

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 MAMBA-3: IMPROVED SEQUENCE MODELING USING STATE SPACE PRINCIPLES

Anonymous authors

Paper under double-blind review

ABSTRACT

The recent scaling of test-time compute for LLMs has restricted the practical deployment of models to those with strong capabilities that can generate high-quality outputs in an inference-efficient manner. While current Transformer-based models are the standard, their quadratic compute and linear memory bottlenecks have spurred the development of sub-quadratic models with linear-scaling compute with constant memory requirements. However, many recent linear-style models lack certain capabilities or lag behind in quality, and even their linear-time inference is not hardware-efficient. Guided by an inference-first perspective, we introduce three core methodological improvements inspired by the state-space model viewpoint of linear models. We combine a: 1) more expressive recurrence derived from discretization , 2) complex-valued state update rule that enables richer state tracking, and 3) multi-input, multi-output formulation together, resulting in a stronger model. Together with architectural refinements, our **Mamba-3** model achieves significant gains across retrieval, state-tracking, and downstream language modeling tasks. Our new architecture sets the Pareto-frontier for performance under a fixed inference budget and outperforms strong baselines in a head-to-head comparison.

1 INTRODUCTION

Test-time compute has emerged as a key driver of progress in AI, with techniques like chain-of-thought reasoning and iterative refinement demonstrating that inference-time scaling can unlock new capabilities (Wu et al., 2025; Snell et al., 2024). This paradigm shift makes inference efficiency (Kwon et al., 2023; Li et al., 2024) paramount, as the practical impact of AI systems now depends critically on their ability to perform large-scale inference during deployment. Model architecture design plays a fundamental role in determining inference efficiency, as architectural choices directly dictate the computational and memory requirements during generation. While Transformer-based models (Vaswani et al., 2017) are the current industry standard, they are fundamentally bottlenecked by linearly increasing memory demands through the KV cache and quadratically increasing compute requirements through the self-attention mechanism. These drawbacks have motivated recent lines of work on sub-quadratic models, e.g., state-space models (SSMs), which, despite utilizing only constant memory and linear compute, have comparable or better performance than their Transformer counterparts. Models that benefit the most from this new scaling paradigm perform well on the following three axes: (i) quality, (ii) capability, and (iii) inference efficiency.

Recent model architectures have tried to strike a balance between the three, but many fall short on at least one of these three axes. In particular, Mamba-2 and Gated DeltaNet (GDN), which have gained significant traction and adoption due to their inference efficiency, made architectural design choices that enable their linear compute requirements but sacrifice quality and capabilities (Dao & Gu, 2024; Yang et al., 2025a). For example, Mamba-2 was developed to improve training speed and simplicity over Mamba-1 (Gu & Dao, 2024), opting out of more expressive parameterizations of the underlying SSM and hindering the quality of the model (Dao & Gu, 2024). Linear attention-style models (Katharopoulos et al., 2020) have also been shown to lack certain capabilities, with poor state-tracking abilities, e.g., determining parity of bit sequences, being one of the most notable (Grazzi et al., 2025; Sarrof et al., 2024). In addition, despite these sub-quadratic models being prized for theoretically efficient inference, these inference algorithms are not hardware efficient. In particular, because these algorithms were developed from a training perspective, their decoding phase has low arithmetic intensity (the ratio of FLOPs to memory traffic), resulting in large portions of hardware remaining idle.

To develop more performant models from an inference-first paradigm, we introduce three core methodological changes on top of Mamba-2, influenced by a SSM-centric viewpoint of sub-quadratic models. While many recent models fall into the linear attention framework (Dao & Gu, 2024; Yang et al., 2025a; Sun et al., 2023), we find that the classical SSM toolbox (Kalman, 1960; Gopal, 1993) leads to natural interpretations and improvements on modeling.

Trapezoidal Discretization. We discretize the underlying continuous-time dynamical system with a trapezoidal methodology. The final recurrence is a more expressive superset of Mamba-2’s recurrence and can be viewed as a convolution. We combine this new discretization with applied biases on the B, C , inspired by Yu & Erichson (2025), and find that their synergy is able to empirically replace the short causal convolution in language modeling which was previously hypothesized to be essential for recurrent models.

Complex-valued State-Space Model. By viewing the underlying SSM of Mamba-3 as complex-valued, we enable a more expressive state update than Mamba-2’s. This change in update rule, designed to be lightweight for training and inference, overcomes the lack of state-tracking ability common in many current linear models. We emphasize that our complex-valued update rule is equivalent to a data-dependent rotary embedding and can be efficiently computed (Su et al., 2023).

Multi-Input, Multi-Output SSM. To improve FLOP-efficiency during decoding, we shift from outer-product-based state update to matrix-multiplication-based state update. In view of the signal processing foundations of SSMs, such a transition exactly coincides with the generalization from a single-input single-output (SISO) sequence dynamic to a multiple-input multiple-output (MIMO) one. Here, we found that MIMO is particularly suitable for inference, as the extra expressivity allows for more compute during state update, without increasing the state size and hence compromising speed.

These three SSM-centric methodological changes are core to our **Mamba-3** mixer primitive. We also make adjustments to the overall architecture to ensure more similarity to the baseline Transformer architecture. Mamba-3 swaps the pre-output projection norm with the more common QK-normalization (Team et al., 2025; OLMo et al., 2025) and makes the short convolution, a common component found in many other sub-quadratic models (Gu & Dao, 2024; Yang et al., 2025a; von Oswald et al., 2025), optional.

We empirically validate our new model on a suite of synthetic and language-modeling tasks.

- **Better Quality.** Mamba-3 matches or outperforms Mamba-2 and other open-source architectures on standard downstream language modeling evaluations. For example, Mamba-3-1.5B’s average accuracy on all downstream tasks is better than that of its Transformer, Mamba-2, and Gated DeltaNet counterparts.
- **New Capabilities.** Mamba-3’s complexification of the SSM state enables the model to solve synthetic state-tracking tasks that Mamba-2 cannot. We empirically demonstrate that the efficient RoPE-like calculation is able to near perfectly solve arithmetic tasks, while Mamba-3 without RoPE and Mamba-2 perform not better than random guessing.
- **Stronger Inference Efficiency.** Mamba-3’s MIMO variant retains the same state size while enabling better hardware utilization compared to standard Mamba-3 and other models. Its improved performance without increased memory requirements pushes the pareto-frontier of inference efficiency.

2 PRELIMINARIES

2.1 NOTATION

Scalars are denoted by plain-text letters (e.g., x, y). Tensors, including vectors and matrices, are denoted by bold letters (e.g., \mathbf{h}, \mathbf{C}). The shape of the tensor can be inferred from the context. We denote the input sequence length as T , the model dimension as D , and the SSM state size as N . For time indices, we use subscripts (e.g., x_t for the input at time t). The Hadamard product between two tensors is denoted by \odot . For a vector of size $\mathbf{v} \in \mathbb{R}^d$, we denote $\text{Diag}(\mathbf{v}) \in \mathbb{R}^{d \times d}$ as the diagonal matrix with the vector \mathbf{v} as the diagonal, and for products of scalars across time steps, we use the notation $\alpha_{t \dots s} = \alpha_{t:s}^\times = \prod_{i=s}^t \alpha_i$.

108 2.2 SSM PRELIMINARIES
109110 State Space Models (SSMs) describe continuous-time linear dynamics via
111

112
$$\dot{\mathbf{h}}(t) = \mathbf{A}(t) \mathbf{h}(t) + \mathbf{B}(t) x(t), \quad y(t) = \mathbf{C}(t)^\top \mathbf{h}(t),$$

113

114 where $\mathbf{h}(t) \in \mathbb{R}^N$ is the hidden state, $x(t) \in \mathbb{R}$ the input, and $\mathbf{A}(t) \in \mathbb{R}^{N \times N}$, $\mathbf{B}(t), \mathbf{C}(t) \in \mathbb{R}^N$. For
115 discrete sequences with step size Δ_t , Euler’s discretization gives the recurrence
116

117
$$\mathbf{h}_t = e^{\Delta_t \mathbf{A}_t} \mathbf{h}_{t-1} + \Delta_t \mathbf{B}_t x_t, \quad y_t = \mathbf{C}_t^\top \mathbf{h}_t.$$

118

119 **Mamba-2’s parameterization.** Mamba-2 (Dao & Gu, 2024) makes the SSM *data-dependent* and
120 hardware-efficient by (i) projecting $A = \mathbf{A} \in \mathbb{R}_{<0}$, and $\mathbf{B}, \mathbf{C} \in \mathbb{R}^N$ from the current token and (ii)
121 choosing transition matrix $A = \mathbf{A}$ as a data-dependent scalar. Writing $\alpha_t := e^{\Delta_t A_t} \in (0, 1)$ and
122 $\gamma_t := \Delta_t$, the update becomes
123

124
$$\mathbf{h}_t = \alpha_t \mathbf{h}_{t-1} + \gamma_t \mathbf{B}_t x_t, \quad y_t = \mathbf{C}_t^\top \mathbf{h}_t.$$

125

126 The scalar $A_t < 0$ is an input-dependent *forget-gate (decay)* α_t , and the parameter *selectivity* Δ_t
127 jointly controls the forget-gate ($\alpha_t = \exp(\Delta_t A_t)$) and the input-gate ($\gamma_t = \Delta_t$): larger Δ_t forgets
128 faster and up-weights the current token more strongly, while smaller Δ_t retains the hidden state with
129 minimal contributions from the current token.
130

131 2.3 STRUCTURED MASKED REPRESENTATION AND STATE SPACE DUALITY

132 Dao & Gu (2024) show that a large class of SSMs admit a *matrix* form that vectorizes the time-step
133 recurrence. For instance, Mamba-2’s recurrence can be vectorized as a masked matrix multiplication,
134

135
$$\mathbf{Y} = (\mathbf{L} \odot \mathbf{C} \bar{\mathbf{B}}^\top) \mathbf{X} = \left(\begin{bmatrix} 1 & & & \\ \alpha_1 & 1 & & \\ \vdots & & \ddots & \\ \alpha_{T-1} & \cdots & \alpha_T & 1 \end{bmatrix} \odot \mathbf{C} \mathbf{B}^\top \right) \mathbf{X}, \quad (1)$$

136

137 where $\mathbf{L} \in \mathbb{R}^{T \times T}$ is the structured mask, $\mathbf{B}, \mathbf{C} \in \mathbb{R}^{T \times N}$, $\mathbf{X} \in \mathbb{R}^{T \times D}$ is the input to the SSM and
138 $\mathbf{Y} \in \mathbb{R}^{T \times D}$ is its output. Within this form, Mamba-2 can be viewed as a type of linear attention by
139 setting $\mathbf{Q} = \mathbf{C}$, $\mathbf{K} = \mathbf{B}$, $\mathbf{V} = \mathbf{X}$ and viewing \mathbf{L} as a causal, data-dependent mask. When all $\alpha = 1$,
140 the expression reduces to (causal) linear attention (Katharopoulos et al., 2020). A more detailed
141 coverage of related linear-time sequence mixers can be found at Appendix A.
142

3 MODEL DESIGN FROM A STATE-SPACE VIEWPOINT

143 We introduce Mamba-3, with three new innovations rooted in classical state-space theory: trape-
144 zoidal discretization for more expressive dynamics, complex-valued state spaces for state-tracking,
145 and multi-input multi-output (MIMO) to improve hardware utilization. These advances address the
146 quality, capability, and efficiency limitations of current sub-quadratic architectures.
147

3.1 TRAPEZOIDAL DISCRETIZATION

148 Structured SSMs are naturally defined as continuous-time dynamical systems that map input func-
149 tions, $x(t) \in \mathbb{R}$, to output functions, $y(t) \in \mathbb{R}$, for time $t > 0$. In sequence modeling, however,
150 the data is only observed at discrete time steps, which then requires applying a *discretization step*
151 to the SSM to transform its continuous-time dynamics into a discrete recurrence. The preliminary
152 step in deriving Mamba-3’s discretization is to apply the Variation of Constants formula (Propo-
153 sition 5), which decomposes the hidden state into an exponentially decay term and a state update term
154 “information” term dependent on the most recent inputs.
155156 The first step in deriving the discretized recurrence is to approximate the “state-update” integral in
157 equation 10. A straightforward choice, used in Mamba-2, is applying *Euler’s rule* (Süli & Mayers,
158 2003), which approximates the integral by holding the (right) endpoint constant throughout the
159 interval (Fig. 1). This yields Mamba-2’s recurrence,
160

161
$$\begin{aligned} \mathbf{h}_t &= e^{\Delta_t A_t} \mathbf{h}_{t-1} + (\tau_t - \tau_{t-1}) e^{(\tau_t - \tau_t) A_t} \mathbf{B}_t x_t \\ &\approx e^{\Delta_t A_t} \mathbf{h}_{t-1} + \Delta_t \mathbf{B}_t x_t. \end{aligned} \quad (2)$$




Figure 1: **Left:** The structured mask induced by the generalized trapezoid rule is a product of the decay and convolutional mask. **Right:** Euler (hold endpoint) vs trapezoidal rule (average endpoints).

