

Enhancing Linear Attention with Residual Learning

Xunhao Lai¹ Jialiang Kang² Jianqiao Lu³ Pengyu Zhao⁴ Tong Lin^{1*}

¹State Key Laboratory of General Artificial Intelligence,
School of Intelligence Science and Technology, Peking University, China

²Peking University, Beijing, China

³The University of Hong Kong, Hong Kong, China ⁴MiniMax, Shanghai, China

laixunhao@pku.edu.cn, jkang@stu.pku.edu.cn, jqlu@cs.hku.hk

pengyuzhao@pku.edu.cn, lintong@pku.edu.cn

Abstract

Linear attention offers a linear-time alternative to self-attention but often struggles to capture long-range patterns. We revisit linear attention through a prediction-correction lens and show that prevalent variants can be written as a combination of a historical prediction and a single-token correction, which creates an expressivity bottleneck. To address this bottleneck, we introduce **Residual Linear Attention (RLA)**, a framework that equips linear attention with an explicit *residual-fitting* mechanism. RLA maintains an auxiliary recurrent state that learns to accumulate residual errors over time and correct the base prediction. We further instantiate a delta-rule version, **Residual Delta Net (RDN)**, incorporating adaptive gating and residual clipping for enhanced correction control and stability. Our implementation leverages highly optimized linear attention kernels and preserves linear time and memory. Across language modeling and recall-intensive evaluations, RLA and RDN consistently outperform their respective baselines and other modern linear-attention methods, narrowing the gap to standard Transformers while retaining linear scaling.

1 Introduction

The Transformer (Vaswani et al., 2017) architecture has become the standard for large language models. However, the quadratic time complexity of its self-attention mechanism remains a critical bottleneck, limiting its application to long sequences (Li et al., 2025). Linear attention has recently emerged as an efficient alternative to standard self-attention, directly addressing its prohibitive quadratic complexity. By reformulating the attention computation into a recurrent process,

these models achieve linear-time training and inference, making them well-suited for processing long sequences. Architectures such as RetNet (Sun et al., 2023) and Mamba (Gu and Dao, 2024; Dao and Gu, 2024) have demonstrated competitive performances. Methods like GLA (Yang et al., 2024a) and DeltaNet (Yang et al., 2024b) offer further improvements by incorporating data-dependent gating and state update rules to manage the flow of information within a single state matrix.

Modern linear attention methods can be unified as learning a direct mapping from keys to values (Sun et al., 2025), a process analogous to test-time training. For example, the delta update rule (Schlag et al., 2021) can be derived from a single step of online gradient descent on a quadratic loss objective. This perspective opens several avenues for improvement. These include exploring different online learning loss functions to derive new update rules (Schlag et al., 2021; Yang et al., 2024b), employing more sophisticated mapping functions, or modifying the online gradient update mechanism (von Oswald et al., 2026; Siems et al., 2025). For instance, recent works like TTT-MLP (Sun et al., 2025) and Titans (Behrouz et al., 2025) utilize a Multi-Layer Perceptron (MLP) as a deep memory module to achieve a more powerful mapping. However, this approach sacrifices the model’s linear recurrence, thereby complicating parallel training.

Building on this perspective, we offer a new interpretation of the attention output. We show that the output of prevalent linear attention models can be decomposed into a base component generated from historical states and a correction term derived solely from the current token (see Section 2.3). Relying on a single token to perform this systematic correction imposes a bottleneck and is detrimental

*Corresponding author.

to the model’s expressive power. To address these issues, we introduce Residual Linear Attention, a framework that enhances linear attention models with an explicit residual fitting mechanism. Rather than depending on a single token for correction, our method employs an auxiliary state matrix to explicitly model and correct systematic prediction errors of the base linear attention. The final output is a combination of the base prediction and this learned error correction. Our approach can be generalized to a unified framework applicable to various linear attention methods, offering a powerful and efficient strategy for building more capable sequence models.

Building upon existing linear attention methods, we propose two variants enhanced with residual fitting: Residual Linear Attention (RLA) and Residual Delta Net (RDN). We evaluate them on a range of benchmarks, including language modeling and recall-intensive tasks. Our results demonstrate that these models outperform their respective baselines and other modern linear attention methods, while our ablation analysis confirms the importance of each key design choice within our framework.

2 Preliminaries

2.1 Linear Attention as a Recurrent Model

Softmax attention mechanisms exhibit quadratic computational complexity with respect to sequence length, constituting a significant bottleneck when processing long sequences. Linear attention (Katharopoulos et al., 2020) architectures address this by removing the softmax function, which allows for a reordering of the computation.

For the t -th token in a sequence, let the query, key, and value vectors be $\mathbf{q}_t \in \mathbb{R}^{d_q \times 1}$, $\mathbf{k}_t \in \mathbb{R}^{d_k \times 1}$, and $\mathbf{v}_t \in \mathbb{R}^{d_v \times 1}$, where d_q , d_k , and d_v are their respective feature dimensions, with $d_q = d_k$. After applying a kernel function $\phi(\cdot)$ to the queries and keys (omitted for simplicity in the notation), the causal linear attention output $\mathbf{o}_t$ can be expressed as:

$$\mathbf{o}_t = \sum_{i=1}^t \mathbf{v}_i \left(\mathbf{k}_i^\top \mathbf{q}_t \right) = \left(\sum_{i=1}^t \mathbf{v}_i \mathbf{k}_i^\top \right) \mathbf{q}_t.$$

By defining a state matrix $\mathbf{S}_t := \sum_{i=1}^t \mathbf{v}_i \mathbf{k}_i^\top \in \mathbb{R}^{d_v \times d_k}$, we arrive at the following recurrent formulation:

$$\mathbf{S}_t = \mathbf{S}_{t-1} + \mathbf{v}_t \mathbf{k}_t^\top, \quad \mathbf{o}_t = \mathbf{S}_t \mathbf{q}_t.$$

This recurrent form maintains constant time and memory complexity per step during inference and facilitates efficient training through chunk-wise parallel algorithms (Yang et al., 2024a). Furthermore, the use of gating mechanisms has led to the development of more variants such as RetNet (Sun et al., 2023), Lightning Attention (Qin et al., 2024a), and Mamba-2 (Dao and Gu, 2024).

2.2 An Online Learning Perspective

The design of the recurrent update rule can be motivated from an online learning perspective (Sun et al., 2025; Liu et al., 2025). In this view, the token sequence is a stream of data points $(\mathbf{k}_t, \mathbf{v}_t)$, and the state matrix $\mathbf{S}$ acts as model parameters. These parameters are updated online to learn the mapping $\mathbf{k} \mapsto \mathbf{v}$, with $\mathbf{S}$ functioning as a memory from which information is retrieved using the query $\mathbf{q}$ via $\mathbf{S}\mathbf{q}$.

