

Differential Mamba

Nadav Schneider^{1,2}, Itamar Zimerman^{3,4}, Eliya Nachmani¹

¹School of Electrical and Computer Engineering, Ben-Gurion University of the Negev

²IAEC

³Tel-Aviv University

⁴IBM Research

nadavsch@post.bgu.ac.il zimerman1@mail.tau.ac.il eliyanac@bgu.ac.il

Abstract

Sequence models like Transformers and RNNs often overallocate attention to irrelevant context, leading to noisy intermediate representations. This degrades LLM capabilities by promoting hallucinations, weakening long-range and retrieval abilities, and reducing robustness. Recent work has shown that differential design can mitigate this issue in Transformers, improving their effectiveness across various applications. In this paper, we explore whether these techniques, originally developed for Transformers, can be applied to Mamba, a recent architecture based on selective state-space layers that achieves Transformer-level performance with greater efficiency. We show that a naive adaptation of differential design to Mamba is insufficient and requires careful architectural modifications. To address this, we introduce a novel differential mechanism for Mamba, empirically validated on language modeling benchmarks, demonstrating improved retrieval, long-context capabilities and superior performance over vanilla Mamba. Finally, we conduct extensive ablation studies and empirical analyses to justify our design choices and provide evidence that our approach effectively mitigates the overallocation problem in Mamba-based models. Our code is publicly available.



<https://github.com/NadavSc/Diff-Mamba>

1 Introduction

Designing enhanced architectures for deep sequence modeling is a pivotal task in the ML community, as sequence models drive breakthroughs such as ChatGPT (Achiam et al., 2023) in NLP and Stable diffusion (Rombach et al., 2022) in computer vision, which are fundamental to modern generative models. However, these models face two major challenges as: efficiency, constrained by the quadratic time complexity in sequence length of the Transformer architecture (Vaswani, 2017),

and robustness, which is hindered by inconsistency, reliability issues, and hallucinations, leading to sub-optimal performance. Our work addresses these challenges by enhancing the robustness of recent efficient architectures with sub-quadratic complexity such as Mamba (Gu and Dao, 2024), making them more reliable and robust.

To address this robustness issue, we aim to reduce the **over-allocation of attention to irrelevant context** across hidden layers, which often leads to noisy representations. Our approach is inspired by Ye et al. (2024), who mitigated this problem in transformers through differential design. This architectural modification was the core of the Diff-Transformer, a transformer variant that demonstrated improved performance, including greater robustness, enhanced retrieval and long-context capabilities compared to the original model.

Our focus is on improving the recently introduced Mamba architecture, which builds upon selective state-space layers (S6). This architecture is known for its efficiency, introducing sub-quadratic complexity in sequence length while also enabling auto-regressive decoding with a complexity that remains independent of sequence length. Beyond these efficiency advantages, recent studies have demonstrated that Mamba-based architectures can match or even surpass the SoTA performance of transformers, even at relatively large scales. For example, Falcon-Mamba 7B (Zuo et al., 2024), a pure Mamba model, matches the performance of LLaMA- 3-8B on several language tasks. Additional notable examples of Mamba’s integration in LLMs include Jamba (Lieber et al., 2024), Zamba (Glorioso et al., 2024), Hymba (Dong et al., 2024), and others (Wang et al., 2024; Waleffe et al., 2024; Ren et al., 2024). More importantly, recent Mamba-based models have demonstrated remarkable performance as reasoning models (Paliotta et al., 2025; Wang et al., 2025), underscoring their central role in the ongoing test-time scaling revolu-

tion, as exemplified by models such as OpenAI’s O1 (Jaech et al., 2024) and DeepSeek’s R1 (Guo et al., 2025).

While the over-allocation problem is a general issue not specific to any architecture, we hypothesize that Mamba-based LLMs have a stronger tendency toward over-allocation compared to transformers. This is primarily due to two factors: (i) Mamba is a softmax-free architecture, meaning it lacks the exponential scaling effect that helps suppress irrelevant attention weights, and (ii) as a state-based model, Mamba operates locally and cannot directly process distant tokens without considering all intermediate tokens, resulting in the dispersion of important tokens among irrelevant ones. This leads to our central research question:

Can differential design be leveraged to improve the robustness of Mamba models?

where we hypothesize that improving robustness by mitigating the over-allocation problem can enhance the model’s general capabilities, improve retrieval and long-context processing, and potentially reduce issues related to consistency and hallucinations.

We provide a positive answer to this question by showing that, although a naive implementation of differential mechanisms does not improve Mamba architectures, a more carefully designed mechanism does. Through a systematic evaluation on language tasks, ablation studies, and empirical analysis, we conclude that our variant is favorable compared to the vanilla Mamba.

Our main contributions are as follows: (i) We present Diff-Mamba, a modification of the Mamba architecture inspired by Diff-Transformer, which mitigates the problem of over-allocating attention scores to irrelevant context and improves the general language modeling abilities of the model. (ii) Through a series of ablation studies and empirical analysis using mechanistic interpretability tools, we justify our design choices and demonstrate that the intermediate representations obtained from our method are indeed less noisy. (iii) Finally, we show that Diff-Mamba demonstrates improved retrieval and long-context capabilities compared to Mamba. This is particularly important, as recurrent LLMs such as Mamba are primarily designed to address the inefficiencies caused by the quadratic complexity of transformers, which becomes especially critical in long-context regimes. Accordingly, Diff-Mamba achieves improvements over Mamba precisely in the domain where such architectures are

most needed (Ben-Kish et al., 2025b).

2 Background

Here we describe the scientific context and introduce the terminology for discussing our method.

2.1 Differential Transformer

Self-Attention Self-Attention is a fundamental component of Transformer architectures (Vaswani, 2017), has significantly shaped recent advances in both NLP and computer vision. This mechanism enables dynamic focus allocation by capturing pairwise token dependencies, allowing the model to determine the relative importance of each token within a sequence. Mathematically, it is defined by:

$$\text{Attn}(Q, K, V) = \alpha V, \quad \alpha = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) \quad (1)$$

In this formulation, Q , K , and V represent the queries, keys, and values, respectively, while d_k denotes the key dimension. Transformers extend this mechanism by employing H parallel attention heads, enabling the model to capture a broader spectrum of dependencies.

Differential Attention To address the problem of over-allocation of attention to irrelevant tokens, Ye et al. (2024), introduced Diff-Transformer, a mechanism that reduces attention noise through differential denoising by splitting each attention head into two, and subtracting one attention map from the other. This mechanism is defined by:

$$\text{DiffAttn}(Q_1, K_1, Q_2, K_2, V) = (\alpha_1 - \lambda \alpha_2)V \quad (2)$$

$$\alpha_i = \text{softmax} \left(\frac{Q_i K_i^T}{\sqrt{d_k}} \right) \quad (3)$$

where λ is a learnable scalar. To better improve the training dynamics, λ is re-parameterized and Group Normalization (Wu and He, 2018) is applied at the end of each head (post-subtraction).

2.2 State-Space Layers

State-space layers were first introduced in Gu et al. (2021) and were later substantially improved by the S4 model (Gu et al., 2022). Since then, they have demonstrated strong performance across a variety of domains, including NLP (Fu et al., 2022; Mehta et al., 2022), audio generation (Goel et al., 2022), image modeling (Baron et al., 2023; Nguyen et al., 2022), long-horizon video understanding (Wang et al., 2023), reinforcement learning (David

et al., 2022; Lu et al., 2024), and speech recognition (Saon et al., 2023). These models implement linear recurrent update rules derived from time-invariant state-space formulations, which can be efficiently computed in parallel using convolutions and with sub-quadratic complexity.

2.3 Mamba and Selective State-Space Layers

A Mamba block processes a signal $U \in \mathbb{R}^{L \times D}$ where D is the hidden dimension and L is the number of tokens. Its core mechanism is the S6 layer, and its forward path formalized by:

$$\begin{aligned} X &= \sigma(\text{Conv1D}(\text{Linear}(U))), \quad Z = \sigma(\text{Linear}(U)) \\ Y &= \text{S6}(X), \quad \hat{Y} = \text{Linear}(Y \otimes Z) \end{aligned} \quad (4)$$

where X is the input to the S6 layer, and to the S6 layers, $X, Z, Y, \hat{Y} \in \mathbb{R}^{L \times D}$. The function σ represents SiLU activation, and \otimes represents element-wise multiplication with the gating branch. Each Mamba block is primarily parameterized by linear and convolutional layers, along with the internal components of the S6 layer described below.