However, Euler’s rule provides only a first-order approximation to the “state-update” integral: local truncation error is $O(\Delta_t^2)$, which accumulates across steps to yield a global error of $O(\Delta_t)$ over the sequence. In contrast, we adopt a *generalized trapezoidal rule*, which provides a second-order accurate approximation of the integral, offering improved accuracy over the Euler’s rule. Specifically, it approximates the integral with a *data-dependent, convex combination of both interval endpoints*. This generalization extends the classical trapezoidal rule (Süli & Mayers, 2003), which simply averages the interval endpoints, by allowing for a *data-dependent convex combination* (Fig. 1).

Proposition 1 (Generalized Trapezoidal Discretization). *Approximating the state-update integral in equation 10 by the general trapezoidal rule yields the recurrence,*

$$\mathbf{h}_t = e^{\Delta_t A_t} \mathbf{h}_{t-1} + (1 - \lambda_t) \Delta_t e^{\Delta_t A_t} \mathbf{B}_{t-1} x_{t-1} + \lambda_t \Delta_t \mathbf{B}_t x_t, \quad (3)$$

$$:= \alpha_t \mathbf{h}_{t-1} + \beta_t \mathbf{B}_{t-1} x_{t-1} + \gamma_t \mathbf{B}_t x_t, \quad (4)$$

where $\lambda_t \in [0, 1]$ is a data-dependent scalar, $\alpha_t := e^{\Delta_t A_t}$, $\beta_t := (1 - \lambda_t) \Delta_t e^{\Delta_t A_t}$, $\gamma_t := \lambda_t \Delta_t$.

Remark 1 (Expressivity). Our scheme is a generalization of a) The classical trapezoid rule which is recovered when $\lambda_t = \frac{1}{2}$. b) Mamba-2’s Euler’s rule, which is recovered when $\lambda_t = 1$.

Remark 2 (Error Rate). This is a second-order discretization with local truncation error $O(\Delta_t^3)$ and global error $O(\Delta_t^2)$ over the sequence under standard stability assumptions, provided that the trapezoidal parameter satisfies $\lambda_t = \frac{1}{2} + O(\Delta_t)$. However, our ablations indicate that not enforcing this constraint is the best for empirical performance. See Appendix B.2,B.3 for details.

3.1.1 TRAPEZOIDAL DISCRETIZATION IS A CONVOLUTIONAL MASK

We can view the generalized trapezoidal discretization as applying a *data-dependent* convolution of size two on the projected input, $\mathbf{B}_t x_t$, to the SSM. We now show that a similar vectorization to Equation (1) holds with the generalized trapezoidal discretization. Unrolling the recurrence starting from $\mathbf{h}_0 = \gamma_0 \mathbf{B}_0 x_0$ results in $\mathbf{h}_T = \alpha_{T-2}(\gamma_0 \alpha_1 + \beta_1) \mathbf{B}_0 x_0 + \dots + \gamma_T \mathbf{B}_T x_T$.

Unrolling these rows shows that the mask induced by the trapezoidal update is no longer a fixed averaging of endpoints (as in the classical trapezoidal rule), but a *data-dependent convex combination* of the two interval endpoints. In the SSD representation, this corresponds to a mask \mathbf{L} :

$$\begin{bmatrix} \gamma_0 \\ (\gamma_0 \alpha_1 + \beta_1) \\ \alpha_2(\gamma_0 \alpha_1 + \beta_1) \\ \vdots \\ \alpha_{T-2}(\gamma_0 \alpha_1 + \beta_1) \end{bmatrix} \gamma_2 \dots \gamma_T = \begin{bmatrix} 1 & & & & & & \gamma_0 \\ \alpha_1 & 1 & & & & & \beta_1 \\ \alpha_2 \alpha_1 & & \ddots & & & & 0 \\ \vdots & & & \ddots & & & \vdots \\ \alpha_{T-1} \alpha_{T-2} & \dots & 1 & & & & 0 \\ \alpha_{T-2}(\gamma_0 \alpha_1 + \beta_1) & \dots & \gamma_T & & & & \dots & \gamma_T \end{bmatrix}. \quad (5)$$

Here, the first factor is precisely the lower-triangular decay mask from Mamba-2, while the second factor encodes the size two convolution induced by the trapezoidal rule through the coefficients (β_t, γ_t) . We provide a rigorous proof for this decomposition in Appendix B.1.

3.2 COMPLEX-VALUED SSMs

Modern SSMs are designed with efficiency as the central goal, motivated by the need to scale to larger models and longer sequences. For instance, successive architectures have progressively simplified the state transition matrix: S4 (Gu et al., 2022a) used complex-valued Normal plus Low Rank (NPLR) matrices, Mamba (Gu & Dao, 2024) reduced this to a diagonal of reals, and Mamba-2 (Dao & Gu, 2024) further simplified it to a single scalar. Although these simplifications largely maintain language modeling performance, recent works (Merrill et al., 2025; Sarrof et al., 2024; Grazzi et al., 2025) have shown that they degrade the capabilities of the model on simple state-tracking tasks such as parity and modular arithmetic, which can be solved by a one-layer LSTM.

This limitation, formalized in Theorem-1 of (Grazzi et al., 2024), arises from restricting the eigenvalues of the transition matrix to real numbers, which cannot represent “rotational” hidden state dynamics. For instance, consider the parity function on binary inputs $\{0, 1\}$, defined as $\sum_t x_t \bmod 2$. This task can be performed using update: $\mathbf{h}_t = \mathbf{R}(\pi x_t) \mathbf{h}_{t-1}$, where $\mathbf{R}(\cdot)$ is a 2-D rotation matrix. Such rotational dynamics cannot be expressed with real eigenvalues.

To recover this capability, we begin with complex SSMs (6), which are capable of representing state-tracking dynamics. We show that, under discretization (Proposition 5), complex SSMs can be formulated as a real SSMs with a *block-diagonal transition matrix composed of 2×2 rotation matrices* (Proposition 2). We then show that this is equivalent to applying *data-dependent rotary embeddings* on both the input and output projections \mathbf{B}, \mathbf{C} respectively. This result establishes a theoretical connection between complex SSMs and data-dependent RoPE embeddings (Proposition 3). Finally, this allows for an efficient implementation of the complex-valued SSM via the “RoPE trick”, enabling efficient complex-valued state transition matrix with minimal computational overhead over real-valued SSMs.

Proposition 2 (Complex-to-Real SSM Equivalence). *Consider a complex-valued SSM*

$$\begin{aligned}\dot{\mathbf{h}}(t) &= \text{Diag}(A(t) + i\boldsymbol{\theta}(t)) \mathbf{h}(t) + (\mathbf{B}(t) + i\hat{\mathbf{B}}(t)) x(t), \\ y(t) &= \text{Re}\left((\mathbf{C}(t) + i\hat{\mathbf{C}}(t))^\top \mathbf{h}(t)\right),\end{aligned}\quad (6)$$

where $\mathbf{h}(t) \in \mathbb{C}^{N/2}$, $\boldsymbol{\theta}(t), \mathbf{B}(t), \hat{\mathbf{B}}(t), \mathbf{C}(t), \hat{\mathbf{C}}(t) \in \mathbb{R}^{N/2}$, and $x(t), A(t) \in \mathbb{R}$. Under Euler discretization, this system is equivalent to a real-valued SSM

$$\begin{aligned}\mathbf{h}_t &= e^{\Delta_t A_t} \mathbf{R}_t \mathbf{h}_{t-1} + \Delta_t \mathbf{B}_t x_t, \\ y_t &= \mathbf{C}_t^\top \mathbf{h}_t,\end{aligned}\quad (7)$$

with state $\mathbf{h}_t \in \mathbb{R}^N$, projections

$$\mathbf{B}_t = \begin{bmatrix} \mathbf{B}_t \\ \hat{\mathbf{B}}_t \end{bmatrix} \in \mathbb{R}^N, \quad \mathbf{C}_t = \begin{bmatrix} \mathbf{C}_t \\ -\hat{\mathbf{C}}_t \end{bmatrix} \in \mathbb{R}^N,$$

and a transition matrix

$$\mathbf{R}_t = \text{Block}\left(\{R(\Delta_t \boldsymbol{\theta}_t[i])\}_{i=1}^{N/2}\right) \in \mathbb{R}^{N \times N}, \quad R(\Theta) = \begin{bmatrix} \cos(\Theta) & -\sin(\Theta) \\ \sin(\Theta) & \cos(\Theta) \end{bmatrix}.$$

The proof is in Appendix C.1.

Proposition 2 shows that the discretized complex SSM has an equivalent real SSM with doubled state dimension (N), and a block-diagonal transition matrix multiplied with a scalar decay, where each 2×2 block is a data-dependent rotation matrix ($e_t^{\Delta_t A} \mathbf{R}_t$). We now show that the rotations can equivalently be absorbed into the input and output projections $\mathbf{B}_t, \mathbf{C}_t$, yielding an equivalent view that *complex SSMs are real SSMs equipped with data-dependent rotary embeddings (RoPE)*.

Proposition 3 (Complex SSM, Data-Dependent RoPE Equivalence). *Under the notation established in Proposition 2, consider the real SSM defined in Eq. 7 unrolled for T time-steps. The output of the above SSM is equivalent to that of a vanilla scalar transition matrix-based SSM (Eq. 2) with a data-dependent rotary embedding applied on the \mathbf{B}, \mathbf{C} components of the SSM defined as:*

$$\mathbf{h}_t = e^{\Delta_t A_t} \mathbf{h}_{t-1} + \left(\prod_{i=0}^t \mathbf{R}_i^\top\right) \mathbf{B}_t x_t, \quad \mathbf{y}_t = \left(\left(\prod_{i=0}^t \mathbf{R}_i^\top\right) \mathbf{C}_t\right)^\top \mathbf{h}_t \quad (8)$$

where the matrix production represents right matrix multiplication, e.g., $\prod_{i=0}^1 \mathbf{R}_i = \mathbf{R}_0 \mathbf{R}_1$. We denote employing the vanilla SSM to compute the Complex SSM as “RoPE trick”.

The proof is in Appendix C.2.

To observe the connection of complex SSMs to RoPE embeddings, note that in the above proposition, the data-dependent rotations \mathbf{R}_i are aggregated across time-steps and applied to \mathbf{C}, \mathbf{B} , which, by the State Space Duality of Dao & Gu (2024), correspond to the Query (\mathbf{Q}) and Key (\mathbf{K}) components of Attention. Analogously, vanilla RoPE (Su et al., 2023) applies *data-independent* rotation matrices, where the rotation angles follow a fixed frequency schedule $\boldsymbol{\theta}[i] = 10000^{-2i/N}$.

270 *Remark 3* (Generality). Proposition 3 extends to the fully general case where the transition is given
 271 by any complex matrix. By the complex diagonalization theorem, such a matrix is unitarily equivalent
 272 to a complex diagonal matrix, $\text{Diag}(\mathbf{A}(t) + i\theta(t))$ with $\mathbf{A}(t) \in \mathbb{R}^N$. However, in practice,
 273 we restrict $\mathbf{A}(t)$ to a scalar, mirroring the simplification from Mamba to Mamba-2, to enable faster
 274 implementation by avoiding GPU memory bottlenecks.

275 **Proposition 4** (Rotary Embedding Equivalence with Trapezoidal Discretization). *Discretizing a*
 276 *complex SSM with the trapezoidal rule (Proposition 1) yields the recurrence*

$$\begin{aligned} \mathbf{h}_t &= \alpha_t \mathbf{h}_{t-1} + \beta_t \left(\prod_{i=0}^{t-1} \mathbf{R}_i^\top \right) \mathbf{B}_{t-1} x_{t-1} + \gamma_t \left(\prod_{i=0}^t \mathbf{R}_i^\top \right) \mathbf{B}_t x_t, \\ \mathbf{y}_t &= \left(\left(\prod_{i=0}^t \mathbf{R}_i^\top \right) \mathbf{C}_t \right)^\top \mathbf{h}_t. \end{aligned} \quad (9)$$

284 Here \mathbf{R}_t is the block-diagonal rotation matrix defined in Proposition 3.

285 The proof is in Appendix C.3.

286 *Remark 4* (RoPE Trick). Complex SSMs discretized with the general trapezoidal rule of a complex
 287 SSM naturally admit the RoPE trick we established for SSMs discretized with Euler’s rule.

288 3.3 MULTI-INPUT, MULTI-OUTPUT

290 During the decoding phase of autoregressive inference, outputs are generated one token at a time, and
 291 performance is typically measured using in *Tokens generated Per Second (TPS)*. In this metric, sub-
 292 quadratic models, such as Mamba-2 (Dao & Gu, 2024), have a significant advantage over standard
 293 Transformer-style attention, since they feature a fixed-size hidden state (Equation (2)) rather than
 294 maintaining a key-value (KV) cache that grows linearly with the sequence length.

295 TPS, however, does not explicitly factor in hardware efficiency, where we aim to be in a compute-
 296 bound regime (as opposed to memory-bound) in order to fully utilize on-chip accelerators. To
 297 better characterize hardware efficiency, we would need to consider the arithmetic intensity of token
 298 generation. Recall that arithmetic intensity is defined as FLOPs divided by the number of input-
 299 output bytes, for a given op. In order to fully utilize both the accelerators and the bandwidth, we
 300 would like the arithmetic intensity to match the ops:byte ratio of the hardware, which in the case
 301 of NVIDIA H100-SXM5, is 295.2 bfloat16 ops per second with respect to the DRAM, and 31.9
 302 bfloat16 ops per second with respect to the SRAM [Fleetwood].