The state update can be interpreted as one step of gradient descent on a loss function $\mathcal{L}(\mathbf{k}, \mathbf{v}; \mathbf{S})$. For instance, applying a single descent step with the loss $\mathcal{L}(\mathbf{k}_t, \mathbf{v}_t; \mathbf{S}) := -\langle \mathbf{S}\mathbf{k}_t, \mathbf{v}_t \rangle$ recovers the standard linear attention update:

$$\begin{aligned} \mathbf{S}_t &= \mathbf{S}_{t-1} - \nabla_{\mathbf{S}} \mathcal{L}(\mathbf{k}_t, \mathbf{v}_t; \mathbf{S}_{t-1}) \\ &= \mathbf{S}_{t-1} + \mathbf{v}_t \mathbf{k}_t^\top. \end{aligned}$$

An alternative update rule can be derived by minimizing a squared error loss, $\mathcal{L}(\mathbf{k}_t, \mathbf{v}_t; \mathbf{S}) := \frac{1}{2} \|\mathbf{S}\mathbf{k}_t - \mathbf{v}_t\|^2$. Performing one step of gradient descent on $\mathbf{S}_{t-1}$ with a data-dependent learning rate β_t yields the delta rule:

$$\begin{aligned} \mathbf{S}_t &= \mathbf{S}_{t-1} - \beta_t \nabla_{\mathbf{S}} \mathcal{L}(\mathbf{k}_t, \mathbf{v}_t; \mathbf{S}_{t-1}) \\ &= \mathbf{S}_{t-1} (I - \beta_t \mathbf{k}_t \mathbf{k}_t^\top) + \beta_t \mathbf{v}_t \mathbf{k}_t^\top. \end{aligned}$$

This formulation enables models like Delta Net (Yang et al., 2024b; Schlag et al., 2021) to achieve fine-grained memory control. Gated Delta Net (Yang et al., 2025) further enhances this approach by incorporating weight decay into the learning process.

2.3 Decomposition into Prediction and Correction

We interpret linear attention through a prediction-correction lens. The standard linear attention output, $\mathbf{o}_t = \mathbf{S}_t \mathbf{q}_t$, can be viewed as the sum of a base prediction from the past state and a correction

based on the current token:

$$\mathbf{o}_t = \underbrace{\mathbf{S}_{t-1}\mathbf{q}_t}_{\text{Base Prediction}} + \underbrace{(\mathbf{v}_t\mathbf{k}_t^\top)\mathbf{q}_t}_{\text{Error Correction}}.$$

We can generalize this decomposition to the form $\mathbf{o}_t = \mathbf{S}_{t-1}\mathbf{q}_t + \mathbf{R}_t\mathbf{q}_t$, where we introduce $\mathbf{R}_t$ as a generalized correction state. This framework provides a unified view of several methods, as shown in Table 1, which differ not only in the design of their associated state update but also in this correction term.

Building on the prediction-correction viewpoint, we introduce a residual fitting framework to enhance linear attention. Our framework learns a more expressive correction term by explicitly fitting on contextual information beyond the current token.

3 Method

This section presents our proposed method, which enhances linear attention through a residual-fitting process. We begin by describing the foundational residual learning framework that underpins our method. Next, we introduce an adaptive correction factor to enhance modeling capabilities and clipping methods to stabilize the residual fitting process. Finally, we present two final variants of our approach.

3.1 Explicit Residual Fitting

As established in Section 2.3, the output of linear attention can be decomposed into a base prediction and a correction term. To learn a more expressive correction, we introduce an auxiliary state, $\mathbf{R}_t$, which modifies the output formulation to $\mathbf{o}_t = \mathbf{S}_{t-1}\mathbf{q}_t + \mathbf{R}_t\mathbf{q}_t$. Crucially, unlike the standard correction shown in Table 1, which is derived solely from the current token, our auxiliary state $\mathbf{R}_t$ is updated recurrently, analogous to the primary state $\mathbf{S}_t$.

Inspired by gradient boosting (Friedman, 2001), an ensemble learning method that iteratively corrects errors made by existing models, we propose leveraging the residual errors in the linear attention mechanism to improve model performance. In this framework, each stage of learning identifies a new function h that minimizes the loss when added to the current model f :

$$h^* = \arg \min_h \mathbb{E}_{\mathbf{k}, \mathbf{v}} [\mathcal{L}(f(\mathbf{k}) + h(\mathbf{k}), \mathbf{v})].$$

Since directly solving for the optimal h is intractable, we approximate the loss function via a first-order Taylor expansion. This allows us to interpret the steepest descent direction in the function space as the negative gradient of the loss with respect to the model’s output. This direction is often referred to as the pseudo-residual $\mathbf{r}$:

$$\mathbf{r} = -\frac{\partial \mathcal{L}(f(\mathbf{k}) + h(\mathbf{k}), \mathbf{v})}{\partial h(\mathbf{k})} \approx -\frac{\partial \mathcal{L}(f(\mathbf{k}), \mathbf{v})}{\partial f(\mathbf{k})}.$$

Hence, the task becomes learning a function $h(\mathbf{k})$ that models this pseudo-residual. For the commonly used $L2$ loss, $\mathcal{L} = \frac{1}{2}\|\mathbf{v} - \hat{\mathbf{v}}\|^2$, the pseudo-residual simplifies to the actual residual error, $\mathbf{r}_t := \mathbf{v}_t - \mathbf{S}_{t-1}\mathbf{k}_t$. We then model the residual error through an auxiliary state $\mathbf{R}_t$. By adopting the online learning framework of linear attention, as described in Section 2, we apply a similar update rule to the auxiliary state, resulting in the following recurrent process:

$$\begin{aligned} \mathbf{S}_t &= \mathbf{S}_{t-1} + \mathbf{v}_t\mathbf{k}_t^\top && \text{(Base Update)} \\ \mathbf{r}_t &= \mathbf{v}_t - \mathbf{S}_{t-1}\mathbf{k}_t && \text{(Residual)} \\ \mathbf{R}_t &= \mathbf{R}_{t-1} + \mathbf{r}_t\mathbf{k}_t^\top && \text{(Residual Update)} \\ \mathbf{o}_t &= \mathbf{S}_{t-1}\mathbf{q}_t + \mathbf{R}_t\mathbf{q}_t && \text{(Output)} \end{aligned}$$

In this formulation, the auxiliary state $\mathbf{R}_t$ accumulates past residual errors and their corresponding keys. This allows it to model and correct for systematic prediction errors made by the base state $\mathbf{S}_{t-1}$, yielding a more expressive output. Substituting the residual error into the auxiliary update makes the interaction between the two states explicit:

$$\mathbf{R}_t = \mathbf{R}_{t-1} + \mathbf{v}_t\mathbf{k}_t^\top - \mathbf{S}_{t-1}\mathbf{k}_t\mathbf{k}_t^\top.$$

Assuming an initially empty auxiliary state, unrolling the recurrence gives $\mathbf{R}_t = \sum_{i=1}^t \mathbf{r}_i\mathbf{k}_i^\top$. Each outer product associates a residual error with its corresponding key, forming a persistent, key-addressed correction memory across tokens. Since each residual $\mathbf{r}_i$ is evaluated against the base state $\mathbf{S}_{i-1}$, the correction memory is explicitly conditioned on the base prediction. Accumulating these associations allows the correction to span multiple historical key directions and attain a rank greater than one. This residual fitting process can be further generalized into a unified framework for boosting linear attention.