S6 The S6 layer is the most popular variant of SSMs, and it employs real, diagonal and selective SSM. Standard real and diagonal SSMs parameterized by a diagonal transition matrix $A \in \mathbb{R}^{N' \times N'}$, input and output matrices $B, C \in \mathbb{R}^{N' \times 1}$ where N' is the state size, and a timescale $\Delta \in \mathbb{R}$. Each channel of such an SSM can be viewed as a mapping from an input scalar sequence x to an output scalar sequence y via the following recurrent rule:

$$\begin{aligned} h_t &= \bar{A}h_{t-1} + \bar{B}x_t, \quad y_k = Ch_t \\ \bar{A} &= f_A(A, \Delta), \quad \bar{B} = f_B(B, \Delta) \end{aligned} \quad (5)$$

where f_A, f_B are discretization functions, and the discrete system matrices are $\bar{A} \in \mathbb{R}^{N' \times N}$ and $\bar{B} \in \mathbb{R}^{N' \times 1}$. The recurrent rule in Eq. 5 can be computed efficiently in parallel on modern hardware accelerators using work-efficient parallel scans (Smith et al., 2022) or a simple scalar convolution via FFTs (Gu et al., 2021). Note that Eq. 5 is a map from \mathbb{R}^L to \mathbb{R}^L , and to process D channels, multiple independent instances are used.

The S6 layer differs from standard SSMs by employing a selective mechanism, where the system matrices are input-dependent. As a result, the system becomes time-invariant, with the per-step system matrices determined by the entire set of channels and then applied to process each channel independently. The entire mechanism can be computed by: $S_B, S_C \in \mathbb{R}^{N' \times D}$, $A \in \mathbb{R}^{D \times N'}$ and $S_\Delta \in \mathbb{R}^{1 \times D}$ to define the time-variant matrices by:

$$\begin{aligned} B_t &= S_B X_{*t}, \quad C_t = S_C X_{*t}, \quad \Delta_t = \text{softplus}(S_\Delta X_{*t}) \\ \bar{A}_t &= \exp(\Delta_t A), \quad \bar{B}_t = \Delta_t B_t \end{aligned} \quad (6)$$

and the time-variant recurrent rule by:

$$h_t = \bar{A}_t h_{t-1} + \bar{B}_t x_t, \quad y_k = C_t h_t \quad (7)$$

We study the ‘many-to-one’ setting, where the model processes an entire input sequence to produce a single output. This regime is widely used in NLP, encompassing both auto-regressive next-token prediction and sequence classification tasks.

2.3.1 Mamba as Implicit Attention

The connection between Mamba and linear attention layers is well-established in (Ali et al., 2024; Dao and Gu, 2024; Sieber et al., 2025). Specifically, it has been shown that the time-variant recurrent update rule of S6 for a single channel (see Eq. 7) can be explicitly unrolled into the linear attention formulation $Y = AX$ when A is an implicit attention matrix defined by:

$$\begin{bmatrix} C_1 \bar{B}_1 & 0 & \cdots & 0 \\ C_2 \bar{A}_2 \bar{B}_1 & C_2 \bar{B}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ C_L \prod_{k=2}^L \bar{A}_k \bar{B}_1 & C_L \prod_{k=3}^L \bar{A}_k \bar{B}_2 & \cdots & C_L \bar{B}_L \end{bmatrix} \quad (8)$$

This perspective is further extended by Zimerman et al. (2024), who generalize the interpretation of S6 as implicit attention from Eq. 8 to encompass most components of the Mamba block, including activations, normalization layers, convolutional layers, and the gate branch, into a unified implicit attention formulation.

Data-Controlled Linear Operators The formulation of data-controlled linear operators was first introduced by Poli et al. (2023), who demonstrated that self-attention can be viewed as an expressive form of such operators. This principle guided the authors in designing the Hyena layer. Subsequently, Ali et al. (2024) showed that S6 layers could also be unified under an implicit variant of this formulation, which Zimerman et al. (2024) further extended to additional architectures, including the entire Mamba block, RWKV (Peng et al., 2023), and Griffin (De et al., 2024). This extended perspective inspired our approach, leading us to interpret differential design as a method to implicitly parameterize less noisy data-controlled linear operators. This insight motivated our decision to apply differential design at the Mamba level.

3 Method

We begin with the simplest implementation of incorporating the differential mechanism into the Mamba block. Our approach is inspired by the Differential Transformer (Ye et al., 2024), which applies subtraction at the attention-level rather than at the transformer level. This mechanism is built on top of self attention and it can be described as:

$$\forall i \in [1, 2] : Q_i = XW_i^Q, K_i = XW_i^K, V = XW^V \quad (9)$$

$$\text{DiffAttn}(x) = \left(\mathcal{S}\left(\frac{Q_1 K_1^T}{\sqrt{d}}\right) - \lambda \mathcal{S}\left(\frac{Q_2 K_2^T}{\sqrt{d}}\right) \right) V \quad (10)$$

Here, \mathcal{S} is the softmax and λ is parameterized to ensure stability and improve training dynamics:

$$\lambda = \exp(\lambda_{q1} \cdot \lambda_{k1}) - \exp(\lambda_{q2} \cdot \lambda_{k2}) - \lambda_{\text{init}} \quad (11)$$

To better adapt this technique to Mamba models, we reinterpret differential attention through the lens of a data-controlled linear operator (Poli et al., 2023), leading to the following formulation:

$$\begin{aligned} \text{DiffAttn}(x) &= AV, \quad A = A_1 - \lambda A_2 \\ A_1 &= \mathcal{S}\left(\frac{Q_1 K_1^T}{\sqrt{d}}\right), \quad A_2 = \mathcal{S}\left(\frac{Q_2 K_2^T}{\sqrt{d}}\right) \end{aligned} \quad (12)$$

here, the matrix A defines a data-controlled linear operator.

3.1 Diff S6

Eq. 12 defines a straightforward approach to implementing differential Mamba by subtracting values obtained from S6 layers instead of attention layers. This builds upon two key similarities between S6 and attention: (i) S6 in Mamba serves the same role as attention in Transformers—capturing interactions between tokens, and (ii) S6 layers have been shown to be an implicit form of causal linear attention.

Thus, incorporating the differential mechanism into the Mamba block can be achieved by:

$$\text{Diff S6}(X) = S6_1(X) - \lambda S6_2(X) \quad (13)$$

where λ is defined similarly to Eq. 11.

Similar to Eq. 12, this formulation can be rewritten in the form of a data-controlled linear operator:

$$\text{Diff S6}(x) = (A_1 - \lambda A_2)X = AX \quad (14)$$

where A_1 and A_2 are the implicit attention matrices of S6 controlled by the system matrices $\bar{A}_{ij}, \bar{B}_{ij}$ and C_{ij} for any time-step $j \in [L]$ and model index $j \in [1, 2]$ defined as follows:

$$\begin{bmatrix} C_{i1}\bar{B}_{i1} & 0 & \cdots & 0 \\ C_{i2}\bar{A}_{i2}\bar{B}_{i1} & C_{i2}\bar{B}_{i2} & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ C_{iL}\prod_{k=2}^L \bar{A}_{ik}\bar{B}_{i1} & C_{iL}\prod_{k=3}^L \bar{A}_{ik}\bar{B}_{i2} & \cdots & C_{iL}\bar{B}_{iL} \end{bmatrix} \quad (15)$$

One crucial difference between Diff S6 and Diff Attention (see Eqs. 13 and 10) is that Diff Attention subtracts elements on the same scale, as softmax produces values in the range [0,1]. In contrast, S6 produces unnormalized and unbounded outputs. To address this discrepancy, we introduce an additional normalization step denoted by \mathbb{N} :

$$\mathbb{N}\text{-Diff S6}(X) = \mathbb{N}(S6_1(X) - \lambda S6_2(X)) \quad (16)$$

For simplicity, we define λ as:

$$\lambda = \text{Sigmoid}(\sum \bar{\lambda}) + \lambda_{\text{init}} \quad (17)$$

where $\bar{\lambda} \in \mathbb{R}^D$ is a learnable parameter used to parameterize λ as a positive and more stable weight.