303 Table 2(a) shows the arithmetic intensity for a single generation in the SSM component of Mamba
 304 (with respect to 2-byte data). We see that it falls far short of a compute-bound regime, and moreover
 305 it is not clear how one can adjust the existing parameters in Mamba to mitigate the lack of hardware
 306 efficiency. We note that this observation applies generally to other sub-quadratic models, such as
 307 causal linear attention.

308 Input	309 Output	310 FLOPs	311 Arithmetic 312 Intensity	308 Input	309 Output	310 FLOPs	311 Arithmetic 312 Intensity
$H_t : (n, p)$	$y_t : (p)$	$5pn$	$\frac{5pn}{2(1 + 2n + p + np)}$ $\approx 2.5 = \Theta(1)$	$H_t : (n, p)$	$y_t : (p)$	$4nrp + 2np$	$\frac{p(4nr + 2n)}{2(1 + 2nr + pr + np)}$ $\approx 2r = \Theta(r)$
$x_t : (p)$				$x_t : (p, r)$	(p, r)		
$a_t : (1)$				$a_t : (1)$			
$b_t : (n)$				$b_t : (n, r)$			
$c_t : (n)$				$c_t : (n, r)$			

315 (a) SISO (2-byte data).

316 (b) MIMO (2-byte data).

317 Figure 2: Arithmetic Intensity for (a) SISO, (b) MIMO. Batch and head dimensions cancel out.

318 In light of this, we made the following simple adjustment to our recurrent relation: instead of trans-
 319 forming the input $\mathbf{x}_t \in \mathbb{R}^p$ to state $\mathbf{H}_t \in \mathbb{R}^{n \times p}$ via an outer product, i.e., $\mathbf{H}_t \leftarrow a_t \mathbf{H}_{t-1} + b_t \otimes \mathbf{x}_t$, we
 320 made such a transformation via a matrix product, i.e., $\mathbf{H}_t \leftarrow a_t \mathbf{H}_{t-1} + \mathbf{B}_t \mathbf{X}_t^\top$, where $\mathbf{B}_t \in \mathbb{R}^{n \times r}$
 321 and $\mathbf{X}_t \in \mathbb{R}^{p \times r}$ are now matrices with an additional rank r . The emission from state to output
 322 similarly acquire an extra rank r , i.e., $\mathbf{Y}_t \in \mathbb{R}^{r \times p} \leftarrow \mathbf{C}_t^\top \mathbf{H}_t$, where $\mathbf{C}_t \in \mathbb{R}^{n \times r}$, $\mathbf{H}_t \in \mathbb{R}^{n \times p}$.
 323 This simple change increases the arithmetic intensity of recurrence, which now scales with the rank

324 r (Figure 2(b)). Hence, by increasing r , arithmetic intensity improves and shifts decode generation
 325 towards a more compute-bound regime. This increase in FLOPs during decode does not compromise
 326 runtime, as the operation is bounded by the I/O of state $\mathbf{H}_t \in \mathbb{R}^{n \times p}$.

327 Moreover, moving from outer-product-based state update to matrix-product-based coincides exactly
 328 with generalizing from SISO to MIMO SSM, with the rank r being the MIMO rank. Such a generalization
 329 recovers a key expressive feature of SSMs in classical literature; indeed, there has been
 330 previous work, namely Smith et al. (2023), that explored MIMO SSM as a drop-in replacement of
 331 attention, albeit not in the context of Mamba and not necessarily with inference in view. **We note**
 332 **that training and prefilling is generally compute bound, resulting in MIMO incurring increased costs**
 333 **during these stages, while decoding, a memory-bound operation, sees very little increase in latency**
 334 **when utilizing MIMO over SISO.**

335 Details of the MIMO formulation for Mamba-3 are provided in Appendix D.

336 3.4 MAMBA-3 ARCHITECTURE

337 The Mamba-3 block retains the overall layout of its predecessor while introducing several key modi-
 338 fications. Most notably, the SSD layer is replaced with the more expressive trapezoidal SSM defined
 339 in Proposition 4. The extra normalization layer, first introduced between Mamba-1 and Mamba-2 for
 340 training stability, is repositioned to follow the \mathbf{B} , \mathbf{C} projection, mirroring the QK-Norm commonly
 341 used in modern Transformers (Henry et al., 2020; Wortsman et al., 2023). Inspired by the findings
 342 of Yu & Erichson (2025), which prove adding channel-specific bias to \mathbf{B} in a blockwise variant
 343 of Mamba-1 grants universal approximation capabilities, Mamba-3 incorporates a head-specific,
 344 channel-wise bias into both the \mathbf{B} and \mathbf{C} components after its normalization. **These learnable bi-**
 345 **ases are data-independent parameters that are initialized to all ones and independent across \mathbf{B} and**
 346 **\mathbf{C} (ablations for bias parameterization can be found in Appendix G).** Our trapezoidal discretization
 347 complements this bias, empirically eliminating the need for the original short causal convolution and
 348 its accompanying activation function (Section 4.3). Mamba-3 employs the SISO SSM by default,
 349 though we view its MIMO variant as a flexible option that can be toggled depending on inference
 350 requirements. The overall architecture follows the Llama design (Grattafiori et al., 2024), alternating
 351 Mamba-3 and SwiGLU blocks with pre-normalization.

352 4 EMPIRICAL VALIDATION

353 We empirically validate our SSM-centric methodological changes through the Mamba-3 model on
 354 a host of synthetic and real world tasks. Section 4.1 compares our SISO-variant of Mamba-3 on
 355 language modeling and retrieval-based tasks, while Section 4.2 demonstrates inference efficiency of
 356 Mamba-3 and MIMO Mamba-3’s benefits over SISO Mamba-3 under fixed inference compute. We
 357 ablate the impact of our new discretization and BC bias on performance and show that complexifica-
 358 tion of the SSM leads capabilities that prior SSMs such as Mamba-2 lacked in Section 4.3.

359 4.1 LANGUAGE MODELING

360 All models are pretrained with 100B tokens of the FineWeb-Edu dataset (Penedo et al., 2024) with
 361 the Llama-3.1 tokenizer (Grattafiori et al., 2024) at a 2K context length with the same standard
 362 training protocol. Training and evaluation details can be found in Appendix E.

363 Across all four model scales, Mamba-3 outperforms popular baselines at various downstream tasks
 364 (Table 1). We highlight that Mamba-3 does not utilize the short convolution that has been empirically
 365 identified as an important component in many performant linear models (Allen-Zhu, 2025).

366 4.1.1 RETRIEVAL CAPABILITIES

367 Beyond standard language modeling, an important measure for linear models is their retrieval ability
 368 — how well they can recall information from earlier in the sequence (Arora et al., 2025a;b). Unlike
 369 attention models, which can freely revisit past context with the growing KV cache, linear models
 370 must compress context into a fixed-size state. This trade-off is reflected in the Transformer baseline’s
 371 substantially stronger retrieval scores. To evaluate Mamba-3 under this lens, Table 2 compares it
 372 against baselines on both real-world and synthetic needle-in-a-haystack (NIAH) tasks (Hsieh et al.,
 373 2024), using our pretrained 1.5B models from Section 4.1. We restrict the task sequence length to
 374 2K tokens to match the training setup and adopt the cloze-style format for our real-world tasks to
 375 mirror the next-token-prediction objective, following Arora et al. (2025b; 2024).

376 Mamba-3 is competitive on real-world associative recall and question-answering but struggles when
 377 extracting information from semi-structured or unstructured data. On synthetic NIAH tasks, how-

378
 379 Table 1: Downstream language modeling evaluations on models trained with 100B FineWeb-Edu
 380 tokens. Best results for each size are **bolded**, and second best are underlined. All models are trained
 381 with the same procedure. Mamba-3 outperforms Mamba-2 and others at every model scale.

Model	FW-Edu ppl ↓	LAMB. ppl ↓	LAMB. acc ↑	HellaS. acc..n ↑	PIQA acc ↑	Arc-E acc ↑	Arc-C acc..n ↑	WinoGr. acc ↑	OBQA acc ↑	Average acc ↑
Transformer-180M	16.89	45.0	<u>32.5</u>	39.0	67.1	59.8	27.9	51.2	21.8	42.8
Gated DeltaNet-180M	<u>16.61</u>	35.9	<u>33.7</u>	<u>40.2</u>	<u>66.8</u>	59.6	28.5	51.2	21.6	<u>43.1</u>
Mamba-2-180M	16.76	41.8	30.9	40.1	<u>66.8</u>	<u>60.1</u>	27.3	52.0	23.2	42.9
Mamba-3-180M (SISO)	16.59	<u>37.7</u>	<u>32.5</u>	40.8	66.1	61.5	27.9	52.0	22.8	43.4
Transformer-440M	13.03	21.2	41.7	50.5	69.9	67.6	34.6	56.7	26.0	49.6
Gated DeltaNet-440M	13.12	19.0	40.4	50.5	70.5	67.5	34.0	55.3	25.8	49.1
Mamba-2-440M	<u>13.00</u>	<u>19.6</u>	<u>40.8</u>	51.7	70.6	<u>68.8</u>	35.0	54.1	26.0	49.6
Mamba-3-440M (SISO)	12.87	19.6	40.2	<u>51.7</u>	<u>71.9</u>	68.9	34.4	55.8	26.0	49.8
Transformer-880M	11.42	15.0	44.7	57.2	72.6	71.6	39.2	57.7	26.8	52.8
Gated DeltaNet-880M	11.39	12.7	<u>47.1</u>	57.5	<u>72.6</u>	<u>72.5</u>	38.8	<u>57.9</u>	30.6	53.9
Mamba-2-880M	<u>11.35</u>	13.8	45.0	58.1	<u>72.5</u>	<u>72.3</u>	38.7	56.8	30.2	53.4
Mamba-3-880M (SISO)	11.23	<u>12.9</u>	47.2	58.8	73.6	72.7	40.2	58.4	30.0	54.4
Transformer-1.5B	10.51	11.1	50.3	60.6	73.8	74.0	40.4	58.7	29.6	55.4
Gated DeltaNet-1.5B	10.51	10.8	<u>49.9</u>	60.5	74.3	73.3	40.4	61.5	30.4	55.7
Mamba-2-1.5B	10.47	12.0	<u>47.8</u>	61.4	73.6	<u>75.3</u>	41.8	57.5	32.6	<u>55.7</u>
Mamba-3-1.5B (SISO)	10.35	<u>10.9</u>	49.4	61.9	73.6	75.9	42.7	59.4	32.0	56.4

396
 397 Table 2: Retrieval capabilities measured by a mixture of real-world and synthetic retrieval tasks. Real-world re-
 398 trieval tasks utilize cloze variants of the original datasets and are truncated to 2K length. Mamba-3 demonstrates
 399 strong associative recall and question-answering but suffers with information extraction of semi-structured and
 400 unstructured data. Mamba-3 has strong needle-in-a-haystack (NIAH) accuracy and generalizes outside its
 401 trained context.

Model (1.5B)	SWDE	SQuAD	FDA	TQA	NQ	Drop	NIAH-Single-1			NIAH-Single-2			NIAH-Single-3		
							1024	2048	4096	1024	2048	4096	1024	2048	4096
Context Length				2048											
Transformer	48.9	46.6	58.4	67.5	31.7	26.4	100.0	100.0	0.0	92.2	100.0	0.0	98.6	99.4	0.0
Gated DeltaNet	32.7	<u>40.0</u>	28.3	63.5	<u>25.7</u>	24.5	100.0	100.0	99.8	100.0	93.8	49.8	83.8	68.4	34.2
Mamba-2	<u>30.7</u>	<u>39.1</u>	<u>23.7</u>	<u>64.3</u>	<u>25.1</u>	<u>28.5</u>	<u>100.0</u>	99.6	62.0	<u>100.0</u>	<u>53.8</u>	11.8	95.8	87.4	13.4
Mamba-3 (SISO)	28.5	40.1	23.4	64.5	26.5	<u>27.4</u>	100.0	100.0	<u>88.2</u>	100.0	95.4	50.6	<u>92.4</u>	<u>81.4</u>	34.2

408
 409 ever, Mamba-3 surpasses or matches baselines on most cases and notably demonstrates markedly
 410 better out-of-distribution retrieval abilities than its Mamba-2 predecessor.

4.2 INFERENCE EFFICIENCY

411 In this section, we investigate our methodological changes in the context of inference performance.
 412 We first present our inference benchmark in Section 4.2.1; we then establish a framework for com-
 413 paring the inference performance in Section 4.2.2. Finally, we focus on the effectiveness of MIMO
 414 in Section 4.2.3.