Method	Output Combination	State Update Rule	Correction Term
LinearAttn	$\mathbf{o}_t = \mathbf{S}_{t-1}\mathbf{q}_t + \mathbf{R}_t\mathbf{q}_t$	$\mathbf{S}_t = \mathbf{S}_{t-1} + \mathbf{v}_t\mathbf{k}_t^\top$	$\mathbf{R}_t = \mathbf{v}_t\mathbf{k}_t^\top$
Mamba2	$\mathbf{o}_t = \alpha_t\mathbf{S}_{t-1}\mathbf{q}_t + \mathbf{R}_t\mathbf{q}_t$	$\mathbf{S}_t = \alpha_t\mathbf{S}_{t-1} + \mathbf{v}_t\mathbf{k}_t^\top$	$\mathbf{R}_t = \mathbf{v}_t\mathbf{k}_t^\top$
Gated LinearAttn	$\mathbf{o}_t = \text{diag}(\alpha_t)\mathbf{S}_{t-1}\mathbf{q}_t + \mathbf{R}_t\mathbf{q}_t$	$\mathbf{S}_t = \text{diag}(\alpha_t)\mathbf{S}_{t-1} + \mathbf{v}_t\mathbf{k}_t^\top$	$\mathbf{R}_t = \mathbf{v}_t\mathbf{k}_t^\top$
DeltaNet	$\mathbf{o}_t = \mathbf{S}_{t-1}\mathbf{q}_t + \beta_t\mathbf{R}_t\mathbf{q}_t$	$\mathbf{S}_t = \mathbf{S}_{t-1}(I - \beta_t\mathbf{k}_t\mathbf{k}_t^\top) + \beta_t\mathbf{v}_t\mathbf{k}_t^\top$	$\mathbf{R}_t = (\mathbf{v}_t - \mathbf{S}_{t-1}\mathbf{k}_t)\mathbf{k}_t^\top$
Gated DeltaNet	$\mathbf{o}_t = \alpha_t\mathbf{S}_{t-1}\mathbf{q}_t + \beta_t\mathbf{R}_t\mathbf{q}_t$	$\mathbf{S}_t = \alpha_t\mathbf{S}_{t-1}(I - \beta_t\mathbf{k}_t\mathbf{k}_t^\top) + \beta_t\mathbf{v}_t\mathbf{k}_t^\top$	$\mathbf{R}_t = (\mathbf{v}_t - \alpha_t\mathbf{S}_{t-1}\mathbf{k}_t)\mathbf{k}_t^\top$

Table 1: Comparison of different linear attention methods with base prediction and error correction.

It is worth noting two special cases of this formulation. If we set $\mathbf{R}_t = \mathbf{v}_t\mathbf{k}_t^\top$, using information only from the current token, our method reduces to standard linear attention (Katharopoulos et al., 2020). If we instead use $\mathbf{R}_t = (\mathbf{v}_t - \mathbf{S}_{t-1}\mathbf{k}_t)\mathbf{k}_t^\top$, the correction mechanism becomes equivalent to a one-step delta rule update (Schlag et al., 2021).

3.2 Adaptive Gating and Correction Factor

To enhance control over the state dynamics, we incorporate learnable gating scalars, a practice common in recent recurrent models (Yang et al., 2025; Dao and Gu, 2024). We introduce a decay factor $\alpha_t \in [0, 1]$ to control the retention of past information, and an update rate $\beta_t \in [0, 1]$ to modulate the influence of the current token. These factors can be applied to both state updates and output combinations:

$$\begin{aligned}\mathbf{S}_t &= \alpha_t\mathbf{S}_{t-1} + \beta_t\mathbf{v}_t\mathbf{k}_t^\top \\ \mathbf{R}_t &= \alpha_t\mathbf{R}_{t-1} + \beta_t\mathbf{r}_t\mathbf{k}_t^\top \\ \mathbf{o}_t &= \alpha_t\mathbf{S}_{t-1}\mathbf{q}_t + \beta_t\mathbf{R}_t\mathbf{q}_t\end{aligned}$$

However, using the same update rate β_t for both states couples the learning of the base representation and the error correction. To achieve more fine-grained control, we introduce a dedicated scalar correction factor, $\gamma_t \in [0, 1]$. This factor decouples the update processes and allows the model to dynamically scale the contribution of the residual correction term. The auxiliary state updates and output computation are given by:

$$\begin{aligned}\mathbf{S}_t &= \alpha_t\mathbf{S}_{t-1} + \beta_t\mathbf{v}_t\mathbf{k}_t^\top \\ \mathbf{R}_t &= \alpha_t\mathbf{R}_{t-1} + \gamma_t\mathbf{r}_t\mathbf{k}_t^\top \\ \mathbf{o}_t &= \alpha_t\mathbf{S}_{t-1}\mathbf{q}_t + \gamma_t\mathbf{R}_t\mathbf{q}_t\end{aligned}$$

This formulation uses the decay and correction factors to dynamically gate the retrieval from the base and auxiliary states, respectively.

3.3 Normalization and Residual Clipping

To ensure computational stability and robustness during training, we introduce two key mechanisms. First, we apply $L2$ normalization to the query and key vectors, improving numerical stability:

$$\begin{aligned}\text{norm}(\mathbf{q}) &= \frac{\mathbf{q}}{\sqrt{\sum_{i=1}^d q_i^2}}, \\ \text{norm}(\mathbf{k}) &= \frac{\mathbf{k}}{\sqrt{\sum_{i=1}^d k_i^2}}.\end{aligned}$$

Second, to prevent large gradient magnitudes, we modify the global loss function to a more robust variant. Let $e_i = v_i - \hat{v}_i$ denote the prediction error for the i -th dimension. This loss function uses the $L2$ loss for small errors and the $L1$ loss for large errors, ensuring that the gradient remains bounded:

$$\begin{aligned}\mathcal{L}(\hat{\mathbf{v}}, \mathbf{v}) &= \sum_{i=1}^d \ell_c(e_i), \\ \ell_c(e_i) &= \begin{cases} \frac{1}{2}e_i^2 & \text{for } |e_i| \leq c \\ c(|e_i| - \frac{1}{2}c) & \text{for } |e_i| > c \end{cases}.\end{aligned}$$

This modified loss function provides a smoother and more stable gradient, particularly for large prediction errors. According to our residual fitting framework, we derive the pseudo-residual as:

$$\mathbf{r}_t = \text{Clip}_{[-c, c]}(\mathbf{v}_t - \mathbf{S}_{t-1}\mathbf{k}_t).$$

By clipping the residuals, we prevent potential instability in the auxiliary state $\mathbf{R}_t$ and keep the error-correction update bounded. This mechanism improves training stability in the evaluated setting, particularly in the presence of large, transient errors in the base model’s predictions.