3.2 Diff-Mamba

However, as detailed in the results section, Diff S6 does not perform well and, in practice, falls short of standard Mamba layers. We suspect this arises from S6 being too simple and not functioning as a general-purpose, expressive mixing alternative to attention layers. Consequently, it fails to leverage the full potential of differential techniques. To address this limitation, we draw inspiration from Zimerman et al. (2024), who demonstrated that the entire Mamba block can function as an alternative mixing mechanism to attention by formulating it as a data-controlled linear operator. In particular, the authors show that Mamba can be reformulated as implicit attention by:

$$\text{Mamba}(X) = AX \quad (18)$$

$$A = \text{SILU}(\text{Linear}(x)) \hat{\alpha} \text{ diag}(\text{Sig}(\text{Conv}(x)))M$$

where M is a matrix representing the convolution layer, and $\hat{\alpha}$ is the linear operator corresponding to S6, as formalized by Ali et al. (2024). This formulation characterizes Mamba as a data-control linear operator with richer and more expressive implicit attention matrices. Building on this insight, we introduce Diff-Mamba, a mechanism that extends the differential approach to the full Mamba block:

$$\text{Diff-Mamba}(X) = \text{Mamba}_1(X) - \lambda \text{Mamba}_2(X) \quad (19)$$

Similar to Eqs. 12 and 14, this can be rewritten as a data-controlled linear operator, defined by:

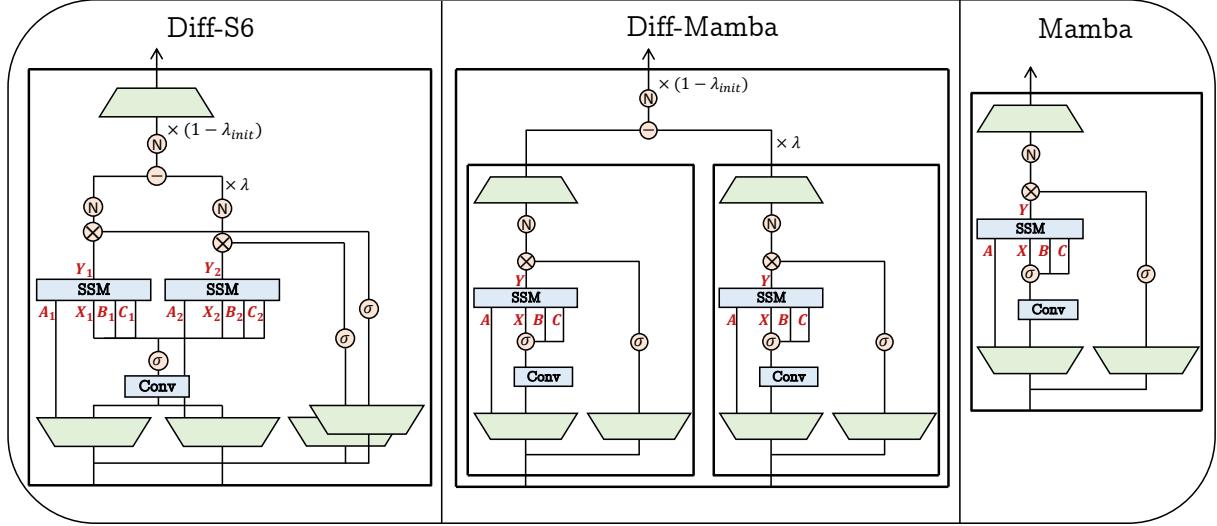


Figure 1: Comparative illustration of our variants Diff-Mamba and Diff-S6 versus the original Mamba architecture, where \otimes is elementwise multiplication, σ is the SILU activation, Linear and Conv1D are standard linear projection and 1-dimensional convolution layers, and N stands for normalizations.

$$\text{Diff-Mamba}(X) = AX, \quad A = A_1 - \lambda A_2 \quad (20)$$

A key distinction between Diff Attention and Diff-Mamba is that the latter applies subtraction across a broader set of components, as illustrated in Figure 1.

Similar to the normalized Diff S6 variant in Eq. 16, we add a normalization term:

$$N\text{-Diff-Mamba}(x) = N(Mamba_1(x) - \lambda Mamba_2(x)) \quad (21)$$

Resulting in the normalized Diff-Mamba mechanism, which is our primary contribution. Finally, following Ye et al. (2024), we multiply the output of all variants by $1 - \lambda_{\text{init}}$.

4 Experiments

In this section, we empirically evaluate the effectiveness of the Diff-Mamba architecture. We begin in Section 4.1 by demonstrating that Diff-Mamba outperforms the original Mamba architecture in small-scale language modeling tasks across multiple datasets. In Section 4.2, we justify our key design decisions through a comprehensive series of ablation studies. Subsequently, in Section 4.3, we use carefully designed synthetic tasks that are predictive of model behavior in large scale settings. Then, in Section 4.4 we train the models at medium scale and test long-context capabilities. Afterwards, in Section 4.5 we showcase the superior retrieval performance of Diff-Mamba relative to Mamba. Finally, in Section 4.6, we utilize tools from the domain of mechanistic interpretability, such as Tuned-lens (Belrose et al., 2023), to examine the internal

Model	Dataset	# Layers	# Params	PPL (↓)
Mamba	Wikitext-103	6	167M	22.357
Diff-Mamba	Wikitext-103	6	169M	22.282
Mamba	Text8	12	255M	20.413
Diff-Mamba	Text8	12	259M	20.012
Mamba	Text8	6	127M	2.416
Diff-Mamba	Text8	6	129M	2.396
Mamba	Text8	12	255M	2.525
Diff-Mamba	Text8	12	259M	2.479
Mamba	Enwik8	6	127M	2.321
Diff-Mamba	Enwik8	6	129M	2.314
Mamba	Enwik8	12	255M	2.422
Diff-Mamba	Enwik8	12	259M	2.381

Table 1: Final performance of Mamba and Diff-Mamba across model sizes. All models were trained for 40 epochs on each dataset. Trends shown in Figure 2.

representations of Diff-Mamba in comparison to Mamba, empirically validating that our differential approach effectively reduces noise in intermediate representations. A full description of our experimental setup, as well as an efficiency analysis, is provided in Appendix A and Appendix B.

4.1 Language Modeling

To evaluate the performance of Diff-Mamba relative to Mamba on general NLP tasks, we train both models from scratch using comparable model sizes and an identical training setup, including the same codebase (Gu et al., 2022), datasets, and hyperparameters. We focus on three widely used benchmarks: WikiText-103, Text8, and Enwik8,

