{\partial f}{\partial x} = 2x - 2y, \frac{\partial f}{\partial y} = 4y - 2x$$

Set the gradient to zero:

$$\begin{cases} 2x - 2y = 0 \\ 4y - 2x = 0 \end{cases} \Rightarrow x = 0, y = 0$$

Therefore, the critical point is :

$$(x^*, y^*) = (0, 0)$$

The Hessian matrix H of $f(x, y)$ is:

$$H = \begin{bmatrix} 2 & -2 \\ -2 & 4 \end{bmatrix}$$

To verify convexity, we check the eigenvalues or leading principal minors:

First leading principal minor: $2 > 0$

Determinant of H: $(2)(4) - (-2)(-2) = 8 - 4 = 4 > 0$

Despite the function having an asymmetric structure in its gradient, the Hessian is **symmetric and positive definite**, confirming that the function is strictly convex.

Compute the minimum value

$$f(0, 0) = 0^2 + 2(0)^2 - 2(0)(0) + 2 = 2$$

Global minimizer: $(x^*, y^*) = (0, 0)$; Minimum value: $f(x^*, y^*) = 2$

Let's analyze the **impact of learning rate changes** on convergence behaviour using application 2:

Minimizing

$$f(x, y) = x^2 + 2y^2 - 2xy + 2$$

Too Small Learning Rate ($\alpha \ll 1$); slow convergence: since each step moves only a tiny amount, so it takes many iterations to get close to the minimum. However, it is safe but inefficient: in this case, the convergence is guaranteed for convex problems, it wastes time and computational resources.

If $\alpha = 0.001$, then starting from (2,2) the updates are minuscule, and you will need thousands of steps to approach the minimizer (0,0).

Optimal Learning Rate is fastest descent without overshooting and in quadratic problems like this one, if the Hessian matrix H is known with the optimal fixed learning rate as:

$$\alpha_{opt} = \frac{2}{\lambda_{max} + \lambda_{min}}$$

where $\lambda_{max} + \lambda_{min}$ are the eigenvalues of H

For our function:

$$H = \begin{bmatrix} 2 & 2 \\ -2 & 4 \end{bmatrix} \Rightarrow \text{Eigenvalues: } \lambda_{min} = 1.17, \lambda_{max} = 4.83$$

$$\Rightarrow \alpha_{opt} = \frac{2}{1.17 + 4.83} \approx 0.33$$

Using $\alpha = 0.3$ to 0.40 would yield rapid convergence.

Too Large Learning Rate ($\alpha > 1$) Overshooting: These updates are too large, causing the algorithm to jump past the minimum. It may cause the iterates to diverge or oscillate wildly and can even increase the objective value over time.

Discussions and Implications for Applications 4.10.1 and 4.10.2

Applications 4.10.1 and 4.10.2 investigated how learning rate strategies impact the convergence behaviour of Accelerated Mirror Gradient Techniques (AMGT) in convex and non-convex optimization tasks.

Application 4.10.1, studied gradient descent on convex quadratic functions, highlighting the significance of spectral properties in ensuring convergence. The study confirmed that convergence was guaranteed when the learning rate η satisfied $0 < \eta < 2/\lambda_{max}$, where $\lambda_{max} = 6$, giving a critical threshold of $\eta = 1/3$. Deviating from this threshold resulted in oscillation or divergence, consistent with prior theoretical analyses (Nesterov, 2004; Boyd & Vandenberghe, 2004). The elliptical shape of the loss function, characterized by eigenvalues 2 and 6, influenced convergence speed based on the chosen η value. These results are consistent with the previous findings on learning rate sensitivity (Bottou *et al.*, 2018; Ruder, 2016) and fulfilled the objective of promoting stability in AMGT's convergence phase through adaptive rate tuning.

Application 4.10.2, focused on the trade-off between speed and stability, highlighting the use of spectral information to select optimal learning rates for faster convergence while maintaining stability. It demonstrated that spectral information such as λ_{max} could be used to select efficient learning rates, achieving faster convergence without sacrificing stability. However, excessive rates still caused divergence, aligning with findings in non-convex optimization literature (Hardt *et al.*, 2016; Laisin and Adigwe, 2025a). The findings highlight the importance of spectral-aware global rate tuning in AMGT, aligning with existing literature on non-convex optimization challenges.

These findings suggest that in deep learning, it is crucial to use adaptive methods such as Adam and RMSProp, which adjust updates based on local gradients (Kingma and Ba, 2015). Additionally, SGD with momentum is known to have better generalization performance (Wilson *et al.*, 2017; Laisin and Adigwe, 2025b). Advanced schedulers like cosine annealing, warm restarts, and layer-wise scaling can mimic ideal convergence patterns (You *et al.*, 2020) to effectively navigate dynamic, non-convex landscapes.

Polyhedral optimization methods also underscore the significance of parameter tuning (Laisin *et al.*, 2025; Laisin and Edike, 2025). Future research directions could explore curvature-aware techniques such as AdaHessian and Shampoo to improve optimization in non-convex scenarios and gain a deeper understanding of non-convex landscapes.