4.2.1 FAST MAMBA-3 KERNELS

415 We complement Mamba-3’s methodological advances with optimized kernels that deliver fast infer-
 416 ence in practical settings. Specifically, we implement a new series of inference kernels for Mamba-
 417 3—using Triton for the forward (prefill) path and CuTe-DSL for decode—and compare their per-
 418 token decode latency against the released Triton kernels for Mamba-2 and Gated DeltaNet (GDN)¹
 419 in Table 3. The evaluation uses the setting: a decode step at batch size 128 on a single H100 for
 420 1.5B-parameter models with model dimension 2048, state dimension $\in \{64, 128\}$ in both FP32 and
 421 BF16 datatypes. Across all configurations, SISO achieves the lowest latency amongst baselines,
 422 while MIMO incurs only a minor overhead relative to SISO. This indicates that our CuTe-DSL de-
 423 code implementation is competitive and that the additional components of Mamba-3 (trapezoidal
 424 update, complex-valued state, and MIMO projections) are lightweight. This supports our overall
 425 inference-first perspective: the Mamba-3 admits **simple, low-latency implementation** while pro-
 426 viding strong empirical performance. A thorough analysis, including prefill and prefill with decode
 427 results are provided in Appendix H.

432	Model	FP32		BF16	
433		$d_{\text{state}} = 64$	$d_{\text{state}} = 128$	$d_{\text{state}} = 64$	$d_{\text{state}} = 128$
434	Mamba-2	0.295	0.409	0.127	0.203
435	Gated DeltaNet	0.344	0.423	0.176	0.257
436	Mamba-3 (SISO)	0.261	0.356	0.106	0.152
437	Mamba-3 (MIMO)	0.285	0.392	0.136	0.185

440
441 Table 3: Latency (in milliseconds) compari-
442 son across models, precision, and d_{state} val-
443 ues. Both Mamba-3 SISO and MIMO are
444 faster than the Mamba-2 and Gated DeltaNet
445 at the commonly used bf16, $d_{\text{state}} = 128$ set-
446 ting.

447 4.2.2 PARETO FRONTIER FOR INFERENCE EFFICIENCY

448 For Mamba and many variants of sub-quadratic models, the generation of tokens during decoding is
449 heavily dominated by memory I/O due to the low arithmetic intensity of computing the recurrent update (c.f. Section 3.3). Furthermore, among the data being transferred, the latent state \mathbf{H}_t dominates
450 in terms of size. Indeed, from Table 3, we see that the runtime scales with d_{state} , which configures
451 the size of the hidden state.

452 As d_{state} dominates the decode runtime for the subquadratic models considered in this paper, we
453 opt to use it as a proxy for inference speed. By plotting the validation perplexity (itself a proxy
454 for model performance) as a function of d_{state} , we aim to formulate a holistic picture about how the
455 subquadratic models can trade off performance with inference speed.

456 Figure 3 shows such a Pareto front for the Mamba variants models considered in this paper. For each
457 data point, we train a 440M parameter model to $2 \times$ Chinchilla optimal tokens on the Fineweb-Edu
458 dataset, where the model is configured with a d_{state} of $\{16, 32, 64, 128\}$. As expected, we observe
459 an inverse correlation between validation loss and d_{state} ; moreover, we noticed a general downward
460 shift on the Pareto front moving from Mamba-2 to Mamba-3. A further downward shift is observed
461 when moving from the SISO variant of Mamba-3 to the MIMO variant of Mamba-3 (where we set
462 the Mimo rank $r = 4$ and decrease our MLP inner dimension to parameter match the SISO variants).
463 We expand the comparison to include the Gated DeltaNet baseline in Figure 7. The results highlight
464 both the expressivity gain coming our methodology change as well as the effectiveness of the MIMO
465 mechanism in improving decoding efficiency.

466 4.2.3 MIMO ENHANCES INFERENCE EFFICIENCY


467 MIMO, with its higher arithmetic intensity, increases the decoding FLOPs without significantly
468 increasing decode runtime (Table 3)² The implication is that any performance gain from MIMO
469 translates into efficiency gain in decoding: a conclusion supported by the downward shift of the
470 MIMO pareto curve we observed in Section 4.2.2.

471 We aim to further verify the gain from MIMO by investigating its language-modeling capabilities.
472 To that end, we train a 440M and 820M parameter MIMO models with MIMO rank $r = 4$ on 100B
473 tokens on Fineweb-Edu (i.e., same setting as the 440M parameter run in Section 4.1; we are currently
474 training the 1.5B model). To ensure the total parameter count equals SISO, we decrease the inner
475 dimension of the MLP layers to compensate for the increase due to the MIMO projections.

476 On both validation perplexity and our suite of language evaluation tasks (Table 6), we see significant
477 gain when moving from SISO to MIMO. Namely, we attain a perplexity gain of 0.16 on the 100B
478 tokens run, and Figure 3 illustrates the downward shift in our validation loss. On the language
479 evaluation front, we see significant gain on most tasks when compared to SISO, resulting in an
480 overall gain of 1.2 point over SISO. This strongly supports MIMO as a SSM-centric technique to
481 improve model quality without compromising decoding speed.

482 ¹Details on each kernel DSL and the exact kernel fusion structure is provided in Appendix H.

483 ²The kernel for MIMO Mamba-3 in fact fuses the MIMO projection, and so the reported wall clock time is
484 actually an overestimate for the pure SSM update.



485 Figure 3: Exploration of state size (inference
486 speed proxy) versus pretraining perplexity (per-
487 formance proxy) across different Mamba variants.
488 Mamba-3 MIMO drives the Pareto frontier without
489 increasing state size.

486
 487 **Table 4: Left:** Ablations on core modeling components of Mamba-3, results on test split of dataset. A
 488 combination of our BC bias and trapezoidal discretization makes the convolution optional. **Right:** Formal
 489 language evaluation (scaled accuracy, %). Higher is better. Models are trained on short sequences and evaluated
 490 on longer lengths to test length generalization. For Gated DeltaNet we report the variant with eigenvalue range
 491 $[-1, 1]$.
 492

Model Variant (SISO)	ppl \downarrow
Mamba-3 – bias – trap	16.68
Mamba-3 – bias	16.49
Mamba-3	15.72
Mamba-3 + conv	15.85

(a) Component ablation (350M).

Model	Parity \uparrow	Arith. w/o \uparrow brackets	Arith. w/ \uparrow brackets
Mamba-3	100.00	98.51	87.75
Mamba-3 (w/o RoPE)	2.27	1.49	0.72
Mamba-3 (w/ Std. RoPE)	1.56	20.70	2.62
Mamba-2	0.90	47.81	0.88
Gated DeltaNet [-1,1]	100.00	99.25	93.50

(b) Performance comparison on formal language tasks. Results show that unlike Mamba-2, Mamba-3 features state tracking ability stemming from data-dependent RoPE embeddings. **We used Mamba-3 (SISO) for these ablations.**

503 4.3 SSM-CENTRIC METHODOLOGICAL ABLATIONS

504 Table 4a ablates the changes made to the core SSM component, mainly the introduction of BC bias
 505 and trapezoidal discretization. We report the pretraining test perplexity on models at the 440M scale,
 506 trained for Chinchilla optimal tokens. We find that the bias and trapezoidal SSM synergize well and
 507 make the short convolution utilized by many current linear models redundant.

508 We empirically demonstrate that data-dependent RoPE in Mamba-3 enables state tracking. Following
 509 Grazzi et al. (2025), we evaluate on tasks from the Chomsky hierarchy—Parity, Modular Arithmetic
 510 (without brackets), and Modular Arithmetic (with brackets)—and report scaled accuracies in
 511 Table 4b. Mamba-3 solves Parity and Modular Arithmetic (without brackets), and nearly closes the
 512 accuracy gap on Modular Arithmetic (with brackets). In contrast, Mamba-3 without RoPE, **Mamba-3 with standard RoPE** (Su et al., 2023), and Mamba-2 fail to learn these tasks. We use the state-
 513 tracking-enabled *Gated DeltaNet* variant of and observe that *Mamba-3* is competitive—matching
 514 parity and approaching its performance on both modular-arithmetic tasks. Experimental settings are
 515 covered in Appendix E.

517 5 CONCLUSION AND FUTURE WORK

518 We introduce Mamba-3, an SSM model with three axes of improvement rooted in SSM principles: (i) *improved quality*, via trapezoidal discretization; (ii) *new capabilities*, through complex
 519 SSMs that recover state-tracking; and (iii) *higher inference efficiency*, with a MIMO formulation
 520 that raises arithmetic intensity. Mamba-3 delivers strong language modeling results and establishes
 521 a new Pareto frontier on the performance-efficiency axes with respect to strong baseline models. A
 522 limitation remains in retrieval, where fixed-state architectures lags attention-based models. We see
 523 **hybrid Mamba-3 architectures** that integrate retrieval mechanisms as a promising path, alongside
 524 broader application of our design principles to linear-time sequence models.

526
 527
 528
 529
 530
 531
 532
 533
 534
 535
 536
 537
 538
 539

540 REFERENCES

541
 542 Zeyuan Allen-Zhu. Physics of Language Models: Part 4.1, Architecture Design and the Magic
 543 of Canon Layers. *SSRN Electronic Journal*, May 2025. <https://ssrn.com/abstract=5240330>.

544
 545 Aryaman Arora, Neil Rathi, Nikil Roashan Selvam, Róbert Csordás, Dan Jurafsky, and Christopher
 546 Potts. Mechanistic evaluation of transformers and state space models, 2025a. URL <https://arxiv.org/abs/2505.15105>.

547
 548 Simran Arora, Aman Timalsina, Aaryan Singhal, Benjamin Spector, Sabri Eyuboglu, Xinyi Zhao,
 549 Ashish Rao, Atri Rudra, and Christopher Ré. Just read twice: closing the recall gap for recurrent
 550 language models, 2024. URL <https://arxiv.org/abs/2407.05483>.

551
 552 Simran Arora, Sabri Eyuboglu, Michael Zhang, Aman Timalsina, Silas Alberti, Dylan Zinsley,
 553 James Zou, Atri Rudra, and Christopher Ré. Simple linear attention language models balance
 554 the recall-throughput tradeoff, 2025b. URL <https://arxiv.org/abs/2402.18668>.

555
 556 Aviv Bick, Kevin Y. Li, Eric P. Xing, J. Zico Kolter, and Albert Gu. Transformers to ssms: Distill-
 557 ing quadratic knowledge to subquadratic models, 2025a. URL <https://arxiv.org/abs/2408.10189>.

558
 559 Aviv Bick, Eric Xing, and Albert Gu. Understanding the skill gap in recurrent language models:
 560 The role of the gather-and-aggregate mechanism, 2025b. URL <https://arxiv.org/abs/2504.18574>.

561
 562 Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. Piqa: Reasoning about
 563 physical commonsense in natural language, 2019. URL <https://arxiv.org/abs/1911.11641>.

564
 565 Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas
 566 Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, David Belanger, Lucy
 567 Colwell, and Adrian Weller. Rethinking attention with performers, 2022. URL <https://arxiv.org/abs/2009.14794>.

568
 569 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and
 570 Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge,
 571 2018. URL <https://arxiv.org/abs/1803.05457>.

572
 573 Tri Dao and Albert Gu. Transformers are ssms: Generalized models and efficient algorithms through
 574 structured state space duality, 2024. URL <https://arxiv.org/abs/2405.21060>.

575
 576 Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner.
 577 Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs, 2019.
 578 URL <https://arxiv.org/abs/1903.00161>.

579
 580 Christopher Fleetwood. Domain specific architectures for ai inference. URL <https://fleetwood.dev/posts/domain-specific-architectures>.

581
 582 Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Fos-
 583 ter, Laurence Golding, Jeffrey Hsu, Alain Le Noac'h, Haonan Li, Kyle McDonell, Niklas Muen-
 584 nighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang
 585 Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. The language model
 586 evaluation harness, 07 2024. URL <https://zenodo.org/records/12608602>.

587
 588 Madan Gopal. *Modern control system theory*. New Age International, 1993.

589
 590 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 591 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan,
 592 Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Ko-
 593 rennev, Arthur Hinsvark, Arun Rao, Aston Zhang, and et. al. The llama 3 herd of models, 2024.
 594 URL <https://arxiv.org/abs/2407.21783>.

594 Riccardo Grazzi, Julien Siems, Simon Schrodi, Thomas Brox, and Frank Hutter. Is mamba capable
 595 of in-context learning?, 2024. URL <https://arxiv.org/abs/2402.03170>.

596

597 Riccardo Grazzi, Julien Siems, Arber Zela, Jörg K. H. Franke, Frank Hutter, and Massimiliano
 598 Pontil. Unlocking state-tracking in linear rnns through negative eigenvalues, 2025. URL <https://arxiv.org/abs/2411.12537>.

599

600 Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces, 2024.
 601 URL <https://arxiv.org/abs/2312.00752>.

602

603 Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured
 604 state spaces, 2022a. URL <https://arxiv.org/abs/2111.00396>.

605

606 Albert Gu, Ankit Gupta, Karan Goel, and Christopher Ré. On the parameterization and initialization
 607 of diagonal state space models. *arXiv preprint arXiv:2206.11893*, 2022b. URL <https://arxiv.org/abs/2206.11893>.

608

609 Ankit Gupta, Albert Gu, and Jonathan Berant. Diagonal state spaces are as effective as structured
 610 state spaces, 2022. URL <https://arxiv.org/abs/2203.14343>.

611

612 Alex Henry, Prudhvi Raj Dachapally, Shubham Pawar, and Yuxuan Chen. Query-key normalization
 613 for transformers, 2020. URL <https://arxiv.org/abs/2010.04245>.

614

615 Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, Yang
 616 Zhang, and Boris Ginsburg. Ruler: What's the real context size of your long-context language
 617 models?, 2024. URL <https://arxiv.org/abs/2404.06654>.