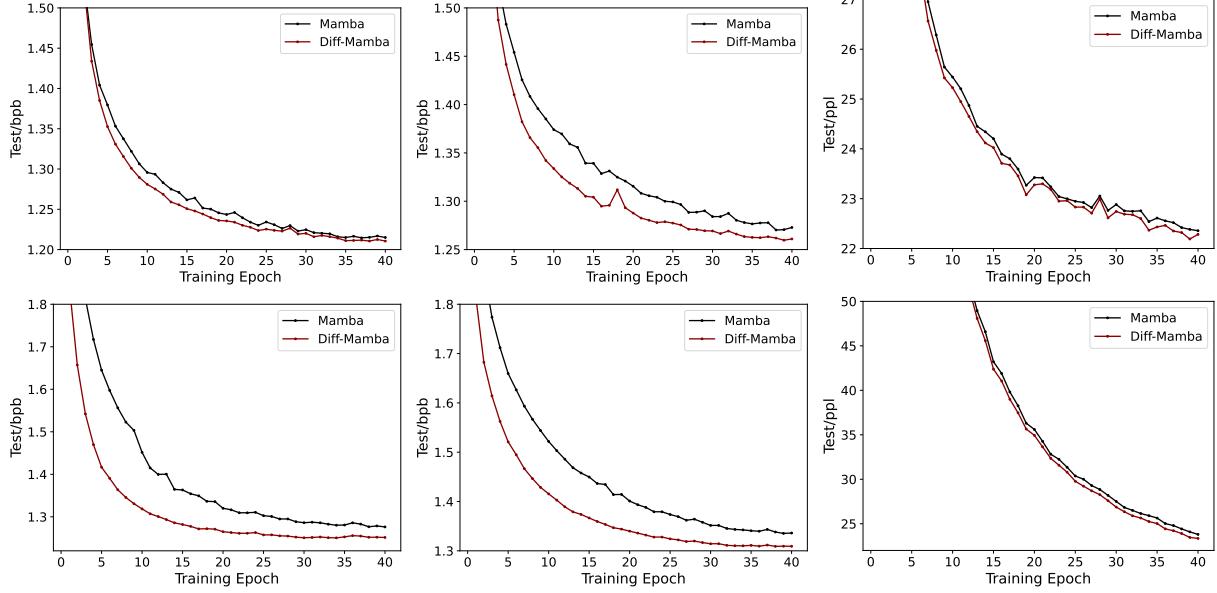


Figure 2: Comparison of test curves through the training for Mamba and Diff-Mamba. The top row shows results for 6-layer models, and the bottom row for 12-layer models. Columns correspond to datasets: Enwik8 (left), Text-8 (center), and WikiText-103 (right).

and experiment with models of varying depth. The final results are reported in Table 1. To provide a more comprehensive view of the optimization process, we include test curves through the training in Figure 2.

As shown in Table 1, Diff-Mamba outperforms Mamba across all evaluated environments, achieving consistently lower loss. In particular, for the 12 layer model, Diff-Mamba improves over Mamba by 0.4 perplexity on WikiText-103, 0.046 bits per byte (bpb) on Text8, and 0.041 bpb on Enwik8. For 6 layer model, Diff-Mamba improves over Mamba by 0.075 on WikiText-103, 0.02 on Text8, and 0.007 on Enwik8. Interestingly, we observe that as the number of layers increases, the differential design in Mamba becomes increasingly effective. A possible explanation for this is that, in the lower layers, the dependencies captured by the implicit attention matrices are shorter and simpler. In these cases, Mamba can manage overallocation effectively without the need for a differential mechanism. However, in the upper layers, the dependencies become more complex, spanning longer ranges (Ben-Kish et al., 2025a) and exhibiting more diverse patterns. This amplifies the impact of overallocation, thereby making the benefits of the differential design more pronounced. Furthermore, the training curves in Figure 2 provide insight into the optimization properties of Diff-Mamba, showing that it consistently outperforms Mamba and achieves faster conver-

Model	w.o Nrm	w. Nrm	# Params
Mamba	2.577	–	127M
Diff-S6	2.520	2.512	128M
Diff-S6 + re. $\bar{\lambda}$	2.529	2.517	128M
Diff-Mamba	2.508	2.493	128M
Diff-Mamba + re. $\bar{\lambda}$	2.517	2.503	128M

Table 2: **Ablation study** comparing (i) the scope at which the differential mechanism is applied (Diff-S6 vs. Diff-Mamba), (ii) the effect of including normalization ("w. Nrm") versus excluding it ("w.o Nrm"), and (iii) the importance of reparameterization for $\bar{\lambda}$ ("re. $\bar{\lambda}$ "). All models were trained on full Text8 with an identical parameter count. Reported values are test perplexity (PPL) on epoch 10. Lower is Better.

gence. We hypothesize that this phenomenon arises from the fact that the differential design reduces the amount of noise, which appears to be critical for improving convergence (Johnson and Zhang, 2013; Zhang et al., 2019).

4.2 Ablations Analysis

To validate our design decisions regarding (i) applying the differential operation at the S6 versus Mamba layer, (ii) incorporating an additional normalization sub-layer before subtraction, and (iii) reparameterization $\bar{\lambda} \in \mathbb{R}^D$ to a scalar, we conducted dedicated ablation experiments on the Text8 benchmark. Table 2 summarizes the results. All models share an identical parameter count and were trained with same hyper-parameters that optimized

for the baseline model. It can be seen that all three design choices are justified, specifically Diff-S6 outperforms Mamba without normalization, while Diff-Mamba outperforms Diff-S6. Incorporating additional normalization leads to improved results, yielding gains of 0.015 and 0.008 perplexity for Diff-Mamba and Diff-S6, respectively. Finally, $\bar{\lambda}$ reparameterization doesn't contribute to better performance demonstrated both in Diff-S6 and Diff-Mamba models.

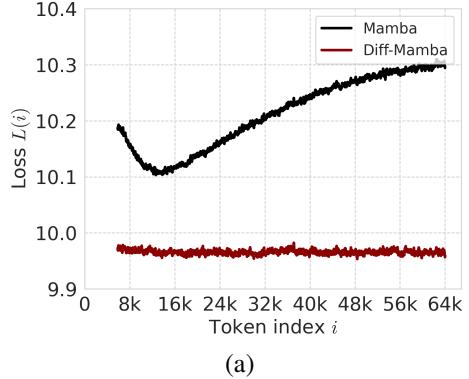
4.3 Unit Tests as Predictors of Scaling

Training LLMs at scale is prohibitively resource-intensive. Nevertheless, prior work has established that carefully designed small-scale capability evaluations (Gupta et al., 2022) are predictive of model behavior in large-scale (Poli et al., 2024). To this end, we follow the MAD pipeline (Poli et al., 2024) and adopt a rigorous suite of synthetic token manipulation tasks, that serve as capability unit tests. These tasks allow us to isolate and assess core mechanisms that underpin scalable performance. We investigate Mamba and Diff-Mamba. Diff-Mamba model consistently demonstrates superior results (Figure 5), providing strong evidence that the architecture is well-positioned to retain and extend these capabilities when trained at larger scales. Both models achieve the highest accuracy on in-context-recall (ICR), noisy ICR, and selective-copying tasks, while Diff-Mamba outperforms Mamba on the rest of the tasks, specifically up to 10.6% improvement in compression, 80% in fuzzy ICR, and 0.2% in memorization. Diff-Mamba surpasses Mamba by 3.7% in total.

4.4 Early Results at Medium Scale

To thoroughly evaluate Diff-Mamba, we trained both Mamba and Diff-Mamba from scratch, each with 370M parameters, on a 50B token subset of The Pile dataset (Gao et al., 2020). After an ablation study (Figure 7 in the Appendix), the most effective Diff-Mamba variant for scalable performance was found to be a repeated alternation of Mamba and Diff-Mamba layers throughout the network. This hybrid variant performs better than both models with fully Mamba or Diff-Mamba layers.

Next, we evaluated the models via zero-shot tests on the LongCrawl64 dataset (Buckman, 2024), a long-sequence subset of RedPajama-v2 (Weber et al., 2024) designed particularly for research on long-context. In Figure 3 (a), per-token loss was calculated over the dataset following (Lin et al.,



(a)

Model	The Pile (↓)	PG19 (↓)
Mamba	11.212	28.064
Diff-Mamba	11.081	26.619

(b)

Figure 3: Diff-Mamba excels in both tests. (a) The x-axis is the token index, and the y-axis is the corresponding per-token loss. (b) PPL results on test set of The Pile and PG19.

2025). In addition, in Figure 3 (b), we calculated PPL on the test subset of The Pile and PG19. Diff-Mamba demonstrates impressive long context capabilities, maintaining per-token loss of around 9.98 across different context lengths, while Mamba increases significantly as the context grows. Furthermore, Diff-Mamba outperforms Mamba by PPL scores of 0.131, and 1.445 on The Pile and PG19 test sets respectively. We consider the positive results at medium scale promising, suggesting that alternating between Mamba and Diff-Mamba layers yields a more effective, scalable, and robust architecture.