Conclusions and Recommendations

Conclusions

This study presented the Adaptive Moment Gradient Thresholding (AMGT) algorithm, a mathematically refined variant of the Adam optimizer, developed to enhance convergence speed, numerical stability, and robustness across both convex and non-convex optimization tasks commonly encountered in deep neural networks. The primary objectives of AMGT were to ensure stable learning dynamics, establish global convergence guarantees, and maintain reliable performance in complex, high-dimensional optimization landscapes.

The proposed AMGT effectively fulfilled these objectives. Theoretical analysis established linear convergence under strong convexity with a rate of $O(k^{-\{\mu/2\}})$, where μ denotes the strong convexity constant—representing a significant improvement over

the standard Adam algorithm, which lacks such guarantees. Moreover, AMGT preserved convergence under inexact gradient conditions, an essential property for stochastic and noisy training environments, and achieved convergence in expectation toward stationary points in non-convex domains. These findings confirm AMGT's capability to deliver stable and consistent optimization performance across diverse problem settings.

Empirical evaluations on benchmark datasets—MNIST, CIFAR-10, and ImageNet—further validated the theoretical claims. AMGT consistently achieved faster and more stable convergence than Adam and related optimizers. On MNIST, it reached 99.2% accuracy with smoother loss trajectories; on CIFAR-10, it reduced total training time by approximately 15% while maintaining superior generalization; and on ImageNet, it exhibited faster early-epoch convergence and greater resilience to learning rate perturbations.

Collectively, these results demonstrate that integrating momentum thresholding with an adaptive step-size decay mechanism yields substantial improvements in training stability and optimization efficiency. The algorithm's flexible structure also supports seamless adaptation to various neural architectures and computational environments, making it particularly suitable for resource-constrained or distributed learning contexts, such as federated and edge computing.

In summary, AMGT bridges the gap between theoretical rigor and practical deployment by providing a theoretically grounded, computationally efficient, and empirically validated optimization framework. Its proven performance across multiple datasets and optimization regimes positions AMGT as a robust, scalable, and general-purpose alternative to existing adaptive optimizers, establishing a strong foundation for future extensions in large-scale and decentralized learning systems.

5.2 Recommendations: Include the following:

Practical Recommendations

1. **Use in Deep Learning:** Replace Adam with AMGT in critical tasks (e.g., NLP, computer vision, or symbolic reasoning) where training stability and convergence are essential.
2. **Learning Rate Scheduling:** Apply a diminishing step size $\alpha k = \frac{1}{\mu k}$ to ensure stable descent; tune μ based on the task.
3. **Combine with Regularization:** Use AMGT with regularization (such as weight decay, dropout, or batch normalization) to maintain bounded gradients and improve generalization (preserve the assumptions under which AMGT performs optimally).

Future Research Directions

1. **Adaptive Thresholding:** Explore momentum thresholding based on curvature or variance to improve non-convex performance.
2. **Second-Order Extensions:** Extend AMGT with second-order info (e.g., via AdaHessian) for sharper or ill-conditioned landscapes.
3. **Robustness to Noise:** Analyze AMGT's behaviour under noisy, non-stationary, or adversarial conditions.

Final remarks

In conclusion, the Adaptive Moment Gradient Thresholding (AMGT) algorithm represents a significant advancement toward more reliable and theoretically grounded optimization in deep learning. By combining rigorous convergence guarantees with strong empirical performance, AMGT offers a robust, scalable, and efficient framework suitable for both academic research and practical machine learning deployment.