618

619 Samy Jelassi, David Brandfonbrener, Sham M. Kakade, and Eran Malach. Repeat after me: Trans-
 620 formers are better than state space models at copying, 2024. URL <https://arxiv.org/abs/2402.01032>.

621

622 Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly
 623 supervised challenge dataset for reading comprehension, 2017. URL <https://arxiv.org/abs/1705.03551>.

624

625 Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. 1960.

626

627 Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are
 628 rnns: Fast autoregressive transformers with linear attention, 2020. URL <https://arxiv.org/abs/2006.16236>.

629

630 Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris
 631 Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion
 632 Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav
 633 Petrov. Natural questions: A benchmark for question answering research. *Transactions of the
 634 Association for Computational Linguistics*, 7:452–466, 2019. doi: 10.1162/tacl_a_00276. URL
 635 <https://aclanthology.org/Q19-1026/>.

636

637 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E.
 638 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model
 639 serving with pagedattention, 2023. URL <https://arxiv.org/abs/2309.06180>.

640

641 Baolin Li, Yankai Jiang, Vijay Gadepally, and Devesh Tiwari. Llm inference serving: Survey of
 642 recent advances and opportunities, 2024. URL <https://arxiv.org/abs/2407.12391>.

643

644 William Merrill, Jackson Petty, and Ashish Sabharwal. The illusion of state in state-space models,
 645 2025. URL <https://arxiv.org/abs/2404.08819>.

646

647 Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct
 648 electricity? a new dataset for open book question answering, 2018. URL <https://arxiv.org/abs/1809.02789>.

648 Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia,
 649 Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord,
 650 Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha
 651 Dziri, Michal Guerquin, and et. al. 2 olmo 2 furious, 2025. URL <https://arxiv.org/abs/2501.00656>.

653 Antonio Orvieto, Samuel L Smith, Albert Gu, Anushan Fernando, Caglar Gulcehre, Razvan Pas-
 654 canu, and Soham De. Resurrecting recurrent neural networks for long sequences, 2023. URL
 655 <https://arxiv.org/abs/2303.06349>.

656 Daniele Paliotta, Junxiong Wang, Matteo Pagliardini, Kevin Y. Li, Aviv Bick, J. Zico Kolter, Albert
 657 Gu, François Fleuret, and Tri Dao. Thinking slow, fast: Scaling inference compute with distilled
 658 reasoners, 2025. URL <https://arxiv.org/abs/2502.20339>.

660 Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Quan Ngoc Pham, Raffaella Bernardi,
 661 Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. The lambada dataset:
 662 Word prediction requiring a broad discourse context, 2016. URL <https://arxiv.org/abs/1606.06031>.

664 Jongho Park, Jaeseung Park, Zheyang Xiong, Nayoung Lee, Jaewoong Cho, Samet Oymak, Kang-
 665 wook Lee, and Dimitris Papailiopoulos. Can mamba learn how to learn? a comparative study on
 666 in-context learning tasks, 2024. URL <https://arxiv.org/abs/2402.04248>.

667 668 Guilherme Penedo, Hynek Kydlíček, Loubna Ben allal, Anton Lozhkov, Margaret Mitchell, Colin
 669 Raffel, Leandro Von Werra, and Thomas Wolf. The fineweb datasets: Decanting the web for the
 670 finest text data at scale, 2024. URL <https://arxiv.org/abs/2406.17557>.

671 Bo Peng, Ruichong Zhang, Daniel Goldstein, Eric Alcaide, Xingjian Du, Haowen Hou, Jiaju Lin,
 672 Jiaxing Liu, Janna Lu, William Merrill, Guangyu Song, Kaifeng Tan, Saiteja Utpala, Nathan
 673 Wilce, Johan S. Wind, Tianyi Wu, Daniel Wuttke, and Christian Zhou-Zheng. Rwkv-7 "goose"
 674 with expressive dynamic state evolution, 2025. URL <https://arxiv.org/abs/2503.14456>.

675 676 Pranav Rajpurkar, Jian Zhang, and Percy Liang. Know what you don't know: Unanswerable ques-
 677 tions for squad. In *ACL 2018*, 2018.

678 679 Yuval Ran-Milo, Eden Lumbroso, Edo Cohen-Karlik, Raja Giryes, Amir Globerson, and Nadav
 680 Cohen. Provable benefits of complex parameterizations for structured state space models, 2024.
 681 URL <https://arxiv.org/abs/2410.14067>.

682 683 Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adver-
 684 sarial winograd schema challenge at scale, 2019. URL <https://arxiv.org/abs/1907.10641>.

685 686 Yash Sarrof, Yana Veitsman, and Michael Hahn. The expressive capacity of state space models: A
 687 formal language perspective, 2024. URL <https://arxiv.org/abs/2405.17394>.

688 689 Imanol Schlag, Kazuki Irie, and Jürgen Schmidhuber. Linear transformers are secretly fast weight
 690 programmers, 2021. URL <https://arxiv.org/abs/2102.11174>.

691 692 Julien Siems, Timur Carstensen, Arber Zela, Frank Hutter, Massimiliano Pontil, and Riccardo
 693 Grazzi. Deltaproduct: Improving state-tracking in linear rnns via householder products, 2025.
 694 URL <https://arxiv.org/abs/2502.10297>.

695 Jimmy T. H. Smith, Andrew Warrington, and Scott W. Linderman. Simplified state space layers for
 696 sequence modeling, 2023. URL <https://arxiv.org/abs/2208.04933>.

697 698 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally
 699 can be more effective than scaling model parameters, 2024. URL <https://arxiv.org/abs/2408.03314>.

700 701 Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. Roformer: En-
 702 hanced transformer with rotary position embedding, 2023. URL <https://arxiv.org/abs/2104.09864>.

702 Yutao Sun, Li Dong, Shaohan Huang, Shuming Ma, Yuqing Xia, Jilong Xue, Jianyong Wang, and
 703 Furu Wei. Retentive network: A successor to transformer for large language models, 2023. URL
 704 <https://arxiv.org/abs/2307.08621>.

705 Endre Süli and David F. Mayers. *An Introduction to Numerical Analysis*. Cambridge University
 706 Press, 2003.

708 Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej,
 709 Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas
 710 Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon,
 711 Etienne Pot, Ivo Penchev, Gaël Liu, and et. al. Gemma 3 technical report, 2025. URL <https://arxiv.org/abs/2503.19786>.

713 M. Tenenbaum and H. Pollard. *Ordinary Differential Equations: An Elementary Textbook for Stu-
 714 dents of Mathematics, Engineering, and the Sciences*. Dover Books on Mathematics. Dover Pub-
 715 lications, 1985. ISBN 9780486649405. URL <https://books.google.com/books?id=iU4zDAAAQBAJ>.

717 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
 718 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information
 719 processing systems*, pp. 5998–6008, 2017. URL <https://arxiv.org/abs/1706.03762>.

721 Johannes von Oswald, Nino Scherrer, Seijin Kobayashi, Luca Versari, Songlin Yang, Maximil-
 722 ian Schlegel, Kaitlin Maile, Yanick Schimpf, Oliver Sieberling, Alexander Meulemans, Rif A.
 723 Sauro, Guillaume Lajoie, Charlotte Frenkel, Razvan Pascanu, Blaise Agüera y Arcas, and João
 724 Sacramento. Mesanet: Sequence modeling by locally optimal test-time training, 2025. URL
 725 <https://arxiv.org/abs/2506.05233>.

726 Mitchell Wortsman, Peter J. Liu, Lechao Xiao, Katie Everett, Alex Alemi, Ben Adlam, John D. Co-
 727 Reyes, Izzeddin Gur, Abhishek Kumar, Roman Novak, Jeffrey Pennington, Jascha Sohl-dickstein,
 728 Kelvin Xu, Jaehoon Lee, Justin Gilmer, and Simon Kornblith. Small-scale proxies for large-scale
 729 transformer training instabilities, 2023. URL <https://arxiv.org/abs/2309.14322>.

731 Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. Inference scaling laws:
 732 An empirical analysis of compute-optimal inference for problem-solving with language models,
 733 2025. URL <https://arxiv.org/abs/2408.00724>.

734 Songlin Yang, Jan Kautz, and Ali Hatamizadeh. Gated delta networks: Improving mamba2 with
 735 delta rule, 2025a. URL <https://arxiv.org/abs/2412.06464>.

737 Songlin Yang, Bailin Wang, Yu Zhang, Yikang Shen, and Yoon Kim. Parallelizing linear trans-
 738 formers with the delta rule over sequence length, 2025b. URL <https://arxiv.org/abs/2406.06484>.

740 Annan Yu and N. Benjamin Erichson. Block-biased mamba for long-range sequence processing,
 741 2025. URL <https://arxiv.org/abs/2505.09022>.

742 Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a ma-
 743 chine really finish your sentence?, 2019. URL <https://arxiv.org/abs/1905.07830>.

745
 746
 747
 748
 749
 750
 751
 752
 753
 754
 755

756 **LLM Usage.** We utilized Large Language Models to polish the writing in our submission as well as
 757 generate latex code for formatting tables and figures.
 758

759 A RELATED WORK

760 **Linear-time sequence mixers.** State-space models (SSMs) provide linear-time sequence mixing
 761 through explicit dynamical states and efficient scan/convolution implementations, offering signifi-
 762 cant computational advantages over quadratic-time attention mechanisms (Gu et al., 2022a; Smith
 763 et al., 2023; Gupta et al., 2022). Mamba-1 (Gu & Dao, 2024) introduced input-dependent selectivity
 764 to SSMs, while Mamba-2 (Dao & Gu, 2024) formalized the connection between SSMs and attention
 765 via structured state-space duality (SSD) (Katharopoulos et al., 2020; Choromanski et al., 2022). De-
 766 spite matching transformers on standard language understanding benchmarks, these recurrent mod-
 767 els exhibit limitations on tasks requiring precise algorithmic reasoning. Recent evaluations identified
 768 gaps in capabilities such as associative retrieval (Bick et al., 2025b; Arora et al., 2025a), exact copy-
 769 ing (Jelassi et al., 2024), and in-context learning (Park et al., 2024; Grazzi et al., 2024). To address
 770 these limitations, DeltaNet enhances linear attention by replacing additive updates with delta-rule
 771 recurrence (Schlag et al., 2021), with recent work developing hardware-efficient, sequence-parallel
 772 training algorithms for this architecture (Yang et al., 2025b). This has catalyzed a broader effort
 773 to improve the algorithmic capabilities of linear-time models through architectural innovations in-
 774 cluding gating mechanisms, improved state transition dynamics, and hybrid approaches (Peng et al.,
 775 2025; Siems et al., 2025; Yang et al., 2025a; Paliotta et al., 2025; Bick et al., 2025a).

776 **Expressivity and state tracking in recurrent mixers.** Recent work characterizes the types of
 777 state that recurrent, constant-memory mixers can maintain, revealing algorithmic deficiencies in
 778 previous SSM-based models. Merrill et al. (2025) show that under finite precision, practical SSMs
 779 collapse to TC^0 , leading to failures on tasks like permutation composition over S_5 unless the primitive
 780 is extended. Similarly, Yu & Erichson (2025) prove that a single-layer Mamba is not a universal
 781 approximator. Several modifications have been proposed to improve expressivity. For instance,
 782 the same work shows that a block-biased variant regains the universal approximation property with
 783 only minor changes, either through block decomposition or a channel-specific bias. Allowing nega-
 784 tive eigenvalues or non-triangular transitions enables linear RNNs—including diagonal and House-
 785 holder/DeltaNet forms—to capture parity and, under mild assumptions, regular languages (Grazzi
 786 et al., 2025). Complex-valued parameterizations provide another avenue for enhanced expressivity.
 787 Diagonal LTI SSMs demonstrate effectiveness for language modeling (Gu et al., 2022b; Orvieto
 788 et al., 2023), with complex variants achieving equivalent functions using smaller, well-conditioned
 789 parameters (Ran-Milo et al., 2024). However, the introduction of selectivity—the central innovation
 790 of modern SSMs (Gu & Dao, 2024)—narrowed the performance gap with Transformers by enabling
 791 input-dependent dynamics and achieving state-of-the-art results on language modeling benchmarks,
 792 leading practitioners to abandon complex states in favor of simpler real-valued architectures. We
 793 extend this line of work by reintroducing complex-valued state evolution that yields a real SSM with
 794 doubled dimensionality and block-diagonal rotations applied to the update rule—analogous through
 795 SSD (Dao & Gu, 2024) to how RoPE (Su et al., 2023) applies complex rotations to queries and
 796 keys in attention. The resulting data-dependent rotational structure expands stable dynamics to in-
 797 clude oscillatory modes, enabling richer states while maintaining constant memory and linear-time
 798 complexity.