Tasks	qa1	1.05	1.12	1.03	1.24	1.64	1.75	1.77	1.64
qa2	1.01	1.18	1.24	1.82	2.07	1.86	2.25	2.67	
qa3	1.03	1.12	1.12	1.56	2.27	2.40	3.21	3.53	
qa4	1.21	1.03	1.11	2.09	3.43	2.39	2.36	3.47	
qa5	1.40	1.52	1.55	1.92	2.85	2.77	3.31	3.05	
	0k	1k	2k	4k	8k	16k	32k	64k	Context Length

Figure 4: **Retrieval Abilities:** Comparison of Diff-Mamba and Mamba models across five retrieval tasks from BABILong. X-axis represents the context length, and y-axis corresponds to the task index. Each cell displays the ratio in which one model outperforms the other. Green cells indicate wins by Diff-Mamba, while red cells indicate wins by Mamba.

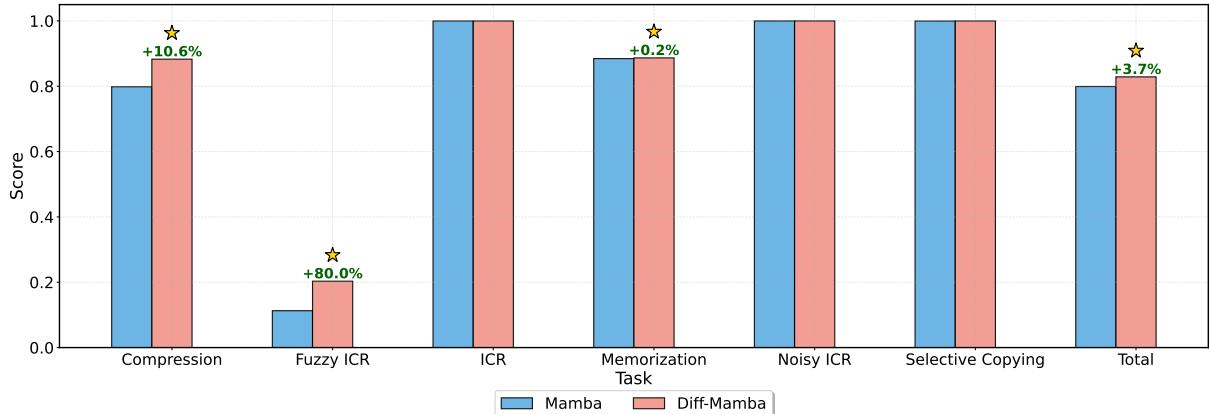


Figure 5: Performance comparison on synthetic token manipulation tasks. We evaluate Mamba and Diff-Mamba architectures across six synthetic capability benchmarks. Diff-Mamba consistently outperforms Mamba, with notable improvements in Fuzzy ICR of 80.0% and Compression with 10.6%. Stars indicate the best model per task.

4.5 Retrieval

The Diff-Transformer exhibits significantly improved retrieval capabilities compared to the original Transformer. Consequently, we conduct experiments across five retrieval tasks from the BABI-Long benchmark (Kuratov et al., 2024) to evaluate whether these enhanced abilities transfer to Diff-Mamba. BABILong comprises a diverse array of reasoning tasks, including fact chaining, simple induction, and deduction, where relevant facts are embedded within lengthy natural language passages of varying context lengths. This setup provides a rigorous benchmark for evaluating models’ ability to retrieve and reason over extended contexts. To ensure a fair comparison, both Diff-Mamba and Mamba (from Section 4.4) were fine-tuned on BABILong tasks with up to 1k tokens to facilitate a more targeted comparison, ensuring that their learning processes were aligned with the same objectives.

Results (Figure 4) are reported on the test set. Each cell displays the ratio in which one model outperforms the other. Green indicates wins by Diff-Mamba, while red indicates wins by Mamba. As evident from the results, Diff-Mamba outperforms Mamba, exhibiting a slower performance degradation and a larger ratio score as context length increases. Diff-Mamba wins consistently, achieving a ratio of up to 3.95. Our findings demonstrate that Diff-Mamba achieves superior retrieval performance, particularly in long-context scenarios.

4.6 Noise in Intermediate Representations

Our empirical analysis in previous sections, as well as the results of Diff-Transformer suggest that the

differential design can mitigate the overallocation problem and improve general performance. To further investigate the underlying causes of this phenomenon, we analyze the model’s internal representations using tools from the field of mechanistic interpretability. In particular, we leverage Tuned-lens (Belrose et al., 2023) - a method designed to examine intermediate representations by training an affine probe to map activations at each layer to the model’s final prediction, thereby enabling layer-wise interpretability and insight into the model’s internal computation. Building on this tool, we measure the signal-to-noise ratio in the hidden representations of Diff-Mamba compared to standard Mamba models. The Tuned-lens tool projects logits from each layer into predictions for the next token on the retrieval task. By measuring the predicted probability of the needle token, we can estimate the signal-to-noise ratio at different layers. Notably, as can be seen in Figure 6, across the majority of layers Diff-Mamba exhibits a higher signal-to-noise ratio compared to Mamba. This difference is especially pronounced in the early to mid layers, where the predicted probability of the needle token in Diff-Mamba is substantially higher. This analysis empirically demonstrates that Diff-Mamba produces less noisy representations, directly aligning with the key principles underlying the differential mechanism, which is designed to mitigate the overallocation problem.

5 Discussion: Why Differential Design

A key question that arises from the empirical findings presented in this paper and in Diff-Transformer is why differential design is so effec-

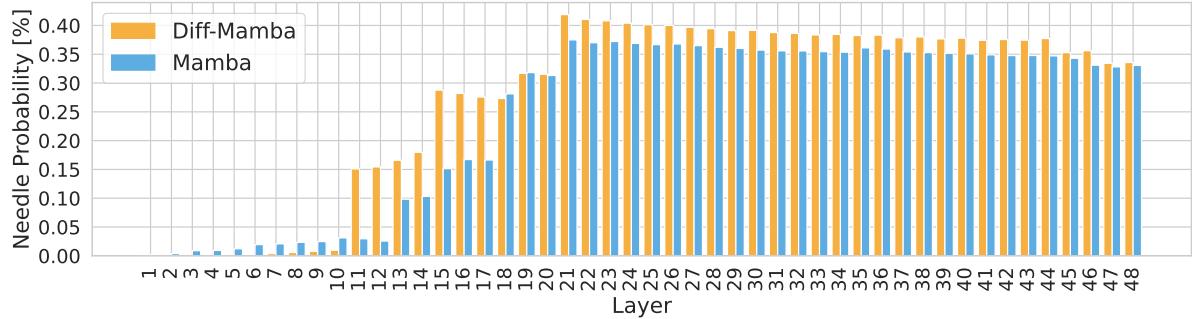


Figure 6: **Measuring Signal-to-Noise Ratio:** The y-axis represents the probability of predicting the desired needle token, where lower values indicate higher noise. The x-axis denotes various layers within the model where intermediate noise is measured. Results show the average needle probabilities in each layer on 1k examples in BABILong questions of 1k-2k tokens.

tive and what the underlying causes of its power are. Our work sheds light on this question by first showing that the problem of over-allocating attention to irrelevant context is not a phenomenon unique to transformers, but rather a general challenge in architectural research. Moreover, while previous work has primarily motivated the differential design through analogies to noise-canceling headphones and differential amplifiers, and supported it with strong empirical performance, we take this a step further. In Figure 6, we present a quantitative analysis using tools from the field of model understanding, showing that the intermediate representations in Diff-Mamba exhibit a higher signal-to-noise ratio compared to their non-differential Mamba counterparts, providing empirical support for the motivation previously proposed in the literature. Yet, the underlying reasons behind the empirical success of differential design remain largely unexplored, and further investigation from both theoretical and empirical perspectives is required to advance progress in this important direction.

6 Conclusions

In this paper, we introduced Diff-Mamba, a variant of the Mamba architecture that leverages differential design principles to mitigate the problem of attention score over-allocation to irrelevant tokens, thereby enhancing overall performance, with a particular emphasis on long-context, retrieval and robustness capabilities. Diff-Mamba advances architectural research on Mamba variants by introducing an inductive bias that enhances the signal-to-noise ratio, a property that may be critical at scale, as hallucinations remain one of the most significant challenges in LLMs.