3.4 Final Formulations

The residual fitting principle is a general technique that can be integrated with various linear attention

backbones. By applying our residual mechanism to both the standard additive update rule and the delta update rule, we derive two powerful variants. This leads to our final models:

Residual Linear Attention (RLA):

$$\begin{aligned} r_t &= \text{Clip}_{[-c,c]}(v_t - S_{t-1}k_t) \\ R_t &= \alpha_t R_{t-1} + \gamma_t r_t k_t^\top \\ S_t &= \alpha_t S_{t-1} + \beta_t v_t k_t^\top \\ o_t &= \alpha_t S_{t-1} q_t + \gamma_t R_t q_t \end{aligned}$$

Residual Delta Net (RDN):

$$\begin{aligned} r_t &= \text{Clip}_{[-c,c]}(v_t - S_{t-1}k_t) \\ R_t &= \alpha_t R_{t-1}(I - \gamma_t k_t k_t^\top) + \gamma_t r_t k_t^\top \\ S_t &= \alpha_t S_{t-1}(I - \beta_t k_t k_t^\top) + \beta_t v_t k_t^\top \\ o_t &= \alpha_t S_{t-1} q_t + \gamma_t R_t q_t \end{aligned}$$

For brevity, the equations omit the L2 normalization and SiLU activation applied to query and key vectors. Regarding the adaptive gates, the decay factor α_t adopts the re-parameterization from Mamba-2 (Dao and Gu, 2024), while β_t and γ_t are computed via a linear projection followed by a sigmoid activation. The structure of our attention block is shown in Figure 1.

4 Experiment

4.1 Setup

Implementation To maximize efficiency, we implement our custom attention kernels in Triton (Tillet et al., 2019), building upon the flash-linear-attention library (Yang and Zhang, 2024). We exploit the fact that our state update rule is identical to linear attention’s, requiring only a minor modification to their kernel: we augment it to return both the attention result and the intermediate residual. This design allows the same highly optimized kernel to be reused across all residual-fitting stages, ensuring high throughput.

Model Settings We evaluate our model against several recent linear attention architectures, including Retentive Network (RetNet) (Sun et al., 2023), Mamba2 (Dao and Gu, 2024), and Gated Delta Net (GDN) (Yang et al., 2025). Additionally, we establish a baseline for RLA by evaluating scalar-gated linear attention (sGLA), a linear attention variant equipped with query-key normalization and scalar gates (α and β). In our main experiments,

we set the clipping threshold to $c = 1$. All models contain approximately 1.5 billion parameters and are trained on 100 billion tokens under identical conditions to ensure a fair comparison.

4.2 Main Results

Kernel Efficiency We benchmark our kernel’s runtime against linear attention baselines and FlashAttention (Dao et al., 2022; Dao, 2024), as shown in Figure 2. Although the residual fitting process adds computational overhead, our method’s runtime scales linearly with sequence length. This makes it significantly faster than FlashAttention, which scales quadratically, on longer sequences. Regarding throughput, our method, like other linear attention mechanisms, maintains a nearly constant high throughput. Conversely, the throughput of the compute-bound FlashAttention degrades rapidly as sequence length increases. The auxiliary state has the same dimensions as the primary state, resulting in twice the per-layer recurrent-state storage of the corresponding backbone. RLA and RDN thus trade a constant-factor increase in recurrent-state storage and computation for a more expressive correction while retaining linear sequence scaling and recurrent inference.

Language Modeling & Commonsense Reasoning

We evaluate RLA and RDN on WikiText (Merity et al., 2017) perplexity and a suite of benchmarks assessing reasoning and commonsense understanding. The reasoning tasks include ARC-Easy, ARC-Challenge (Clark et al., 2018), PIQA (Bisk et al., 2020), and MMLU (Hendrycks et al., 2021), while commonsense understanding is evaluated on HellaSwag (Zellers et al., 2019), Winogrande (Sakaguchi et al., 2020), SocialIQA (Sap et al., 2019), and LAMBADA (Paperno et al., 2016). Our main results, summarized in Table 2, show that our proposed residual learning variants, RLA and RDN, achieve a consistent improvement in perplexity over their respective baselines, sGLA and GDN. In addition, our models outperform other leading linear attention methods across multiple benchmarks and deliver performance competitive with a standard Transformer.

Recall-intensive tasks To evaluate memory capacity, we benchmark our model on the recall-intensive tasks from Arora et al. (2024). In addition, we also directly evaluate the model’s retrieval ability using the "Needle-in-a-Haystack"




Figure 1: The architecture of our proposed model. The model structure (left) consists of N stacked blocks. The detailed Attention Block (right) illustrates our core mechanism. Our primary contribution, the explicit residual fitting process, is highlighted in purple dash lines. This path computes the clipped residual $r_t = \text{Clip}(v_t - S_{t-1}k_t)$, which is then modulated by a dedicated correction factor $\gamma_t = \sigma(W_\gamma x)$ to dynamically correct the base prediction from the model’s primary state. The model also utilizes gates $\alpha_t = \exp(-a \text{softplus}(W_\alpha x + b))$ and $\beta_t = \sigma(W_\beta x)$ to control the state dynamics, where a and b are learnable scalars.




Figure 2: Comparison of attention kernel computation time (left) and model throughput (right) with respect to sequence length.

task (NIAH) (gkamradt, 2023), which requires retrieving key-value pairs inserted at varying depths within a long document. These benchmarks are challenging for linear attention models because their finite state-space creates an information bottleneck, as shown in Table 3. Results demonstrate that our proposed RLA and RDN consistently outperform their corresponding baselines, with particularly strong gains on the DROP and FDA benchmarks. Furthermore, they substantially outperform other models on the NIAH task, highlighting an enhanced capacity for information recall.