References

- Alistarh, D., Grubic, D., Li, J., Tomioka, R., & Vojnovic, M. (2017). QSGD: Communication-efficient SGD via gradient quantization and encoding. *Advances in Neural Information Processing Systems*, 30.
- Bottou, L., Curtis, F. E., & Nocedal, J. (2018). *Optimization methods for large-scale machine learning*. SIAM Review, 60(2), 223–311. <https://doi.org/10.1137/16M1080173>
- Cauchy, A. L. (1847). *Méthode générale pour la résolution des systèmes d'équations simultanées*. Comptes Rendus, 25, 536–538.
- Chen, J., Zhang, H., Xu, X., & Yin, W. (2018). **On the convergence of a class of Adam-type algorithms for non-convex optimization**. *arXiv preprint arXiv:1808.02941*. <https://arxiv.org/abs/1808.02941>.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT* (pp. 4171–4186). <https://doi.org/10.48550/arXiv.1810.04805>.
- Duchi, J., Hazan, E., & Singer, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12, 2121–2159.
- Ghadimi, S., & Lan, G. (2013). Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4), 2341–2368. <https://doi.org/10.1137/130905661>.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). *Deep residual learning for image recognition*. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770–778). <https://doi.org/10.1109/CVPR.2016.90>.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1412.6980>.
- Laisin, M., Edike, C. and Bright O. Osu (2024); The construction of rational polyhedron on an $n \times n$ board with some application on integral polyhedral. TIJER, ISSN 2349-9249, Vol 11, Issue 11, www.tijer.org.
- Laisin, M. & Edike, C. (2025). Hybrid Optimization with Integer Constraints: Modeling and Solving Problems Using Simplex Techniques. *Global Online Journal of Academic Research (GOJAR)*, 4(2), 22-38. <https://klamidas.com/gojar-v4n2-2025-02/>.
- Laisin, M., & Adigwe, R. U. (2025a). Implementation and comparative analysis of AMGT method in Maple 24: Convergence performance in optimization problems. *Global Online Journal of Academic Research (GOJAR)*, 4(52), 26–40. <https://klamidas.com/gojar-v4n1-2025-02/>
- Laisin, M., Edike, C., & Ujumadu, R. N. (2025). Characterizing Boundedness and Solution Size in Rational Linear Programming and Polyhedrall Optimization. *Global Online Journal of Academic Research (GOJAR)*, 4(2), 63-76. <https://klamidas.com/gojar-v4n2 2025-04/>.
-
- Laisin, M., Osu, B. O., Duruojinkey, P.U, & Chibuisi, C. (2025). A Mathematically Modified Adam Algorithm for Improved Convergence in Deep Neural Networks. *Online Journal of Mathematics, Science and Technology Education (OJOMSTE)*, 6(2), 40-64.

- Laisin, M. & Adigwe, R. U. (2025b). Gradient Descent Convergence: From Convex Optimization to Deep Learning. *SOLVANGLE*, 1(1), 7-26. [https://klamidas.com/solvangle v1n1-2025-01/](https://klamidas.com/solvangle/v1n1-2025-01/)
- Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., & Han, J. (2020). On the variance of the adaptive learning rate and beyond. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1908.03265>.
- Luo, L., Xiong, Y., Liu, Y., & Sun, X. (2019). Adaptive gradient methods with dynamic bound of learning rate. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1902.09843>.
- Nesterov, Y. (2004). *Introductory lectures on convex optimization: A basic course*. Springer Science & Business Media.
- Nesterov, Y. E. (1983). A method for solving the convex programming problem with convergence rate $O(1/k^2)$. *Soviet Mathematics Doklady*, 27(2), 372–376.
- Nedic, A., & Bertsekas, D. P. (2001). Incremental subgradient methods for nondifferentiable optimization. *SIAM Journal on Optimization*, 12(1), 109–138. <https://doi.org/10.1137/S1052623499362822>.
- Polyak, B. T. (1964). Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5), 1–17.
- Reddi, S. J., Kale, S., & Kumar, S. (2018). *On the convergence of Adam and beyond*. In *Proceedings of ICLR*. <https://doi.org/10.48550/arXiv.1904.09237>.
- Robbins, H., & Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3), 400–407. <https://doi.org/10.1214/aoms/1177729586>.
- Schmidt, M., Le Roux, N., & Bach, F. (2011). Convergence rates of inexact proximal-gradient methods for convex optimization. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems* (Vol. 24). Curran Associates, Inc.
- Tieleman, T., & Hinton, G. (2012). Lecture 6.5—RMSProp: Divide the gradient by a running average of its recent magnitude. In *COURSERA: Neural Networks for Machine Learning*. https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf.
- Wilson, A. C., Roelofs, R., Stern, M., Srebro, N., & Recht, B. (2017). *The marginal value of adaptive gradient methods in machine learning*. In *Proceedings of NeurIPS* (pp. 4148–4158). https://proceedings.neurips.cc/paper_files/paper/2017/file/5d44ee6f2c3f71b73125876103c8f6c4-Paper.pdf.
- You, Y., Reddi, S., Hseu, J., Kumar, S., Bhojanapalli, S., Song, X., ... & Hsieh, C.-J. (2020). Large batch optimization for deep learning: Training BERT in 76 minutes. *International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1904.00962>.
- Zhuang, J., Tang, T., Ding, Y., Tatikonda, S., Dvornik, N., Papademetris, X., & Duncan, J. S. (2020). **AdaBelief optimizer: Adapting stepsizes by the belief in observed gradients**. *Advances in Neural Information Processing Systems*, 33, 18795–18806. <https://proceedings.neurips.cc/paper/2020/hash/1ede3d44f3efc4098a5a5ea0f4f74c30-Abstract.html>.

ETHICAL AND SCIENTIFIC PRINCIPLES STATEMENT OF RESPONSIBILITY

The author(s) declare that ethical principles and scientific citation principles were adhered to throughout the preparation of this study. OJOMSTE assumes no responsibility if a contrary finding occurs; all responsibility belongs to the authors.

STATEMENT OF RESEARCHERS' CONTRIBUTION TO THE ARTICLE

First author contribution rate: 25%

Second author contribution rate: 25%

3rd author contribution rate: 25%

4th author contribution rate: 25%