7 Limitations

Although the experimental results show promising improvements over language modeling and retrieval tasks, we did not provide a rigorous theoretical framework explaining precisely why differential designs improve Transformer or Mamba-based LLMs. Developing such a theoretical justification is left as future work. Additionally, our results are limited to small-to-medium scale experiments due to constraints imposed by our academic budget. Finally, it remains an open question whether differential design principles can be effective in other domains, beyond NLP tasks, for instance, to domains such as computer vision, graph modeling, or time-series analysis.

Ethics Statement

Our work aims to improve the robustness of Mamba models by increasing the signal-to-noise ratio through differential design. In the broader context, this contribution supports the development of more efficient, reliable, and trustworthy LLMs. Additionally, our findings highlight when and how differential design is effective and offer practical guidance for incorporating inductive biases that lead to less noisy architectures. Thus, we conclude that our approach contributes to the principled design of safer and more reliable LLMs.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Ameen Ali, Itamar Zimerman, and Lior Wolf. 2024. The hidden attention of mamba models. *arXiv preprint arXiv:2403.01590*.

Ethan Baron, Itamar Zimerman, and Lior Wolf. 2023. 2-d ssm: A general spatial layer for visual transformers. *arXiv preprint arXiv:2306.06635*.

Nora Belrose, Zach Furman, Logan Smith, Danny Halawi, Igor Ostrovsky, Lev McKinney, Stella Biderman, and Jacob Steinhardt. 2023. Eliciting latent predictions from transformers with the tuned lens. *arXiv preprint arXiv:2303.08112*.

Assaf Ben-Kish, Itamar Zimerman, Shady Abu-Hussein, Nadav Cohen, Amir Globerson, Lior Wolf, and Raja Giryes. 2025a. **Decimamba: Exploring the length extrapolation potential of mamba**. In *The Thirteenth International Conference on Learning Representations*.

Assaf Ben-Kish, Itamar Zimerman, M Jehanzeb Mirza, James Glass, Leonid Karlinsky, and Raja Giryes. 2025b. Overflow prevention enhances long-context recurrent llms. *arXiv preprint arXiv:2505.07793*.

Jacob Buckman. 2024. Longcrawl64: A Long-Context Natural-Language Dataset.

Tri Dao and Albert Gu. 2024. Transformers are ssms: Generalized models and efficient algorithms through structured state space duality. *arXiv preprint arXiv:2405.21060*.

Shmuel Bar David, Itamar Zimerman, Eliya Nachmani, and Lior Wolf. 2022. Decision s4: Efficient sequence-based rl via state spaces layers. In *The Eleventh International Conference on Learning Representations*.

Soham De, Samuel L Smith, Anushan Fernando, Aleksandar Botev, George Cristian-Muraru, Albert Gu, Ruba Haroun, Leonard Berrada, Yutian Chen, Srivatsan Srinivasan, and 1 others. 2024. Griffin: Mixing gated linear recurrences with local attention for efficient language models. *arXiv preprint arXiv:2402.19427*.

Xin Dong, Yonggan Fu, Shizhe Diao, Wonmin Byeon, Zijia Chen, Ameya Sunil Mahabaleshwarkar, Shih-Yang Liu, Matthijs Van Keirsbilck, Min-Hung Chen, Yoshi Suhara, and 1 others. 2024. Hymba: A hybrid-head architecture for small language models. *arXiv preprint arXiv:2411.13676*.

Daniel Y Fu, Tri Dao, Khaled K Saab, Armin W Thomas, Atri Rudra, and Christopher Ré. 2022. Hungry hungry hippos: Towards language modeling with state space models. *arXiv preprint arXiv:2212.14052*.

Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, and 1 others. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.

Paolo Glorioso, Quentin Anthony, Yury Tokpanov, James Whittington, Jonathan Pilault, Adam Ibrahim, and Beren Millidge. 2024. Zamba: A compact 7b ssm hybrid model. *arXiv preprint arXiv:2405.16712*.

Karan Goel, Albert Gu, Chris Donahue, and Christopher Ré. 2022. It's raw! audio generation with state-space models. In *International Conference on Machine Learning*, pages 7616–7633. PMLR.

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 *The International Conference on Learning Representations (ICLR)*.

Albert Gu, Isys Johnson, Karan Goel, Khaled Saab, Tri Dao, Atri Rudra, and Christopher Ré. 2021. Combining recurrent, convolutional, and continuous-time models with linear state space layers. *Advances in neural information processing systems*, 34:572–585.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shiron Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

Ankit Gupta, Harsh Mehta, and Jonathan Berant. 2022. Simplifying and understanding state space models with diagonal linear rnns. *arXiv preprint arXiv:2212.00768*.

Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, and 1 others. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.

Rie Johnson and Tong Zhang. 2013. **Accelerating stochastic gradient descent using predictive variance reduction**. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.

Yuri Kuratov, Aydar Bulatov, Petr Anokhin, Ivan Rodkin, Dmitry Sorokin, Artyom Sorokin, and Mikhail Burtsev. 2024. **Babilong: Testing the limits of llms with long context reasoning-in-a-haystack**. In *Advances in Neural Information Processing Systems*, volume 37, pages 106519–106554. Curran Associates, Inc.

Opher Lieber, Barak Lenz, Hofit Bata, Gal Cohen, Jhonathan Osin, Itay Dalmedigos, Erez Safahi,

Shaked Meirom, Yonatan Belinkov, Shai Shalev-Shwartz, and 1 others. 2024. Jamba: A hybrid transformer-mamba language model. *arXiv preprint arXiv:2403.19887*.

Zhixuan Lin, Evgenii Nikishin, Xu He, and Aaron Courville. 2025. *Forgetting transformer: Softmax attention with a forget gate*. In *The Thirteenth International Conference on Learning Representations*.

Chris Lu, Yannick Schroecker, Albert Gu, Emilio Parisotto, Jakob Foerster, Satinder Singh, and Feryal Behbahani. 2024. Structured state space models for in-context reinforcement learning. *Advances in Neural Information Processing Systems*, 36.

Harsh Mehta, Ankit Gupta, Ashok Cutkosky, and Behnam Neyshabur. 2022. Long range language modeling via gated state spaces. *arXiv preprint arXiv:2206.13947*.

Eric Nguyen, Karan Goel, Albert Gu, Gordon Downs, Preety Shah, Tri Dao, Stephen Baccus, and Christopher Ré. 2022. S4nd: Modeling images and videos as multidimensional signals with state spaces. *Advances in neural information processing systems*, 35:2846–2861.

Daniele Paliotta, Junxiong Wang, Matteo Pagliardini, Kevin Y Li, Aviv Bick, J Zico Kolter, Albert Gu, François Fleuret, and Tri Dao. 2025. Thinking slow, fast: Scaling inference compute with distilled reasoners. *arXiv preprint arXiv:2502.20339*.

Bo Peng, Eric Alcaide, Quentin Anthony, Alon Albalak, Samuel Arcadinho, Huanqi Cao, Xin Cheng, Michael Chung, Matteo Grella, Kranthi Kiran GV, and 1 others. 2023. Rwkv: Reinventing rnns for the transformer era. *arXiv preprint arXiv:2305.13048*.

Michael Poli, Stefano Massaroli, Eric Nguyen, Daniel Y Fu, Tri Dao, Stephen Baccus, Yoshua Bengio, Stefano Ermon, and Christopher Ré. 2023. Hyena hierarchy: Towards larger convolutional language models. *arXiv preprint arXiv:2302.10866*.

Michael Poli, Armin W Thomas, Eric Nguyen, Praagaash Ponnusamy, Björn Deiseroth, Kristian Kersting, Taiji Suzuki, Brian Hie, Stefano Ermon, Christopher Ré, and 1 others. 2024. Mechanistic design and scaling of hybrid architectures. *arXiv preprint arXiv:2403.17844*.

Liliang Ren, Yang Liu, Yadong Lu, Yelong Shen, Chen Liang, and Weizhu Chen. 2024. Samba: Simple hybrid state space models for efficient unlimited context language modeling. *arXiv preprint arXiv:2406.07522*.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.

George Saon, Ankit Gupta, and Xiaodong Cui. 2023. Diagonal state space augmented transformers for speech recognition. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.