4.3 Ablation Study

In this section, we present a series of ablation studies to verify the contributions of key components. We first quantify the advantage of our learned residual fitting approach over a predefined correction. Next, we investigate the importance of using a dedicated correction factor, followed by an analysis of the necessity of the gated mechanism for combining the base prediction and the correction. Finally, we examine the effect of normalization and residual clipping.

Residual Fitting To validate the importance of accumulating past errors, we test a variant that uses a simpler, predefined correction term. In this ab-

Model	Wiki ppl	LAMB ppl	ARC-C acc_n	ARC-E acc	Hella. acc_n	LAMB acc	MMLU acc_n	PIQA acc	SIQA acc	Wino acc	Avg
Transformer	17.33	19.53	30.2	55.7	49.0	44.3	31.1	70.4	37.9	53.5	46.51
RetNet	18.86	27.62	29.4	56.0	40.5	36.5	30.3	69.9	36.0	50.9	43.69
Mamba2	18.42	20.80	29.1	<u>57.4</u>	46.2	42.7	30.9	69.4	37.6	<u>51.0</u>	45.54
sGLA	17.63	18.06	30.1	56.9	46.4	44.6	30.9	70.4	36.7	50.3	45.79
GDN	<u>17.27</u>	15.76	<u>30.8</u>	54.0	46.8	44.0	<u>31.3</u>	<u>71.3</u>	<u>38.1</u>	51.9	46.03
RLA (ours)	17.35	<u>15.59</u>	30.6	56.6	48.1	<u>46.3</u>	31.0	70.7	38.5	49.7	<u>46.44</u>
RDN (ours)	16.57	14.93	32.1	58.7	<u>47.7</u>	48.7	31.6	71.7	37.7	49.5	47.20

Table 2: Language modeling, reasoning, and commonsense understanding results. We report perplexity (lower is better) and accuracy (higher is better). The **bold/underlined** numbers indicate the first/second best values of the linear attention models in each column.

Model	DROP	FDA	NQ	SQD	SWDE	TQA	NIAH	Avg
Transformer	26.9	63.0	29.8	35.7	68.7	43.6	70.5	48.31
RetNet	26.2	35.3	20.8	31.5	44.1	39.9	65.7	37.64
Mamba2	26.2	41.2	23.6	<u>33.0</u>	<u>63.4</u>	<u>43.2</u>	67.2	42.54
sGLA	25.7	<u>51.8</u>	24.8	31.6	<u>63.4</u>	41.5	76.6	45.06
GDN	25.9	47.4	26.8	32.4	63.3	43.5	75.7	44.99
RLA (ours)	<u>26.5</u>	51.5	25.2	33.6	64.4	42.3	83.6	<u>46.73</u>
RDN (ours)	27.8	57.5	<u>26.3</u>	32.5	<u>63.4</u>	43.1	<u>79.2</u>	47.11

Table 3: Accuracy on recall-intensive benchmarks. The **bold/underlined** numbers indicate the first/second best values of the linear attention models in each column.

Model	Wiki	LAMB	GSM8k	HumanEval
RLA	18.76	23.44	3.92	9.61
RLA w/o fitting	20.19	22.50	6.85	16.23

Table 4: Ablation study of the residual fitting process, comparing perplexity across various datasets. The best results are shown in **bold**.

lation, we replace our persistent auxiliary state, R_t , with a stateless correction derived only from the current residual, $R_t = (v_t - S_{t-1}k_t)k_t^\top$. As demonstrated in Table 4, the variant lacking explicit residual fitting underperforms our full method. Although this ablated variant maintains competitive performance on some benchmarks, it exhibits a substantial increase in perplexity on both the training and evaluation sets. This performance drop extends to specialized domains, with a substantial degradation in its math and code abilities, as measured by perplexity on GSM8k (Cobbe et al., 2021) and HumanEval (Chen et al., 2021). This demonstrates the critical role of the auxiliary state in accumulating past residuals to refine the model’s output effectively.

Dedicated Correction Factor We analyze the benefit of using a dedicated correction factor, γ , by comparing our full models against variants where γ is tied to the update factor β . In Figure 3a, the models with an independent γ consistently achieve lower evaluation loss, with the RDN variant showing greater improvement. This trend extends to downstream performance, as demonstrated by the results in Figure 3b, which also show that the dedicated correction factor yields performance gains across multiple benchmarks. Notably, our foundational architecture, which does not require an additional γ , still marks a notable improvement over the baseline linear attention method.

Gated Output Combination We conducted an ablation study to examine the effect of the output combination strategy. Specifically, we compared our full model, which uses a gated combination of the base prediction and error correction ($o_t = \alpha_t S_{t-1}q_t + \gamma_t R_t q_t$), with a variant that employs simple addition ($o_t = S_{t-1}q_t + R_t q_t$). As shown in Table 5, the removal of the gating mechanism resulted in only marginal performance changes. These results suggest that, while the gated output combination provides performance benefits in certain scenarios, the primary advantage of our




Figure 3: Ablation study on the correction factor γ . Using a dedicated γ consistently lowers validation loss compared to tying it to β . The evaluation uses the same benchmarks as in Section 4.2, divided into three task types, and confirms that a dedicated γ improves performance across several categories.

model stems from the residual fitting process applied to the auxiliary state R_t , rather than the specific weights used to integrate the correction term.

Normalization and Residual Clipping Finally, we investigate the importance of normalization and residual clipping. We perform an ablation study on RLA by removing normalization and clipping. As shown in Figure 4, both components are crucial for stable training; their removal leads to unbounded activations and degraded performance. In contrast, the RDN model is largely insensitive to residual clipping. This robustness is attributable to the inherent stability of its delta rule update, which maintains a consistent loss curve without residual clipping (Figure 4b).

5 Related Works

Sequence modeling has been historically dominated by Recurrent Neural Networks (RNNs) (Lipton et al., 2015), including variants like Long Short-Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997) and Gated Recurrent Units (GRU) (Cho et al., 2014). While effective, their inherently sequential nature impedes training parallelization. The Transformer architecture (Vaswani et al., 2017) overcame this limitation, emerging as the de facto standard for sequence modeling. However, its self-attention mechanism, with a computational complexity quadratic in sequence length, poses a significant bottleneck for long-context applications.

To address these challenges, recent research has revisited Linear RNNs as a foundation for efficient Transformer alternatives. By formulating sequence processing as a linear recurrence, these models


achieve both parallelizable training and linear-time inference. Early explorations in this domain, such as S4 (Gu et al., 2022), LRU (Orvieto et al., 2023), and RetNet (Sun et al., 2023), utilized structured state transition matrices. A subsequent performance leap was achieved by incorporating data-dependent dynamics. Models like Mamba (Gu and Dao, 2024; Dao and Gu, 2024), HGRN (Qin et al., 2023, 2024b), and Gated Linear Attention (Yang et al., 2024a) leverage input-dependent gating to dynamically control state transitions, thereby enhancing their expressive power.