799 B TRAPEZOIDAL DISCRETIZATION

800 **Proposition 5** (Variation of Constants (Tenenbaum & Pollard, 1985)). *Consider the linear SSM*

$$\dot{\mathbf{h}}(t) = A(t) \mathbf{h}(t) + \mathbf{B}(t) x(t),$$

801 where $\mathbf{h}(t) \in \mathbb{R}^N$, $A(t) \in \mathbb{R}$ is a scalar decay, and $\mathbf{B}(t)x(t) \in \mathbb{R}^N$. For Δ_t discretized time grid
 802 $\tau_t = \tau_{t-1} + \Delta_t$, the hidden state satisfies

$$\mathbf{h}_t \approx e^{\Delta_t A_t} \mathbf{h}_{t-1} + \int_{\tau_{t-1}}^{\tau_t} e^{(\tau_t - \tau) A_t} \mathbf{B}(\tau) x(\tau) d\tau. \quad (10)$$

803 *Proof.* Since $A(t)$ is scalar, the homogeneous system $\dot{\mathbf{h}}(t) = A(t)\mathbf{h}(t)$ has solution

$$\mathbf{h}(t) = \phi(t, s) \mathbf{h}(s), \quad \phi(t, s) = \exp\left(\int_s^t A(\xi) d\xi\right).$$

810 The Variation of Constants formula gives us,
 811

$$812 \quad \mathbf{h}(t) = \phi(t, s) \mathbf{h}(s) + \int_s^t \phi(t, \tau) \mathbf{B}(\tau) x(\tau) d\tau.$$

814 Setting $(s, t) = (t_{k-1}, t_k)$ yields the exact \mathbf{h}_t given \mathbf{h}_{t-1} . We approximate $\int_s^t A(\xi) d\xi$ by setting
 815 $A(\tau) \approx A_k$ over $[t_{k-1}, t_k]$, which gives us,
 816

$$817 \quad \phi(t_k, t_{k-1}) = \exp \left(\int_s^t A(\xi) d\xi \right) \approx \exp \left(\int_s^t A_k d\xi \right) = e^{\Delta_k A_k},$$

819 Substituting these approximations in the Variation of Constants integral, we get the approximation
 820

$$821 \quad \mathbf{h}_t \approx e^{\Delta_t A_t} \mathbf{h}_{t-1} + \int_{\tau_{t-1}}^{\tau_t} e^{(\tau_t - \tau) A_t} \mathbf{B}(\tau) x(\tau) d\tau.$$

□

825 B.1 TRAPEZOID DISCRETIZATION'S MASK MATRIX

826 *Proof.* When viewing the tensor contraction form, let us call $C = (T, N), B = (S, N), L =$
 827 $(T, S), X = (S, P)$ based on the Mamba-2 paper. With this decomposition of our mask, we can
 828 view $L = \text{contract}(TZ, ZS \rightarrow TS)(L_1, L_2)$.

829 The original contraction can be seen as
 830

$$\text{contract}(TN, SN, TS, SP \rightarrow TP)(C, B, L, X)$$

832 We can now view it as

$$833 \quad \text{contract}(TN, SN, TJ, JS, SP \rightarrow TP)(C, B, L_1, L_2, X)$$

834 This can be broken into the following:
 835

$$\begin{aligned} 836 \quad Z &= \text{contract}(SN, SP \rightarrow SNP)(B, X) \\ 837 \quad Z' &= \text{contract}(JS, SNP \rightarrow JNP)(L_2, Z) \\ 838 \quad H &= \text{contract}(TJ, JNP \rightarrow TNP)(L_1, Z') \\ 839 \quad Y &= \text{contract}(TN, TNP \rightarrow TP)(C, H) \end{aligned}$$

840 Thus, we can view this step: $\text{contract}(ZS, SNP \rightarrow ZNP)(L_2, Z)$ as a conv of size two applied on
 841 Bx with the traditional SSD $L = L_1$ matrix. □
 842

843 B.2 TRAPEZOIDAL DISCRETIZATION ERROR RATE

844 **Standard assumptions.** We assume that: $A(t), \mathbf{B}(t), x(t)$ are bounded and C^2 on each timestep,
 845 so that $g(\tau)$ has two bounded derivatives; the map $\mathbf{h} \mapsto A(t)\mathbf{h} + \mathbf{B}(t)x(t)$ is Lipschitz in \mathbf{h} which
 846 is true for linear systems; λ_t lies in a bounded interval so that the update is zero-stable.
 847

848 *Proof.* Let $g(\tau) := e^{(t_k - \tau) A_k} \mathbf{B}(\tau) x(\tau)$ denote the integrand in the second term of Proposition 5.
 849 Since $A(t), \mathbf{B}(t), x(t)$ are C^2 on $[t_{k-1}, t_k]$, the function g has two bounded derivatives. A second-
 850 order Taylor expansion of g around t_{k-1} gives us,
 851

$$852 \quad \int_{t_{k-1}}^{t_k} g(\tau) d\tau = \Delta_t g(t_{k-1}) + \frac{\Delta_t^2}{2} g'(t_{k-1}) + \frac{\Delta_t^3}{6} g''(t_{k-1}) + O(\Delta_t^4).$$

854 Recall that the trapezoidal approximation to this integral is given by,
 855

$$856 \quad Q_\lambda = \Delta_t \left[(1 - \lambda_t) g(t_{k-1}) + \lambda_t g(t_k) \right].$$

858 Expanding $g(t_k)$ using Taylor expansion: $g(t_k) = g(t_{k-1}) + \Delta_t g'(t_{k-1}) + \frac{\Delta_t^2}{2} g''(t_{k-1}) + O(\Delta_t^3)$.
 859 Substituting this into Q_λ ,
 860

$$\begin{aligned} 861 \quad Q_\lambda &= \Delta_t \left[(1 - \lambda_t) g(t_{k-1}) + \lambda_t g(t_k) \right] \\ 862 \\ 863 \quad &= \Delta_t g(t_{k-1}) + \lambda_t \Delta_t^2 g'(t_{k-1}) + \lambda_t \frac{\Delta_t^3}{2} g''(t_{k-1}) + O(\Delta_t^4). \end{aligned}$$

864 Hence, the error is given by:
 865

$$866 \int_{t_{k-1}}^{t_k} g(\tau) d\tau - Q_\lambda = \left(\frac{1}{2} - \lambda_t\right) \Delta_t^2 g'(t_{k-1}) + \left(\frac{1}{6} - \frac{\lambda_t}{2}\right) \Delta_t^3 g''(t_{k-1}) + O(\Delta_t^4).$$

868 Under the assumption that $\lambda_t = \frac{1}{2} + c_t \Delta_t$, where $c_t = O(1)$, then $\frac{1}{2} - \lambda_t = -c_t \Delta_t = O(\Delta_t)$ and
 869 thus the Δ_t^2 term is $O(\Delta_t^3)$. Therefore,
 870

$$871 \int_{t_{k-1}}^{t_k} g(\tau) d\tau - Q_\lambda = O(\Delta_t^3),$$

873 which yields an $O(\Delta_t^3)$ local truncation error. Since the update $\mathbf{h}_k = e^{\Delta_t A_k} \mathbf{h}_{k-1} + Q_\lambda$ is linear
 874 and zero-stable for bounded λ_t , standard numerical ODE results imply an $O(\Delta_t^2)$ global error. \square
 875

876 B.3 TRAPEZOIDAL PARAMETERIZATION

878 Parameterization	879 Form of λ_t	880 ppl ↓
881 Default	$\sigma(u_t)$	15.72
882 Fixed 1/2	$\frac{1}{2}$	15.76
883 No trapezoid (Euler)	1	15.81

884 Table 5: **Ablations on λ_t parameterization in the trapezoidal update.**

885 **Setting:** All runs use the Mamba-3 (SISO) 440M model trained at Chinchilla scale, with the other
 886 architectural and optimization hyperparameters being the same as in Table 1.

887 The default model uses a data-dependent gate $\lambda_t = \sigma(u_t)$, where u_t is a learned projection of the
 888 current input token. In Table 5, we try different parameterizations for λ_t and find that the default pa-
 889 rameterization empirically performs the best. Hence we choose the simpler default parameterization
 890 that does *not* enforce the $O(\frac{1}{2} + \Delta_t)$.
 891

892 C COMPLEX SSM PROOFS

893 C.1 PROOF OF PROPOSITION 2

894 **Proposition 2** (Complex-to-Real SSM Equivalence). *Consider a complex-valued SSM*

$$895 \begin{aligned} \dot{\mathbf{h}}(t) &= \text{Diag}(A(t) + i\boldsymbol{\theta}(t)) \mathbf{h}(t) + (\mathbf{B}(t) + i\hat{\mathbf{B}}(t)) x(t), \\ 896 y(t) &= \text{Re}\left((\mathbf{C}(t) + i\hat{\mathbf{C}}(t))^\top \mathbf{h}(t)\right), \end{aligned} \quad (6)$$

897 where $\mathbf{h}(t) \in \mathbb{C}^{N/2}$, $\boldsymbol{\theta}(t), \mathbf{B}(t), \hat{\mathbf{B}}(t), \mathbf{C}(t), \hat{\mathbf{C}}(t) \in \mathbb{R}^{N/2}$, and $x(t), A(t) \in \mathbb{R}$. Under Euler
 898 discretization, this system is equivalent to a real-valued SSM

$$899 \mathbf{h}_t = e^{\Delta_t A_t} \mathbf{R}_t \mathbf{h}_{t-1} + \Delta_t \mathbf{B}_t x_t, \quad (7)$$

$$900 y_t = \mathbf{C}_t^\top \mathbf{h}_t,$$

901 with state $\mathbf{h}_t \in \mathbb{R}^N$, projections

$$902 \mathbf{B}_t = \begin{bmatrix} \mathbf{B}_t \\ \hat{\mathbf{B}}_t \end{bmatrix} \in \mathbb{R}^N, \quad \mathbf{C}_t = \begin{bmatrix} \mathbf{C}_t \\ -\hat{\mathbf{C}}_t \end{bmatrix} \in \mathbb{R}^N,$$

903 and a transition matrix

$$904 \mathbf{R}_t = \text{Block}\left(\{R(\Delta_t \boldsymbol{\theta}_t[i])\}_{i=1}^{N/2}\right) \in \mathbb{R}^{N \times N}, \quad R(\Theta) = \begin{bmatrix} \cos(\Theta) & -\sin(\Theta) \\ \sin(\Theta) & \cos(\Theta) \end{bmatrix}.$$

905 *Proof.* We first present the derivation for $N = 2$; the block-diagonal structure for general even N
 906 follows by grouping pairs of coordinates.

907 Let $h_t + i\hat{h}_t$ denote the complexified hidden state, with parameters $A(t) + i\theta(t)$ and $B(t) + i\hat{B}(t)$ for
 908 the transition and input, respectively. By the variation of constants formula (Proposition 5), applying
 909 zero-order hold and Euler’s rule over a step $[t_{k-1}, t_k]$ gives
 910

$$911 h_k + i\hat{h}_k = e^{\Delta_t (A_t + i\theta_t)} (h_{k-1} + i\hat{h}_{k-1}) + \Delta_t (B_t + i\hat{B}_t) x_t.$$

918 Expanding the exponential,
 919

$$920 \quad e^{\Delta_t(A_t + i\theta_t)} = e^{\Delta_t A_t} \left(\cos(\Delta_t \theta_t) + i \sin(\Delta_t \theta_t) \right),$$

922 so in real coordinates $\mathbf{h}_t = \begin{bmatrix} h_t \\ \hat{h}_t \end{bmatrix} \in \mathbb{R}^2$ the recurrence becomes
 923
 924

$$925 \quad \mathbf{h}_t = e^{\Delta_t A_t} \underbrace{\begin{bmatrix} \cos(\Delta_t \theta_t) & -\sin(\Delta_t \theta_t) \\ \sin(\Delta_t \theta_t) & \cos(\Delta_t \theta_t) \end{bmatrix}}_{R(\Delta_t \theta_t)} \mathbf{h}_{t-1} + \Delta_t \begin{bmatrix} B_t \\ \hat{B}_t \end{bmatrix} x_t.$$

929 Stacking across $N/2$ such pairs yields the block-diagonal transition
 930

$$931 \quad \mathbf{h}_t = e^{\Delta_t A_t} \text{Block}(\{R(\Delta_t \theta_t[i])\}_{i=1}^{N/2}) \mathbf{h}_{t-1} + \Delta_t \begin{bmatrix} \mathbf{B}_t \\ \hat{\mathbf{B}}_t \end{bmatrix} x_t.$$

934 For the output,
 935

$$936 \quad y_t = \text{Re}((\mathbf{C}_t + i\hat{\mathbf{C}}_t)^\top (h_t + i\hat{h}_t)) = \begin{bmatrix} \mathbf{C}_t \\ -\hat{\mathbf{C}}_t \end{bmatrix}^\top \mathbf{h}_t,$$

938 which defines the real projection $\mathbf{C}_t \in \mathbb{R}^N$ in the proposition. This proves the equivalence between
 939 complex SSM and the real block-diagonal system with rotations. \square
 940

941 C.2 PROOF OF PROPOSITION 3

942 **Proposition 3** (Complex SSM, Data-Dependent RoPE Equivalence). *Under the notation established
 943 in Proposition 2, consider the real SSM defined in Eq. 7 unrolled for T time-steps. The output of
 944 the above SSM is equivalent to that of a vanilla scalar transition matrix-based SSM (Eq. 2) with a
 945 data-dependent rotary embedding applied on the \mathbf{B}, \mathbf{C} components of the SSM defined as:*

$$946 \quad \mathbf{h}_t = e^{\Delta_t A_t} \mathbf{h}_{t-1} + \left(\prod_{i=0}^t \mathbf{R}_i^\top \right) \mathbf{B}_t x_t, \quad \mathbf{y}_t = \left(\left(\prod_{i=0}^t \mathbf{R}_i^\top \right) \mathbf{C}_t \right)^\top \mathbf{h}_t \quad (8)$$

950 where the matrix production represents right matrix multiplication, e.g., $\prod_{i=0}^1 \mathbf{R}_i = \mathbf{R}_0 \mathbf{R}_1$. We
 951 denote employing the vanilla SSM to compute the Complex SSM as “RoPE trick”.