Jerome Sieber, Carmen Amo Alonso, Alexandre Dider, Melanie Zeilinger, and Antonio Orvieto. 2025. Understanding the differences in foundation models: Attention, state space models, and recurrent neural networks. *Advances in Neural Information Processing Systems*, 37:134534–134566.

Jimmy TH Smith, Andrew Warrington, and Scott W Linderman. 2022. Simplified state space layers for sequence modeling. *arXiv preprint arXiv:2208.04933*.

A Vaswani. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.

Roger Waleffe, Wonmin Byeon, Duncan Riach, Brandon Norick, Vijay Korthikanti, Tri Dao, Albert Gu, Ali Hatamizadeh, Sudhakar Singh, Deepak Narayanan, and 1 others. 2024. An empirical study of mamba-based language models. *arXiv preprint arXiv:2406.07887*.

Jue Wang, Wentao Zhu, Pichao Wang, Xiang Yu, Linda Liu, Mohamed Omar, and Raffay Hamid. 2023. Selective structured state-spaces for long-form video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6387–6397.

Junxiong Wang, Tushaar Gangavarapu, Jing Nathan Yan, and Alexander M Rush. 2024. Mambabyte: Token-free selective state space model. *arXiv preprint arXiv:2401.13660*.

Junxiong Wang, Wen-Ding Li, Daniele Paliotta, Daniel Ritter, Alexander M Rush, and Tri Dao. 2025. M1: Towards scalable test-time compute with mamba reasoning models. *arXiv preprint arXiv:2504.10449*.

Maurice Weber, Daniel Y. Fu, Quentin Anthony, Yonatan Oren, Shane Adams, Anton Alexandrov, Xiaozhong Lyu, Huu Nguyen, Xiaozhe Yao, Virginia Adams, Ben Athiwaratkun, Rahul Chalamala, Kezhen Chen, Max Ryabinin, Tri Dao, Percy Liang, Christopher Ré, Irina Rish, and Ce Zhang. 2024. Red-pajama: an open dataset for training large language models. *NeurIPS Datasets and Benchmarks Track*.

Yuxin Wu and Kaiming He. 2018. *Group normalization. Preprint*, arXiv:1803.08494.

Tianzhu Ye, Li Dong, Yuqing Xia, Yutao Sun, Yi Zhu, Gao Huang, and Furu Wei. 2024. Differential transformer. *arXiv preprint arXiv:2410.05258*.

Biao Zhang and Rico Sennrich. 2019. Root mean square layer normalization. *Advances in neural information processing systems*, 32.

Michael Zhang, James Lucas, Jimmy Ba, and Geoffrey E Hinton. 2019. Lookahead optimizer: k steps forward, 1 step back. *Advances in neural information processing systems*, 32.

Itamar Zimerman, Ameen Ali, and Lior Wolf. 2024. A unified implicit attention formulation for gated-linear recurrent sequence models. *arXiv e-prints*, pages arXiv-2405.

Jingwei Zuo, Maksim Velikanov, Dhia Eddine Rhaiem, Ilyas Chahed, Younes Belkada, Guillaume Kunsch, and Hakim Hacid. 2024. Falcon mamba: The first competitive attention-free 7b language model. *arXiv preprint arXiv:2410.05355*.

A Experimental Setup

All experiments were conducted with Mamba-2 on an L40s GPU using PyTorch, on publicly available datasets.

A.1 Language Modeling

We train all models with a maximum sequence length of 512 for 40 epochs using 3 seeds (0, 42, 77) with settings as in Table 3. Diff-Mamba uses 1024 channels and 64 heads, while Mamba uses 1024 channels and 128 heads. In Diff-Mamba, we reduce the number of parameters in the linear projections of each Mamba block by half to achieve a similar number of parameters to Mamba. Following (Ye et al., 2024), the parameter λ_{init} in Eq. 11 is defined as $\lambda_{\text{init}} = 0.8 - 0.6 \cdot \exp(-0.3(i_{\text{layer}} - 1))$. In addition, the normalization we use in Eq. 21 is RMS-Norm (Zhang and Sennrich, 2019).

Model	Dataset	# Layers	Dropout	Batch	lr
Mamba	Wikitext-103	6	0.25	100	5e-4
Diff-Mamba	Wikitext-103	6	0.25	100	5e-4
Mamba	Wikitext-103	12	0.5	50	5e-4
Diff-Mamba	Wikitext-103	12	0.5	50	5e-4
Mamba	Text8	6	0.4	170	5e-4
Diff-Mamba	Text8	6	0.4	170	5e-4
Mamba	Text8	12	0.5	50	5e-5
Diff-Mamba	Text8	12	0.5	50	5e-5
Mamba	Enwik8	6	0.4	170	5e-4
Diff-Mamba	Enwik8	6	0.4	170	5e-4
Mamba	Enwik8	12	0.5	80	1e-4
Diff-Mamba	Enwik8	12	0.5	80	1e-4

Table 3: Configuration of Mamba and Diff-Mamba across model sizes. All models were trained for 40 epochs on each dataset. Test trends are presented in Figure 2.

A.2 Ablations Analysis

To validate our architecture design, an ablation study has been done. To match the number of parameters, each architecture has a different inner parameter division. Additional configuration and models architecture details are presented in Table 4 and Table 5.

A.3 Unit Tests as Predictors of Scaling

We follow the MAD pipeline (Poli et al., 2024) and adopt a rigorous suite of synthetic token manipulation tasks, that serve as capability unit tests. We investigate Mamba and Diff-Mamba. Each model has

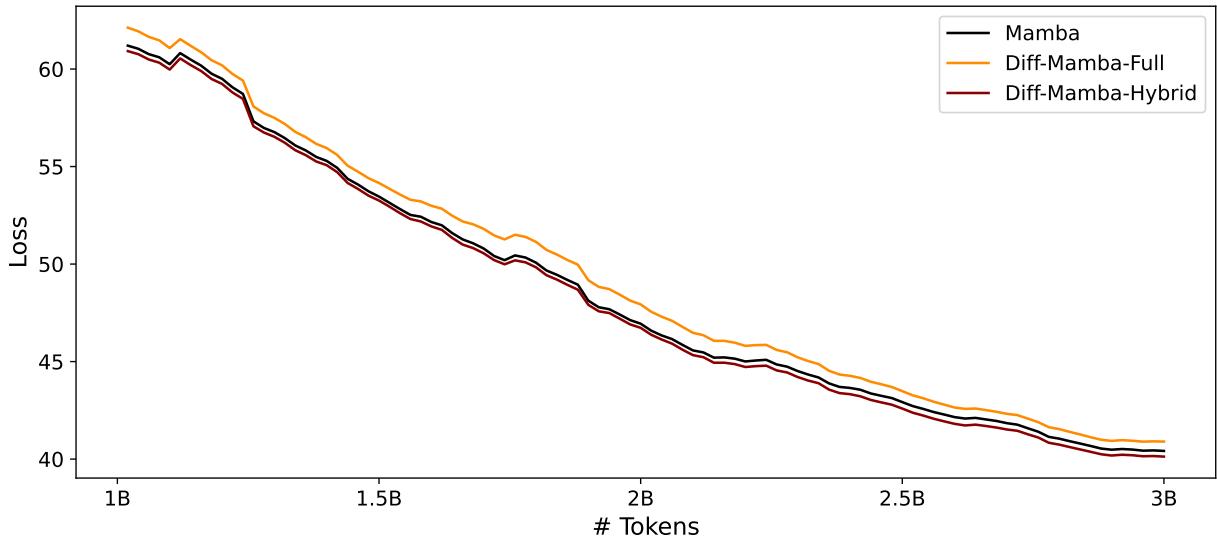


Figure 7: Loss training curve of Diff-Mamba variants compared to Mamba through pre-training on The Pile. Diff-Mamba-Hybrid indicates alternating layers of Mamba and Diff-Mamba, while Diff-Mamba-Full indicates only Diff-Mamba layers model. The trends were smoothed for display adjustment.

Params	Values
Layers	6
lr	5e-4
Warmup Steps	1000
Steps	15,000
Dropout	0.4
Batch Size	170

Table 4: Hyperparameters used for the ablation trainings.