More advanced methods have introduced the delta learning rule, which reframes the state update from a simple gated decay to a fine-grained memory correction. This approach, exemplified by DeltaNet (Yang et al., 2024b; Schlag et al., 2021) and Gated DeltaNet (Yang et al., 2025), enables more precise dynamic memory modifications. This mechanism can be interpreted from an online learning perspective, where the state update is framed as an optimization process, as explored in TTT (Sun et al., 2025). This viewpoint has inspired further work aimed at discovering and improving the intrinsic learning algorithms within sequence models (von Oswald et al., 2023, 2026).


Concurrent research has focused on increasing the expressivity of the state transition. For instance, RWKV-7 (Peng et al., 2025) employs a diagonal-plus-low-rank structure, while DeltaProduct (Siems et al., 2025) generalizes DeltaNet by performing multiple update steps per token. These updates enrich the transition of a single state within each token. Our residual fitting framework explores a complementary direction by maintaining a correc-

	Output Combination	ARC-C	ARC-E	Hella.	LAMB	MMLU	PIQA	SIQA	Wino	Avg
RLA	$\alpha_t S_{t-1} q_t + \gamma_t R_t q_t$	28.5	53.6	42.0	36.8	29.8	68.6	38.1	49.2	43.30
	$S_{t-1} q_t + R_t q_t$	28.9	55.2	41.5	37.0	30.1	68.9	37.3	51.6	43.81
RDN	$\alpha_t S_{t-1} q_t + \gamma_t R_t q_t$	29.8	55.2	42.5	39.9	30.4	69.1	40.4	49.9	44.65
	$S_{t-1} q_t + R_t q_t$	27.8	56.3	41.9	39.7	29.4	70.0	37.7	48.9	43.96

Table 5: Ablation study of different output combination methods. All models share the same 50B tokens pretraining and hyperparameters, with the best results for each method shown in **bold**.



(a) L2 norm of attention output



(b) Training loss

Figure 4: Ablation study on normalization and residual clipping. RLA variants without normalization or clipping exhibit exploding activation norms, indicating training instability. This instability leads to a higher training loss, highlighting that both components are crucial for stable training and better performance. In contrast, residual clipping has a negligible impact on the RDN training process.

tion state that accumulates residual information across tokens. Recent architectures such as Titans (Behrouz et al., 2025) and Miras (Behrouz et al., 2026) further increase memory capacity through non-linear deep memory parameterized by an MLP. RLA and RDN retain matrix-valued linear-attention states and closed-form recurrent updates, supporting associative scans and optimized linear-attention kernels.

6 Conclusion

In this paper, we introduced Residual Linear Attention, a framework that enhances linear attention

models with an explicit residual fitting process. Our method leverages an auxiliary state to correct the predictive errors of the base model, thereby building more robust and accurate contextual representations. The framework is highly adaptable and can be applied to various linear attention methods. Our experiments demonstrated this versatility, showing that our approach consistently outperforms their respective baselines. While this improvement comes at the cost of additional computation for the fitting process, balancing this trade-off offers a promising direction for future research.

Limitations

The residual fitting process introduces an auxiliary state of the same dimensions as the primary state, effectively doubling per-layer state memory and increasing wall-clock time compared to standard linear attention. All experiments are conducted at a single scale of approximately 1.5B parameters trained on 100B tokens; how the benefits transfer to substantially larger or smaller models remains unexplored. Additionally, the clipping threshold c is fixed at 1 without a systematic study of its sensitivity across different settings. Further analysis could compare against widened or dual-state baselines with matched recurrent-state capacity and probe how the auxiliary state retains and forgets residual information. Finally, our evaluation is limited to autoregressive language modeling and NLP benchmarks; generalization to other modalities such as vision or speech has not been investigated.

References

- Simran Arora, Aman Timalina, Aaryan Singhal, Benjamin Spector, Sabri Eyuboglu, Xinyi Zhao, Ashish Rao, Atri Rudra, and Christopher Ré. 2024. Just read twice: closing the recall gap for recurrent language models. *arXiv preprint arXiv:2407.05483*.
- Ali Behrouz, Meisam Razaviyayn, Peilin Zhong, and Vahab Mirrokni. 2026. *It’s all connected: A journey through test-time memorization, attentional bias,*

- retention, and online optimization. In *International Conference on Learning Representations*.
- Ali Behrouz, Peilin Zhong, and Vahab Mirrokni. 2025. [Titans: Learning to memorize at test time](#). In *Advances in Neural Information Processing Systems*, volume 38.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. 2020. [PIQA: Reasoning about physical commonsense in natural language](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7432–7439.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. [Evaluating large language models trained on code](#). *arXiv preprint arXiv:2107.03374*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1724–1734. Association for Computational Linguistics.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *arXiv preprint arXiv:2110.14168*.
- Tri Dao. 2024. [FlashAttention-2: Faster attention with better parallelism and work partitioning](#). In *International Conference on Learning Representations*.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in neural information processing systems*, 35:16344–16359.
- Tri Dao and Albert Gu. 2024. [Transformers are SSMS: Generalized models and efficient algorithms through structured state space duality](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 10041–10071. PMLR.
- Jerome H Friedman. 2001. [Greedy function approximation: a gradient boosting machine](#). *Annals of statistics*, 29(5):1189–1232.
- gkamradt. 2023. Llmtest needle in a haystack - pressure testing llms. https://github.com/gkamradt/LLMTest_NeedleInAHaystack.
- Albert Gu and Tri Dao. 2024. [Mamba: Linear-time sequence modeling with selective state spaces](#). In *First Conference on Language Modeling*.
- Albert Gu, Karan Goel, and Christopher Ré. 2022. [Efficiently modeling long sequences with structured state spaces](#). In *International Conference on Learning Representations*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural computation*, 9(8):1735–1780.
- Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. 2020. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pages 5156–5165. PMLR.
- Haoyang Li, Yiming Li, Anxin Tian, Tianhao Tang, Zhanchao Xu, Xuejia Chen, Nicole Hu, Wei Dong, Qing Li, and Lei Chen. 2025. [A survey on large language model acceleration based on kv cache management](#). *Transactions on Machine Learning Research*.
- Zachary C Lipton, John Berkowitz, and Charles Elkan. 2015. A critical review of recurrent neural networks for sequence learning. *arXiv preprint arXiv:1506.00019*.
- Bo Liu, Rui Wang, Lemeng Wu, Yihao Feng, Peter Stone, and Qiang Liu. 2025. [Longhorn: State space models are amortized online learners](#). In *International Conference on Learning Representations*.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *International Conference on Learning Representations*.
- Antonio Orvieto, Samuel L Smith, Albert Gu, Anushan Fernando, Caglar Gulcehre, Razvan Pascanu, and Soham De. 2023. Resurrecting recurrent neural networks for long sequences. In *International Conference on Machine Learning*, pages 26670–26698. PMLR.
- Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. 2016. [The lambada dataset: Word prediction requiring a broad discourse context](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1525–1534. Association for Computational Linguistics.

