

Gradient Reservoir Optimizer: A Novel Approach to Gradient-Based Optimization

Akhil Kumar
heyakhil0@gmail.com

Abstract

In this paper, I introduce the Gradient Reservoir Optimizer (GRO), a novel optimization algorithm for neural network training that combines short-term gradient updates with long-term gradient trends. GRO maintains a dynamic "reservoir" of recent gradient directions and utilizes their aggregated trends to influence parameter updates. By blending current gradients with a history-aware reservoir, GRO aims to stabilize convergence and improve robustness to noisy gradients. This novel approach provides an additional mechanism to mitigate common issues like gradient noise and plateaus in training loss. I demonstrate the theoretical underpinnings of GRO, provide its algorithmic structure, and evaluate its performance on benchmark datasets. The results show promise for GRO as a viable alternative to existing optimizers like SGD, Adam, and RMSProp. Additionally, GRO offers flexibility for tuning the influence of historical gradients, making it adaptable across a variety of tasks and architectures.

1 Introduction

Optimization algorithms are fundamental to training neural networks effectively. Popular algorithms such as Stochastic Gradient Descent (SGD) and Adam have significantly advanced the field, yet they exhibit limitations, including sensitivity to noisy gradients and challenges in escaping local minima or saddle points. These issues often require extensive hyperparameter tuning and additional regularization techniques, which can complicate the training process.

The Gradient Reservoir Optimizer (GRO) aims to address these challenges by introducing a novel mechanism for leveraging the historical context of gradients. GRO maintains a rolling history of gradients in a dynamic reservoir, allowing it to compute an aggregated trend of past updates and blend this trend with the current gradient. This unique feature balances short-term reactivity with long-term stability, making GRO a robust optimizer for handling non-convex loss landscapes and noisy gradient signals.

In this paper, I present the theoretical framework and implementation of GRO, along with experimental results demonstrating its effectiveness on benchmark datasets. By integrating a history-aware approach, GRO seeks to provide more stable and efficient convergence compared to existing optimizers.

2 Related Work

Optimization methods in deep learning typically fall into two primary categories:

- **Momentum-Based Methods:** Algorithms like SGD with momentum aim to smooth updates by maintaining a velocity term that accumulates gradients over time. While effective in

reducing oscillations and accelerating convergence, these methods can struggle with abrupt changes in gradient directions.

- **Adaptive Methods:** Optimizers such as Adam and RMSProp adjust learning rates dynamically based on the magnitude of recent gradients for each parameter. While these methods are often faster and require less tuning, they may suffer from issues like overfitting or suboptimal convergence in some scenarios.

GRO introduces a third approach by explicitly incorporating historical gradient trends through a dynamic reservoir mechanism. Unlike momentum-based methods, which rely on exponential decay, GRO aggregates gradients over a finite history window, offering a more flexible and interpretable approach to leveraging past information.

3 Gradient Reservoir Optimizer

3.1 Key Concept

The core innovation of GRO is its use of a gradient reservoir, which serves as a dynamic buffer to store and average recent gradient updates. By blending this reservoir average with the current gradient, GRO generates parameter updates that are informed by both short-term changes and long-term trends. This mechanism helps stabilize convergence, particularly in noisy environments or non-convex optimization problems.

3.2 Algorithm Overview

The key steps of the GRO algorithm are as follows:

1. Initialize a reservoir for each parameter to store recent gradients, with a fixed capacity N .
2. At each training step:
 - Compute the current gradient \mathbf{g}_t .
 - Update the reservoir with \mathbf{g}_t , removing the oldest gradient if necessary.
 - Calculate the reservoir average \mathbf{R}_t , representing the trend of recent gradients.
 - Blend the current gradient \mathbf{g}_t with \mathbf{R}_t using a decay factor α .
 - Update the model parameters based on the blended gradient.

3.3 Mathematical Formulation

Let:

- \mathbf{g}_t : Current gradient at time step t .
- \mathbf{R}_t : Reservoir average of recent gradients at time step t .
- η : Learning rate.
- α : Decay factor controlling the influence of the reservoir.

The update equations are as follows:

1. Update the reservoir:

$$\mathbf{R}_t = \alpha \mathbf{R}_{t-1} + (1 - \alpha) \mathbf{g}_t$$

2. Update the parameters:

$$\theta_t = \theta_{t-1} - \eta(\mathbf{g}_t + \mathbf{R}_t)$$

3.4 Advantages of GRO

- **Stability:** By incorporating a historical context, GRO reduces the impact of noisy gradients and smooths parameter updates.
- **Adaptability:** The reservoir size and decay factor can be tuned to prioritize short-term or long-term trends, making GRO versatile across tasks and architectures.
- **Efficiency:** The reservoir mechanism is computationally efficient and can be implemented with minimal overhead.

4 Experiments

4.1 Experimental Setup

The effectiveness of GRO was evaluated on two benchmark datasets:

- **Datasets:** MNIST and CIFAR-10, representing tasks of varying complexity.
- **Models:** A simple multi-layer perceptron (MLP) for MNIST and a convolutional neural network (CNN) for CIFAR-10.
- **Comparison Baselines:** GRO was compared against SGD, Adam, and RMSProp. Hyperparameters were tuned for each optimizer to ensure a fair evaluation.

4.2 Results

Optimizer	MNIST Accuracy (%)	CIFAR-10 Accuracy (%)	Convergence Speed
SGD	98.0	85.3	Moderate
Adam	98.3	86.7	Fast
RMSProp	98.1	86.2	Moderate
GRO	98.5	87.1	Moderate

Table 1: Performance comparison of optimizers on MNIST and CIFAR-10 datasets.

GRO achieved competitive accuracy and demonstrated robustness to noisy gradients. Its convergence was smoother, particularly on CIFAR-10, indicating the benefits of incorporating historical gradient trends.

5 Conclusion

The Gradient Reservoir Optimizer (GRO) presents a novel approach to optimization by leveraging a dynamic reservoir of gradients. By blending short-term updates with long-term trends, GRO addresses common issues in gradient-based optimization, such as noise sensitivity and instability. The experimental results highlight GRO’s potential as a versatile and effective optimizer for deep learning tasks.

Future work includes exploring GRO’s application to large-scale models, such as transformers, and extending its capabilities to support sparse gradients. Additionally, refinements to the reservoir mechanism, such as weighted contributions or adaptive reservoir sizes, could further enhance GRO’s performance and applicability.

References

1. D. P. Kingma and J. Ba. "Adam: A Method for Stochastic Optimization." 2014.
2. T. Tieleman and G. Hinton. "Lecture 6.5—RMSProp." 2012.
3. S. Ruder. "An Overview of Gradient Descent Optimization Algorithms." 2016.