952
 953 *Proof.* Consider the SSM

$$954 \quad \mathbf{h}_t = e^{\Delta_t A_t} \mathbf{R}_t \mathbf{h}_{t-1} + \mathbf{B}_t x_t, \quad \mathbf{y}_t = \mathbf{C}_t^\top \mathbf{h}_t, \quad (11)$$

956 where (as in Proposition 3) $A_t \in \mathbb{R}$ is a scalar (so that $e^{\Delta_t A_t}$ is a scalar and commutes with rotations), and \mathbf{R}_t is block-diagonal orthogonal/unitary, hence $\mathbf{R}_t^{-1} = \mathbf{R}_t^\top$.
 957

958 Unrolling the recurrence with the convention that an empty product is the identity,
 959

$$960 \quad \mathbf{h}_t = \sum_{i=0}^t \left(\prod_{s=i+1}^t e^{\Delta_s A_s} \mathbf{R}_s \right) \mathbf{B}_i x_i. \quad (12)$$

963 Thus
 964

$$965 \quad \mathbf{y}_t = \mathbf{C}_t^\top \mathbf{h}_t = \sum_{i=0}^t \mathbf{C}_t^\top \left(\prod_{s=i+1}^t e^{\Delta_s A_s} \mathbf{R}_s \right) \mathbf{B}_i x_i. \quad (13)$$

968 Using unitarity property,
 969

$$970 \quad \prod_{s=i+1}^t \mathbf{R}_s = \left(\prod_{s=0}^t \mathbf{R}_s \right) \left(\prod_{s=0}^i \mathbf{R}_s \right)^{-1} = \left(\prod_{s=0}^t \mathbf{R}_s \right) \left(\prod_{s=0}^i \mathbf{R}_s^\top \right).$$

972 Since $e^{\Delta_s A_s}$ are scalars, they commute with rotations; hence
973

$$974 \quad \mathbf{y}_t = \sum_{i=0}^t \mathbf{C}_t^\top \left(\prod_{s=0}^t \mathbf{R}_s \right) \left(\prod_{s=i+1}^t e^{\Delta_s A_s} \right) \left(\prod_{s=0}^i \mathbf{R}_s^\top \right) \mathbf{B}_i x_i \quad (14)$$

$$977 \quad = \left(\left(\prod_{s=0}^t \mathbf{R}_s^\top \right) \mathbf{C}_t \right)^\top \sum_{i=0}^t \left(\prod_{s=i+1}^t e^{\Delta_s A_s} \right) \left(\prod_{s=0}^i \mathbf{R}_s^\top \right) \mathbf{B}_i x_i. \quad (15)$$

980 Define the rotated parameters $\bar{\mathbf{C}}_t := (\prod_{s=0}^t \mathbf{R}_s^\top) \mathbf{C}_t$ and $\bar{\mathbf{B}}_i := (\prod_{s=0}^i \mathbf{R}_s^\top) \mathbf{B}_i$. Then
981

$$982 \quad \mathbf{y}_t = \bar{\mathbf{C}}_t^\top \sum_{i=0}^t \left(\prod_{s=i+1}^t e^{\Delta_s A_s} \right) \bar{\mathbf{B}}_i x_i. \quad (16)$$

985 Equivalently, introducing the rotated state $\tilde{\mathbf{h}}_t := (\prod_{s=0}^t \mathbf{R}_s^\top) \mathbf{h}_t$,

$$986 \quad \tilde{\mathbf{h}}_t = e^{\Delta_t A_t} \tilde{\mathbf{h}}_{t-1} + \bar{\mathbf{B}}_t x_t, \quad \mathbf{y}_t = \bar{\mathbf{C}}_t^\top \tilde{\mathbf{h}}_t, \quad (17)$$

988 \square

989 C.3 PROOF OF PROPOSITION 4

990 **Proposition 4** (Rotary Embedding Equivalence with Trapezoidal Discretization). *Discretizing a
991 complex SSM with the trapezoidal rule (Proposition 1) yields the recurrence*

$$993 \quad \mathbf{h}_t = \alpha_t \mathbf{h}_{t-1} + \beta_t \left(\prod_{i=0}^{t-1} \mathbf{R}_i^\top \right) \mathbf{B}_{t-1} x_{t-1} + \gamma_t \left(\prod_{i=0}^t \mathbf{R}_i^\top \right) \mathbf{B}_t x_t,$$

$$996 \quad \mathbf{y}_t = \left(\left(\prod_{i=0}^t \mathbf{R}_i^\top \right) \mathbf{C}_t \right)^\top \mathbf{h}_t. \quad (9)$$

999 Here \mathbf{R}_t is the block-diagonal rotation matrix defined in Proposition 3.

1000 *Proof.* We begin from the complex SSM (as in Prop. 2)

$$1002 \quad \dot{\mathbf{h}}(t) = \text{Diag}(\mathbf{A}(t) + i\boldsymbol{\theta}(t)) \mathbf{h}(t) + (\mathbf{B}(t) + i\hat{\mathbf{B}}(t)) x(t),$$

$$1004 \quad y(t) = \text{Re}((\mathbf{C}(t) + i\hat{\mathbf{C}}(t))^\top \mathbf{h}(t)),$$

1006 where $\mathbf{A}(t) \in \mathbb{R}$ is a scalar and $\boldsymbol{\theta}(t), \mathbf{B}(t), \hat{\mathbf{B}}(t), \mathbf{C}(t), \hat{\mathbf{C}}(t) \in \mathbb{R}^{N/2}$.

1007 Recall from Prop. 5,

$$1009 \quad \mathbf{h}_t \approx e^{\Delta_t(A_t + i\boldsymbol{\theta}_t)} \mathbf{h}_{t-1} + \int_{\tau_{t-1}}^{\tau_t} e^{(\tau_t - \tau)(A_t + i\boldsymbol{\theta}_t)} (\mathbf{B}(\tau) + i\hat{\mathbf{B}}(\tau)) x(\tau) d\tau.$$

1012 Applying Prop. 1 to the above integral, we get

$$1013 \quad \mathbf{h}_t = e^{\Delta_t(A_t + i\boldsymbol{\theta}_t)} \mathbf{h}_{t-1} + \beta_t e^{i\Delta_t \boldsymbol{\theta}_t} (\mathbf{B}_{t-1} + i\hat{\mathbf{B}}_{t-1}) x_{t-1} + \gamma_t (\mathbf{B}_t + i\hat{\mathbf{B}}_t) x_t, \quad (18)$$

1014 where m

$$1015 \quad \alpha_t := e^{\Delta_t A_t}, \quad \beta_t := (1 - \lambda_t) \Delta_t e^{\Delta_t A_t}, \quad \gamma_t := \lambda_t \Delta_t,$$

1017 Since $e^{\Delta_t(A_t + i\boldsymbol{\theta}_t)} = \alpha_t e^{i\Delta_t \boldsymbol{\theta}_t}$ and as shown in Prop. 2, multiplication by $e^{i\Delta_t \boldsymbol{\theta}_t}$ is a block-diagonal
1018 rotation in real coordinates, we get the real N -dimensional recurrence
1019

$$1020 \quad \mathbf{h}_t = \alpha_t \mathbf{R}_t \mathbf{h}_{t-1} + \beta_t \mathbf{R}_t \mathbf{B}_{t-1} x_{t-1} + \gamma_t \mathbf{B}_t x_t, \quad (19)$$

$$1021 \quad \mathbf{y}_t = \mathbf{C}_t^\top \mathbf{h}_t,$$

1023 where $\mathbf{R}_t = \text{Block}(\{R(\Delta_t \boldsymbol{\theta}_t[i])\}_{i=1}^{N/2})$ where $R(\Theta) = \begin{bmatrix} \cos \Theta & -\sin \Theta \\ \sin \Theta & \cos \Theta \end{bmatrix}$, and projections
1024

1025 $\mathbf{B}_t = \begin{bmatrix} \mathbf{B}_t \\ \hat{\mathbf{B}}_t \end{bmatrix}$, $\mathbf{C}_t = \begin{bmatrix} \mathbf{C}_t \\ -\hat{\mathbf{C}}_t \end{bmatrix}$. Note that \mathbf{R}_t is orthogonal, so $\mathbf{R}_t^{-1} = \mathbf{R}_t^\top$.




Figure 4: Contrasting Mamba-2 and Mamba-3 Architectures: Key updates include trapezoidal discretization, data-dependent RoPE embeddings, MIMO projections, QK normalization, and learnable biases.

We define the following,

$$\tilde{\mathbf{h}}_t := \left(\prod_{s=0}^t \mathbf{R}_s^\top \right) \mathbf{h}_t, \quad \bar{\mathbf{B}}_t := \left(\prod_{s=0}^t \mathbf{R}_s^\top \right) \mathbf{B}_t, \quad \bar{\mathbf{C}}_t := \left(\prod_{s=0}^t \mathbf{R}_s^\top \right) \mathbf{C}_t.$$

Left-multiplying equation 19 by $\prod_{s=0}^t \mathbf{R}_s^\top$ and using $\mathbf{R}_t^\top \mathbf{R}_t = I$,

$$\begin{aligned} \tilde{\mathbf{h}}_t &= \alpha_t \tilde{\mathbf{h}}_{t-1} + \beta_t \bar{\mathbf{B}}_{t-1} x_{t-1} + \gamma_t \bar{\mathbf{B}}_t x_t, \\ y_t &= \bar{\mathbf{C}}_t^\top \tilde{\mathbf{h}}_t. \end{aligned}$$

This is a vanilla scalar-transition SSM with data-dependent rotary embeddings absorbed into \mathbf{B}, \mathbf{C} via cumulative products of \mathbf{R}_s^\top . \square

D MIMO FOR MAMBA-3

With hindsight from Mamba and with inference in mind, we propose the following MIMO formulation:

Mamba with MIMO. With a given batch, head, and sequence position t , consider the input $\mathbf{U}_t \in \mathbb{R}^D$. Also denote $P, R \in \mathbb{N}$ as the head dimension and MIMO rank, respectively. We first obtain SSM parameters via a set of projections defined in terms of tensor contraction notation as follows:

$$\begin{aligned} \mathbf{B}_t &= \text{contract}(DNR, D \rightarrow NR)(\mathbf{W}_B, \mathbf{U}_t) & \mathbf{C}_t &= \text{contract}(DNR, D \rightarrow NR)(\mathbf{W}_C, \mathbf{U}_t), \\ \mathbf{X}'_t &= \text{contract}(PD, D \rightarrow P)(\mathbf{W}_{X'}, \mathbf{U}_t) & \mathbf{X}_t &= \text{contract}(PR, P \rightarrow PR)(\mathbf{W}_X, \mathbf{X}'_t), \end{aligned}$$

where $\mathbf{W}_B, \mathbf{W}_C, \mathbf{W}_{X'}, \mathbf{W}_X$ are model parameters. Additionally, we obtain the residual term \mathbf{Z}_t in the same manner as \mathbf{X}_t with weights $\mathbf{W}_{Z'}$ and \mathbf{W}_Z . The state update and the SSM output is then computed via the following MIMO SSM:

$$\mathbf{H}_t = a_t \mathbf{H}_{t-1} + \mathbf{B}_t \mathbf{X}_t^\top \in \mathbb{R}^{N \times P}, \quad \mathbf{Y}_t = \mathbf{H}_t^\top \mathbf{C}_t \in \mathbb{R}^{P \times R}.$$

The intermediate output \mathbf{Y}'_t is obtained via some residual function ϕ , $\mathbf{Y}'_t \leftarrow \phi(\mathbf{Y}_t, \mathbf{Z}_t)$. Finally, the layer output $\mathbf{O}_t \in \mathbb{R}^D$ is computed via the following down projections:

$$\mathbf{O}'_t = \text{contract}(PR, R \rightarrow P)(\mathbf{W}_{O'}, \mathbf{Y}'_t) \quad \mathbf{O}_t = \text{contract}(P, PD \rightarrow D)(\mathbf{W}_O, \mathbf{O}'_t).$$

1080 This formulation enhances the existing Mamba3 architecture by providing a lightweight parameterization that transforms the set of independent SISO SSMs within each head into a set of MIMO SSMs. Here, we note that the hardware-efficient chunking technique employed by Mamba2 for pre-training can be applied with little change, as the MIMO dimension r is orthogonal to the sequence dimension.
 1081
 1082
 1083
 1084
 1085

E EXPERIMENTAL DETAILS

1086 **Language Modeling.** Our pretraining procedures follow that of Dao & Gu (2024)'s section D.2.
 1087 All models at each scale follow the same procedure and were trained with bfloat16. The Mamba
 1088 family of models were trained using the standard expand factor of 2 and a dstate of 128 and head
 1089 dimension of 64. The Transformer baselines follows Dao & Gu (2024), and the Gated DeltaNet
 1090 baselines follow (Yang et al., 2025a). We utilize the Llama-3.1 tokenizer (Grattafiori et al., 2024)
 1091 for all models.
 1092
 1093