- Bo Peng, Ruichong Zhang, Daniel Goldstein, Eric Alcaide, Xingjian Du, Haowen Hou, Jiaju Lin, Jiaxing Liu, Janna Lu, William Merrill, Guangyu Song, Kaifeng Tan, Saiteja Utpala, Nathan Wilce, Johan S. Wind, Tianyi Wu, Daniel Wuttke, and Christian Zhou-Zheng. 2025. [RWKV-7 “goose” with expressive dynamic state evolution](#). In *Second Conference on Language Modeling*.
- Zhen Qin, Weigao Sun, Dong Li, Xuyang Shen, Weixuan Sun, and Yiran Zhong. 2024a. [Various lengths, constant speed: Efficient language modeling with lightning attention](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 41517–41535. PMLR.
- Zhen Qin, Songlin Yang, Weixuan Sun, Xuyang Shen, Dong Li, Weigao Sun, and Yiran Zhong. 2024b. [HGRN2: Gated linear RNNs with state expansion](#). In *First Conference on Language Modeling*.
- Zhen Qin, Songlin Yang, and Yiran Zhong. 2023. Hierarchically gated recurrent neural network for sequence modeling. *Advances in Neural Information Processing Systems*, 36:33202–33221.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. [Winogrande: An adversarial winograd schema challenge at scale](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8732–8740.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. 2019. [Social IQa: Commonsense reasoning about social interactions](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 4463–4473. Association for Computational Linguistics.
- Imanol Schlag, Kazuki Irie, and Jürgen Schmidhuber. 2021. Linear transformers are secretly fast weight programmers. In *International conference on machine learning*, pages 9355–9366. PMLR.
- Julien Siems, Timur Carstensen, Arber Zela, Frank Hutter, Massimiliano Pontil, and Riccardo Grazi. 2025. [DeltaProduct: Improving state-tracking in linear RNNs via householder products](#). In *Advances in Neural Information Processing Systems*, volume 38.
- Yu Sun, Xinhao Li, Karan Dalal, Jiarui Xu, Arjun Vikram, Genghan Zhang, Yann Dubois, Xinlei Chen, Xiaolong Wang, Sanmi Koyejo, Tatsunori Hashimoto, and Carlos Guestrin. 2025. [Learning to \(Learn at test time\): RNNs with expressive hidden states](#). In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 57503–57522. PMLR.
- Yutao Sun, Li Dong, Shaohan Huang, Shuming Ma, Yuqing Xia, Jilong Xue, Jianyong Wang, and Furu Wei. 2023. Retentive network: A successor to transformer for large language models. *arXiv preprint arXiv:2307.08621*.
- Philippe Tillet, Hsiang-Tsung Kung, and David Cox. 2019. [Triton: an intermediate language and compiler for tiled neural network computations](#). In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, pages 10–19.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Johannes von Oswald, Nino Scherrer, Seijin Kobayashi, Luca Versari, Songlin Yang, Maximilian Schlegel, Kaitlin Maile, Yanick Schimpf, Oliver Sieberling, Alexander Meulemans, Guillaume Lajoie, Rif A. Saurous, Charlotte Frenkel, Razvan Pascanu, Blaise Aguera y Arcas, and João Sacramento. 2026. [MesaNet: Sequence modeling by locally optimal test-time training](#). In *International Conference on Learning Representations*.
- Johannes von Oswald, Maximilian Schlegel, Alexander Meulemans, Seijin Kobayashi, Eyvind Niklasson, Nicolas Zuchet, Nino Scherrer, Nolan Miller, Mark Sandler, Blaise Aguera y Arcas, Max Vladymyrov, Razvan Pascanu, and João Sacramento. 2023. [Uncovering mesa-optimization algorithms in transformers](#). *arXiv preprint arXiv:2309.05858*.
- Songlin Yang, Jan Kautz, and Ali Hatamizadeh. 2025. [Gated delta networks: Improving Mamba2 with delta rule](#). In *International Conference on Learning Representations*.
- Songlin Yang, Bailin Wang, Yikang Shen, Rameswar Panda, and Yoon Kim. 2024a. [Gated linear attention transformers with hardware-efficient training](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 56501–56523. PMLR.
- Songlin Yang, Bailin Wang, Yu Zhang, Yikang Shen, and Yoon Kim. 2024b. Parallelizing linear transformers with the delta rule over sequence length. *Advances in neural information processing systems*, 37:115491–115522.
- Songlin Yang and Yu Zhang. 2024. [Fla: A triton-based library for hardware-efficient implementations of linear attention mechanism](#).
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a machine really finish your sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800. Association for Computational Linguistics.

A Unified Formula for Residual Linear Attention

This section details how our residual fitting process can be framed as a form of online gradient boosting, providing a unified and extensible framework.

A.1 Preliminary: Gradient Boosting in a Functional Space

Gradient boosting is an ensemble technique that sequentially adds new models to correct the errors of previous ones. In a gradient boosting framework, the objective at each stage is to find a new function h that minimizes the loss when added to the current function f :

$$h^* = \arg \min_h \mathbb{E}_{\mathbf{k}, v} [\mathcal{L}(f(\mathbf{k}) + h(\mathbf{k}), v)].$$

As finding the optimal function h is generally infeasible, we approximate the objective using a first-order Taylor expansion of the loss:

$$\begin{aligned} \mathcal{L}(f(\mathbf{k}) + h(\mathbf{k}), v) \\ \approx \mathcal{L}(f(\mathbf{k}), v) + h(\mathbf{k}) \frac{\partial \mathcal{L}(f(\mathbf{k}), v)}{\partial f(\mathbf{k})}. \end{aligned}$$

The direction of steepest descent in this functional space is the negative gradient of the loss with respect to the function's output. This target is often called a pseudo-residual, $\mathbf{r}$:

$$\mathbf{r} = -\frac{\partial \mathcal{L}(f(\mathbf{k}) + h(\mathbf{k}), v)}{\partial h(\mathbf{k})} \approx -\frac{\partial \mathcal{L}(f(\mathbf{k}), v)}{\partial f(\mathbf{k})}.$$

The objective thus becomes learning a function $h(\mathbf{k})$ that fits this pseudo-residual. After learning, the boosted function is updated as $f(\mathbf{k}) \leftarrow f(\mathbf{k}) + h(\mathbf{k})$, where $h(\mathbf{k})$ is the error correction term.