Model	# Channels	# Heads	Expand
Mamba	1024	128	2
Diff-S6	864	108	2
Diff-S6 + re. λ	864	108	2
Diff-Mamba	1024	64	1
Diff-Mamba + re. λ	1024	64	1

Table 5: Models architecture details used for the ablation study.

four layers, consisting of the core blocks. Hyperparameters and a set of tests were applied similarly to the MAD pipeline, including different numbers of examples for training, learning rates, sequence lengths, and vocabulary size upon tasks such as compression, ICR, and selective-copying.

A.4 Early Results at Medium Scale

We pre-trained both the Mamba and Diff-Mamba architectures and conducted an ablation study on the Diff-Mamba model to evaluate the impact of architectural variations. The results indicate that interleaving Mamba and Diff-Mamba layers through-

out the network yeilds improved loss performance during pre-training, compared to using fully Diff-Mamba architecture (Figure 7). The pre-training was performed on 8 GPUs over 50 billion tokens using the GPTNeoX tokenizer. The Mamba model contains 368M parameters, Diff-Mamba-Hybrid has 375M parameters, and Diff-Mamba-Full configuration comprises 382M parameters. Although Diff-Mamba-Hybrid has 7 million fewer parameters than Diff-Mamba-Full, it achieves better performance. The models were trained with 48 layers, a learning rate of 1.5e-3, 10,000 warm-up steps, the AdamW optimizer with a weight decay of 0.1, a batch size of 1M tokens per step, and a maximum sequence length of 2048 tokens.

A.5 Retrieval

For retrieval experiments, we use 50B checkpoint of Diff-Mamba and Mamba trained on The Pile, as described in Section 4.4 and finetune those on BABILong tasks, followed by an evaluation on a test set. For statistically significant results, we finetuned and evaluated with 3 seeds.

A.5.1 BABILong Fine-tuning

The models were fine-tuned on 90% examples up to 1k tokens inside the dataset across all tasks, which is approximately equivalent to 17k examples. The other 10% remained for test. Following (Kuratov et al., 2024), we employ a preprocessing step that shapes each sample as follows: "<context>{input}</context> Question:{question} Answer:" and the loss

was calculated on the answer label only. Since the training tiny size, three seeds have been tested and averaged through the results. Original scores for each model are presented in Figure 8. The configuration for the training is in Table 6.

Params	Values
Layers	48
lr	3e-4
Max Length	2048
Steps	500
Batch Size	6
Warmup Steps	50
Optimizer	AdamW
Weight Decay	0.1

Table 6: Hyperparameters used for the BABILong fine-tuning

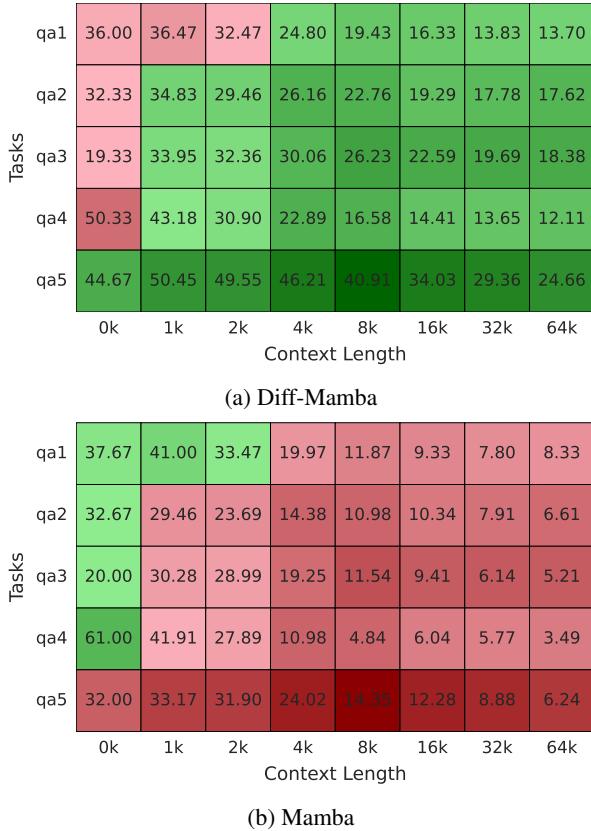


Figure 8: Comparison of original scores of Diff-Mamba and Mamba models after fine-tuning on BABILong. Values are the percentage of correct examples, while green indicates a larger score than the other model, and red indicates the opposite.

A.6 Noise in Intermediate Representations

To measure the signal-to-noise ratio, we employ Tuned-lens (Belrose et al., 2023). Specifically,

we train linear projections to learn the appropriate transformation over intermediate representations across layers and we demonstrate it using the BABILong finetuned models from Section 4.5. The probe for finetuned BABILong models have been trained on the validation set of The Pile for 3 seeds with a similar configuration to the original Tuned-lens paper. Evaluation was conducted on a test set derived from the BABILong benchmark, consisting of 1,000 examples with context lengths ranging from 1k to 2k tokens. These examples were sampled from the first five tasks, with 200 examples per task. For each example, the needle probability was extracted from every layer.

B Efficiency Benchmark

We evaluate the computational efficiency of Diff-Mamba-Hybrid (375M params) compared to Mamba (368M params) with respect to inference speed, memory footprint, and forward-pass latency. These measurements provide a comprehensive view of the trade-offs introduced by the architectural modifications.

B.1 End-to-End Inference

We calculate end-to-end inference throughput across varying batch sizes, measured on an L40S GPU with 48GB of memory (Figure 9). The benchmark was conducted using a prompt length of 2048 tokens and generating 128 new tokens with untrained models of Mamba, Diff-Mamba-Full, Diff-Mamba-Hybrid, and Transformer for a baseline. Throughput (tokens/s) is computed as $batch\ size \times 128 / inference\ time$. Both fully and hybrid Diff-Mamba exhibits lower throughput compared to the original Mamba at small batch sizes; however, this gap narrows considerably as the batch size increases, and becomes only a 12% difference with a batch size of 64 for Diff-Mamba-Hybrid.

B.2 Memory Benchmark

GPU memory usage across different batch sizes is reported in Table 7. Each batch consists of sequences of length 2048, and memory was measured after a forward pass. Diff-Mamba incurs only a modest increase in memory consumption relative to Mamba, with a difference of less than 1GB even at batch size 32 (44.30GB vs. 43.80GB). This result suggests that the additional computational layers introduced in Diff-Mamba achieve efficiency improvements without imposing prohibitive memory

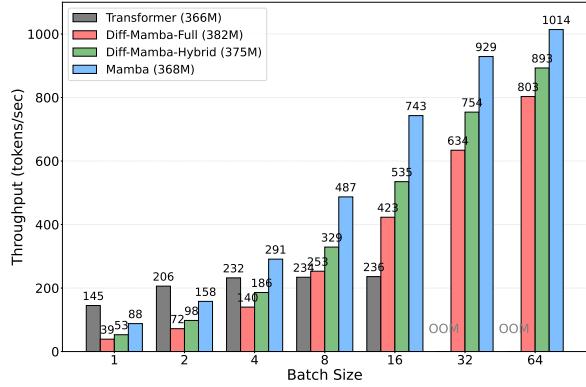


Figure 9: Inference throughput on L40s 48GB (prompt length 2048).

demands, thereby maintaining practical feasibility in resource-constrained environments.

Batch Size	Mamba (GB)	Diff-Mamba (GB)
1	3.92	3.95
2	5.16	5.25
4	7.72	7.84
8	12.84	13.03
16	23.09	23.43
32	43.80	44.30

Table 7: Memory usage comparison across different batch sizes.

B.3 Forward-Pass Latency

We further examine forward-pass latency across varying sequence lengths, summarized in Table 8. Forward-pass time serves as a direct proxy for training efficiency, since it closely aligns with per-step training duration. The results demonstrate that Diff-Mamba sustains comparable per-step efficiency to Mamba, introducing only marginal latency overhead at extended sequence lengths.

Seq Length	Mamba (s)	Diff-Mamba (s)
512	0.063	0.076
1024	0.079	0.090
2048	0.073	0.087
4096	0.147	0.172
8192	0.271	0.282

Table 8: Forward time comparison for different sequence lengths.