1094 We utilize LM Evaluation Harness (Gao et al., 2024) to test the zero-shot language modeling ca-
 1095 pabilities of our pretrained model on LAMBADA (OpenAI version) (Paperno et al., 2016), Hel-
 1096 laSwag (Zellers et al., 2019), PIQA (Bisk et al., 2019), Arc-Easy/Arc-Challenge (Clark et al., 2018),
 1097 WinoGrande (Sakaguchi et al., 2019), and OpenBookQA (Mihaylov et al., 2018).
 1098
 1099

1100 **Real-World and Synthetic Retrieval.** For our real-world retrieval tasks, we evaluate on the com-
 1101 mon suite consisting of SWDE (Arora et al., 2025b), SQuAD (Rajpurkar et al., 2018), FDA (Arora
 1102 et al., 2025b), TriviaQA (Joshi et al., 2017), NQ (Kwiatkowski et al., 2019), and DROP (Dua et al.,
 1103 2019). We utilize the cloze-formatted version of the aforementioned tasks provided by Arora et al.
 1104 (2025b; 2024), as the original datasets are in a question-answering format, making it challenge for
 1105 solely pretrained models. All tasks were truncated to match the training context length. The syn-
 1106 thetic NIAH tasks (Hsieh et al., 2024) were also run with LM Evaluation Harness.
 1107
 1108

1109 **State-Tracking Synthetics.** Training follows a sequence length curriculum that progresses from 3
 1110 -40 to 160, evaluated at 256. Each curriculum runs for 10^4 steps with batch size 256. We use 1 layer
 1111 models for Parity and 3 layer models for Modular-arithmetic tasks. The state size is chosen to be
 1112 64, and we sweep $d_{\text{model}} \in \{32, 64\}$ and 8 learning rates logarithmically spaced between 10^{-4} and
 1113 10^{-2} , reporting the best validation accuracy.
 1114

F ADDITIONAL EXPERIMENTAL RESULTS




Figure 5: Pretrained 1.5B models' performance on the held-out FineWeb-Edu test set at varying context lengths. Mamba-3 exhibits strong length extrapolation while Mamba-2 falters at longer contexts.

1134 Table 6: Downstream language modeling evaluations on parameter-matched pretrained models, in-
 1135 cluding Mamba-3 MIMO. Mamba-3 MIMO’s average accuracy on all tasks is more than 1 percent-
 1136 age point better than the next best (Mamba-3 SISO).

Model	FW-Edu ppl ↓	LAMB. ppl ↓	LAMB. acc ↑	HellaS. acc_n ↑	PIQA acc ↑	Arc-E acc ↑	Arc-C acc_n ↑	WinoGr. acc ↑	OBQA acc ↑	Average acc ↑
Transformer-440M	13.03	21.2	41.7	50.5	69.9	67.6	34.6	56.7	26.0	49.6
Gated DeltaNet-440M	13.12	<u>19.0</u>	<u>40.4</u>	50.5	70.5	67.5	34.0	55.3	25.8	49.1
Mamba-2-440M	13.00	19.6	40.8	<u>51.7</u>	70.6	68.8	35.0	54.1	26.0	49.6
Mamba-3-440M	<u>12.87</u>	19.6	40.2	<u>51.7</u>	71.9	<u>68.9</u>	<u>34.4</u>	55.8	<u>26.0</u>	<u>49.8</u>
Mamba-3-MIMO-440M	12.72	17.1	43.4	52.8	<u>70.8</u>	69.6	35.6	<u>56.3</u>	28.4	51.0
Transformer-880M	11.42	15.0	44.7	57.2	72.6	71.6	39.2	57.7	26.8	52.8
Gated DeltaNet-880M	11.39	12.7	47.1	57.5	72.6	72.5	38.8	57.9	30.6	53.9
Mamba-2-880M	11.35	13.8	45.0	58.1	72.5	72.3	38.7	56.8	30.2	53.4
Mamba-3-880M	<u>11.23</u>	12.9	<u>47.2</u>	58.8	<u>73.6</u>	<u>72.7</u>	<u>40.2</u>	<u>58.4</u>	30.0	54.4
Mamba-3-MIMO-880M	<u>11.11</u>	11.8	49.5	<u>59.2</u>	<u>73.7</u>	<u>74.7</u>	<u>41.2</u>	59.9	28.6	55.3




Figure 6: Mamba-3 demonstrates superior performance compared to strong baselines like Mamba-2, Llama, and Gated Deltanet. These are 440M models, trained and evaluated on FineWeb-Edu.

We also compare the effectiveness of state size usage of Mamba variants to a Gated DeltaNet baseline in Figure 7. We highlight the difficulty of directly comparing GDN versus Mamba-style models due to the differing head structure, multi-head compared to multi-value respectively. Our experiments hold GDN’s v_expand to 2 and decrease the head dimension accordingly to vary the relative total state size. Similar to Figure 3, we train 440M models to $2 \times$ Chinchilla tokens and sweep across $d_{\text{state}} = \{32, 64, 128\}$ for the Mamba models and $d_{\text{head dim}} = \{32, 64, 128\}$ for GDN. We parameter match all models.




Figure 7: Exploration of state size (inference speed proxy) versus pretraining perplexity (performance proxy). Mamba-3 and Mamba-3 MIMO continue set the Pareto frontier.

G ARCHITECTURE ABLATIONS

We explore our model architecture’s ablation in this section. All models are trained at the 440M scale to Chinchilla optimal number of tokens ($20 \times$ tokens to parameters) with the same experimental procedures as our pretrained models as covered in Appendix E unless otherwise stated.

B,C Bias Parameterization. The Mamba-3 model’s separate B and C biases are head-specific and channel-wise and added to both B and C after the QK-Norm. While the biases in the final Mamba-3 model are trainable, data-independent parameters and initialized to all ones, we explore various bias parameterizations in Table 7a. We find our models are not very sensitive to the initialization of the biases as long as they are positive. We choose the all-ones initialization due to its simplicity.

We also explore the impact removing the B or C bias on performance in Table 7b (bias is initialized with our default parameterization when utilized). Unlike in Yu & Erichson (2025), which finds that B bias by itself is able to improve performance on Mamba-1, our experiments find that only having B bias hurts performance slightly and that B and C biases have synergistic properties.

Bias Init.	Trainable	ppl \downarrow	B Bias	C Bias	ppl \downarrow
1.0	✓	15.72	✗	✗	16.52
0.0	✓	16.57	✓	✗	16.68
1.0	✗	15.80	✗	✓	15.98
$\mathcal{U}(0, 1)$	✓	15.76	✓	✓	15.69
$\mathcal{U}(-1, 1)$	✓	16.07			

(a) Effect of parameterization of the B and C bias on model performance, measured by pretraining perplexity. We find our default initialization of all ones (first row) provides the best performance, but performance is not sensitive as long as biases are positive.

(b) Applying a bias to both B and C leads to the best performance. Only applying B bias (Block-Biased (Yu & Erichson, 2025) Mamba-3 variant) does not provide significant gains over the no-bias baseline.

Table 7: Ablations on B , C bias initialization (left) and presence (right) for Mamba-3.

H INFERENCE KERNEL LATENCY ANALYSIS

H.1 KERNEL IMPLEMENTATIONS AND FUSION STRUCTURE

In Table 3, we detail the DSL (Triton, CuTe, PyTorch) and the fusion level of the kernels used in our latency analysis. For Mamba-2 and Gated DeltaNet (GDN), we directly use the publicly released Triton kernels from the respective authors. For Mamba-3, we implement new inference kernels with a comparable fusion structure: the forward uses a Triton kernel fused with rotary position embeddings, while the decode path uses a CuTe kernel fused with gating and MIMO projection.

In Tables 8 and 9, we abbreviate IP = input projection, Conv = 1D convolution, Gate = gating, OP = output projection. Colors indicate implementation backend (Torch, Triton, CuTe).

1242
1243
1244 Table 8: Kernel DSL and fusion structure for **forward** (prefill) kernels.
1245
1246
1247
1248
1249

Model (Forward)	Kernel DSL	Fusion Level
Mamba-2	Triton	IP, Conv , SSM , Gate+OP
Gated DeltaNet	Triton	IP, Conv , Chunked Delta , Gate+OP
Mamba-3 (SISO)	Triton	IP, SSM+Rotary , Gate+OP
Mamba-3 (MIMO)	Triton	IP, SSM+Rotary , Gate+OP

1250
1251 Table 9: Kernel DSL and fusion structure for **decode** kernels.
1252

Model (Decode)	Kernel DSL	Fusion Level
Mamba-2	Triton	IP, Conv , SSM , Gate+OP
Gated DeltaNet	Triton	IP, Conv , Recurrent Delta , Gate+OP
Mamba-3 (SISO)	CuTe + Triton	IP, Rotary , SSM+Gate+OP
Mamba-3 (MIMO)	CuTe + Triton	IP, Rotary , SSM+Gate+OP+MIMO

1260
1261 H.2 EXTENDED PREFILL AND PREFILL+DECODE LATENCY MEASUREMENTS1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
22010
22011
22012
22013
22014
22015
22016
22017
22018
22019
22020
22021
22022
22023
22024
22025
22026
22027
22028
22029
22030
22031
22032
22033
22034
22035
22036
22037
22038
22039
22040
22041
22042
22043
22044
22045
22046
22047
22048
22049
22050
22051
22052
22053
22054
22055
22056
22057
22058
22059
22060
22061
22062
22063
22064
22065
22066
22067
22068
22069
22070
22071
22072
22073
22074
22075
22076
22077
22078
22079
22080
22081
22082
22083
22084
22085
22086
22087
22088
22089
22090
22091
22092
22093
22094
22095
22096
22097
22098
22099
220100
220101
220102
220103
220104
220105
220106
220107
220108
220109
220110
220111
220112
220113
220114
220115
220116
220117
220118
220119
220120
220121
220122
220123
220124
220125
220126
220127
220128
220129
220130
220131
220132
220133
220134
220135
220136
220137
220138
220139
220140
220141
220142
220143
220144
220145
220146
220147
220148
220149
220150
220151
220152
220153
220154
220155
220156
220157
220158
220159
220160
220161
220162
220163
220164
220165
220166
220167
220168
220169
220170
220171
220172
220173
220174
220175
220176
220177
220178
220179
220180
220181
220182
220183
220184
220185
220186
220187
220188
220189
220190
220191
220192
220193
220194
220195
220196
220197
220198
220199
220200
220201
220202
220203
220204
220205
220206
220207
220208
220209
220210
220211
220212
220213
220214
220215
220216
220217
220218
220219
220220
220221
220222
220223
220224
220225
220226
220227
220228
220229
220230
220231
220232
220233
220234
220235
220236
220237
220238
220239
220240
220241
220242
220243
220244
220245
220246
220247
220248
220249
220250
220251
220252
220253
220254
220255
220256
220257
220258
220259
220260
220261
220262
220263
220264
220265
220266
220267
220268
220269
220270
220271
220272
220273
220274
220275
220276
220277
220278
220279
220280
220281
220282
220283
220284
220285
220286
220287
220288
220289
220290
220291
220292
220293
220294
220295
220296
220297
220298
220299
220300
220301
220302
220303
220304
220305
220306
220307
220308
220309
220310
220311
220312
220313
220314
220315
220316
220317
220318
220319
220320
220321
220322
220323
220324
220325
220326
220327
220328
220329
220330
220331
220332
220333
220334
220335
220336
220337
220338
220339
220340
220341
220342
220343
220344
220345
220346
220347
220348
220349
220350
220351
220352
220353
220354
220355
220356
220357
220358
220359
220360
220361
220362
220363
220364
220365
220366
220367
220368
220369
220370
220371
220372
220373
220374
220375
220376
220377
220378
220379
220380
220381
220382
220383
220384
220385
220386
220387
220388
220389
220390
220391
220392
220393
220394
220395
220396
220397
220398
220399
220400
220401
220402
220403
220404
220405
220406
220407
220408
220409
220410
220411
220412
220413
220414
220415
220416
220417
220418
220419
220420
220421
220422
220423
220424
220425
220426
220427
220428
220429
220430
220431
220432
220433
220434
220435
220436
220437
220438
220439
220440
220441
220442
220443
220444
220445
220446
220447
220448
220449
220450
220451
220452
220453
220454
220455
220456
220457
220458
220459
220460
220461
220462
220463
220464
220465
220466
220467
220468
220469
220470
220471
220472
220473
220474
220475
220476
220477
220478
220479
220480
220481
220482
220483
220484
220485
220486
220487
220488
220489
220490
220491
220492
220493
220494
220495
220496
220497
220498
220499
220500
220501
220502
220503
220504
220505
220506
220507
220508
220509
220510
220511
220512
220513
220514
220515
220516
220517
220518
220519
220520
220521
220522
220523
220524
220525
220526
220527
220528
220529
220530
220531
220532
220533
220534
220535
220536
220537
220538
220539
220540
220541
220542
220543
220544
220545
220546
220547
220548
220549
220550
220551
220552
220553
220554
220555
220556
220557
220558
220559
220560
220561
220562
220563
220564
220565
220566
220567
220568
220569
220570
220571
220572
220573
220574
220575
220576
220577
220578
220579
220580
2205