A.2 Unified Residual Linear Attention

From an online learning perspective, linear attention can be viewed as learning a mapping $f(\mathbf{k}; \mathbf{S})$, where the state matrix $\mathbf{S}$ is incrementally updated via online gradient descent to minimize a loss $\mathcal{L}(f(\mathbf{k}; \mathbf{S}), v)$. We enhance this model by incorporating principles from gradient boosting. This involves employing an auxiliary state matrix $\mathbf{R}$ for iterative refinement. In this framework, state matrix $\mathbf{R}$ is updated at timestep t to learn the mapping from the key $\mathbf{k}_t$ to the pseudo-residual $\mathbf{r}_t$ of the prior prediction. This correction process results in a stronger mapping function.

We can decouple the learning objective into two parts: (1) A global objective, defined by an arbitrary, differentiable outer loss $\mathcal{L}_{\text{outer}}$, which sets

the overall key-to-value mapping goal. (2) A local objective, defined by a simple inner loss $\mathcal{L}_{\text{inner}}$, which governs how each individual state matrix is updated. While the target pseudo-residual $\mathbf{r}$ can be complex, the task for each state is deliberately kept simple. Ignoring decay factors and learning rates for clarity, the general recurrence is:

$$\begin{aligned} \mathbf{r}_t &= -\frac{\partial \mathcal{L}_{\text{outer}}(f(\mathbf{k}_t; \mathbf{S}_{t-1}), v_t)}{\partial f(\mathbf{k}_t; \mathbf{S}_{t-1})} \\ \mathbf{R}_t &= \mathbf{R}_{t-1} - \frac{\partial \mathcal{L}_{\text{inner}}(f(\mathbf{k}_t; \mathbf{R}_{t-1}), \mathbf{r}_t)}{\partial \mathbf{R}_{t-1}} \\ \mathbf{S}_t &= \mathbf{S}_{t-1} - \frac{\partial \mathcal{L}_{\text{inner}}(f(\mathbf{k}_t; \mathbf{S}_{t-1}), v_t)}{\partial \mathbf{S}_{t-1}} \\ \mathbf{o}_t &= f(\mathbf{q}_t; \mathbf{S}_{t-1}) + f(\mathbf{q}_t; \mathbf{R}_t) \end{aligned}$$

The gating mechanism and correction strength can also be easily incorporated into the framework. This framework allows two simple inner update rules to approximate a more complex global objective.

B Residual Fitting with Huber Loss

As previously established, our framework can accommodate complex global loss functions, as their complexity is confined to the pseudo-residual calculation. This allows us to use a more robust alternative to the standard L2 loss, such as the Huber loss function. Let $e_i = v_i - \hat{v}_i$ denote the prediction error for the i -th dimension. Then,

$$\begin{aligned} \mathcal{L}_{\text{huber}}(\hat{v}, v) &= \sum_{i=1}^d \ell_c(e_i), \\ \ell_c(e_i) &= \begin{cases} \frac{1}{2}e_i^2 & \text{for } |e_i| \leq c \\ c(|e_i| - \frac{1}{2}c) & \text{for } |e_i| > c \end{cases} \end{aligned}$$

This function uses the L2 loss for small errors and the L1 loss for large errors, making it a more robust alternative. Directly applying this loss yields a non-linear update rule that is difficult to parallelize:

$$\begin{aligned} \mathbf{S}_t &= \mathbf{S}_{t-1} - \frac{\partial \mathcal{L}_{\text{huber}}(v_t - \mathbf{S}_{t-1}\mathbf{k}_t)}{\partial \mathbf{S}_{t-1}} \\ &= \mathbf{S}_{t-1} + \text{Clip}_{[-c, c]}(v_t - \mathbf{S}_{t-1}\mathbf{k}_t)\mathbf{k}_t^\top \end{aligned}$$

Our residual fitting framework elegantly avoids this problem. The complexity is isolated within the

pseudo-residual calculation, while the core state update remains simple. The pseudo-residual for the Huber Loss is $\mathbf{r}_t = \text{Clip}_{[-c,c]}(\mathbf{v}_t - f(\mathbf{k}_t; \mathbf{S}_{t-1}))$, then the inner update rule is only responsible for fitting this target pseudo-residual. Using an inner loss of $\mathcal{L}_{\text{inner}}(f(\mathbf{k}), \mathbf{v}) = -\langle \mathbf{v}, f(\mathbf{k}) \rangle$ yields the following recurrence:

$$\mathbf{r}_t = \text{Clip}_{[-1,1]}(\mathbf{v}_t - \mathbf{S}_{t-1}\mathbf{k}_t)$$

$$\mathbf{R}_t = \mathbf{R}_{t-1} + \mathbf{r}_t\mathbf{k}_t^\top$$

$$\mathbf{S}_t = \mathbf{S}_{t-1} + \mathbf{v}_t\mathbf{k}_t^\top$$

This equivalence provides the theoretical motivation for our clipping mechanism; it is an efficient implementation of a robust Huber loss objective, which leads to a more stable residual fitting process. This principle can be generalized by selecting other robust loss functions. For instance, the Log-Cosh loss has a negative gradient with respect to the prediction $f(\mathbf{k})$ that is equivalent to applying a tanh function to the residual:

$$\mathcal{L}_{\log\text{-cosh}}(f(\mathbf{k}), \mathbf{v}) = \log \cosh(\mathbf{v} - f(\mathbf{k})),$$

$$\mathbf{r}_t = \tanh(\mathbf{v}_t - f(\mathbf{k}_t; \mathbf{S}_{t-1})).$$

This can be viewed as a smooth alternative to the clipped residual from the Huber loss. The ability to easily substitute such loss functions demonstrates our framework’s modularity, allowing for the integration of powerful learning objectives while maintaining computational efficiency.

C Model Structure and Training Hyper Parameters

We evaluate several model architectures, with full specifications detailed in Table 6. Our comparison includes a standard Transformer and Transformers with pure linear attention. For a fair comparison, all models are trained using an identical set of hyperparameters, which are listed in Table 7. We initialize the model weights from a normal distribution with a constant standard deviation and use the AdamW optimizer with a cosine learning rate schedule for training.

Property	Transformer	Linear Attn
Total Params	1.51B	1.55B
Hidden Size	2048	2048
Intermediate Size	8192	8192
Attention Heads	32	16
GQA Groups	4	16
Head Dimension	128	128
Softmax Attn Layers	16	0
Linear Attn Layers	0	16

Table 6: Model Configuration

Peak LR	Min LR	Batch Size	Warmup
3e-4	3e-5	4M	0.5B

Total Tokens	Weight Decay	Grad Clip	Init Std
100B	0.1	1.0	0.006

Table 7: Training hyperparameters