Improved State Mixing in Higher-order and Block Diagonal Linear Recurrent Networks

Igor Dubinin [1 2] **Antonio Orvieto** [* 3 4 5] **Felix Effenberger** [* 1 2]

Abstract

Linear recurrent networks (LRNNs) and linear state space models (SSMs) promise computational and memory efficiency on sequence modeling tasks, yet their diagonal state transitions limit expressivity. Dense and/or nonlinear architectures (e.g., LSTMs) on the other hand are provably more expressive, but computationally costly. Here, we explore how expressivity in LRNNs can be increased via richer state mixing across time and channels while maintaining competitive efficiency. Specifically, we introduce two structured LRNN architectures: (i) Higher-order Linear Recurrent Units (H-LRU), which generalize recurrences to arbitrary order, mixing multiple past states, and (ii) Block-Diagonal LRUs (BD-LRU), which enable dense intra-block channel mixing. To ensure stable training, we introduce a selective gate normalization scheme that allows for scalable window and block sizes. To maintain efficiency, we utilize a parallel-scan implementation that keeps the throughput competitive with diagonal LRNNs for moderate orders (H-LRU) and block sizes (BD-LRU). Consistent with prior theoretical studies on the limitations of diagonal models, we empirically demonstrate in both synthetic sequence modeling and language modeling that our architectures significantly benefit from the increased expressivity of structured state mixing. Our results show that the structure of state mixing is a critical driver of performance in LRNNs, offering a practical pathway to closing the efficiency–expressivity gap in linear sequence models.

[1]Natural Intelligence (NISYS), Frankfurt am Main [2]Ernst Strüngmann Institut, Frankfurt am Main [3]ELLIS Institute Tübingen [4]MPI for Intelligent Systems [5]Tübingen AI Center. Correspondence to: Igor Dubinin <igor.dubinin@esi-frankfurt.de>.

Proceedings of the 43rd International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

1. Introduction

Recent studies have highlighted fundamental limitations of linear recurrent networks (LRNNs) by showing that the structure of the state-transition matrix results in a trade-off between efficiency and expressivity (Merrill & Sabharwal, 2023; Cirone et al., 2024; Merrill et al., 2024). Architectures based on diagonal matrices enable an efficient implementation but are inherently limited in expressive power, while dense models are provably more expressive yet computationally prohibitive. To bridge this gap, several LRNN architectures have been proposed: efficient structured architectures such as ones with diagonal-plus-low-rank matrices (Yang et al., 2024a; Peng et al., 2025) and their products (Siems et al., 2025), ones based on approximations of dense matrices at test time (Sun et al., 2024; Movahedi et al., 2025; von Oswald et al., 2025), and other solutions that are *de facto* equivalent to block-diagonal architectures (e.g., oscillatory blocks (Rusch & Rus, 2024) and complex-valued states (Orvieto et al., 2023; De et al., 2024)). Together, these studies suggest that exploring the configuration space between diagonal and dense transition matrices may yield more expressive LRNN models.

When designing block-diagonal recurrences, the immediate issue one faces is that of dynamical stability and forward pass normalization – a crucial element that is well studied and discussed in diagonal LRNNs (Orvieto et al., 2023; Wang & Li, 2023; Zucchet & Orvieto, 2024), yet requires additional care in non-diagonal linear architectures where eigenvalues are not readily available. Traditionally, stability has been ensured by parameterizations that constrain eigenvalues of the transition matrix inside the complex unit disk (Arjovsky et al., 2016; Helfrich et al., 2018), a strategy that effectively mitigates vanishing and exploding gradients. More recently, similar conditions have been applied to derive efficient reparameterizations that ensure stability in diagonal linear recurrent units (Orvieto et al., 2023; De et al., 2024). In both selective and non-selective SSMs (designed in continuous-time), stability is achieved by exponential parametrization, resulting from zero-order-hold discretization techniques (Gu et al., 2021; Gu & Dao, 2023). Finally, in LRNNs with diagonal-plus-low-rank transition matrices, normalization arises naturally from the structure of generalized Householder transformations (Yang et al.,

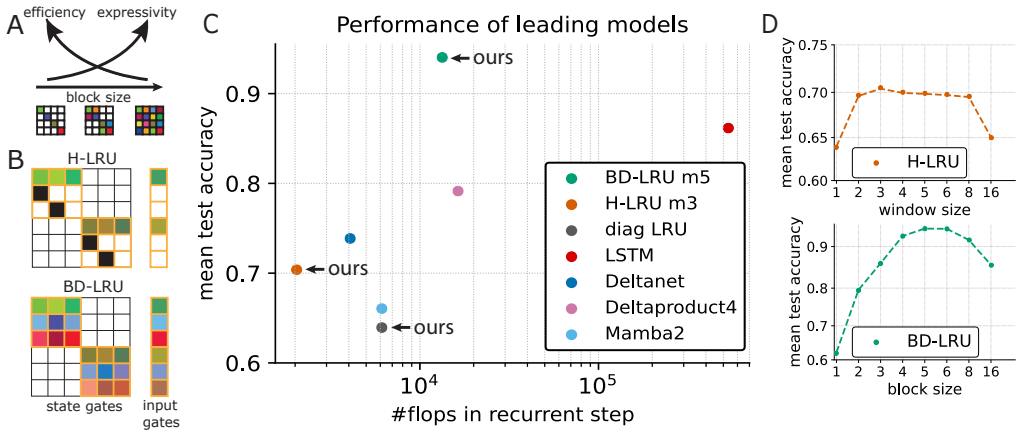

Figure 1. Structure and performance of the proposed H-LRU and BD-LRU architectures. **A.** A schematic illustration of the theoretically predicted trade-off between expressivity and efficiency of block-diagonal linear recurrent networks. **B.** Schematic illustration of the gating mechanisms in block-diagonal form, showing both the state gates that constitute the state-transition matrix and the input gates that act on external inputs. The structure of the gates' selectivity is color-coded: white squares indicate fixed zero gates, black squares indicate fixed identity gates, other colors indicate active selective gates; similar color palettes indicate row-wise normalization. **C.** Summary of the performance of the proposed and the baseline models. The x-axis indicates the number of FLOPs per recurrent step. The y-axis denotes the mean test accuracy over all considered synthetic tasks (compression, selective copying, in-context recall, permutation) of the overall best-performing model configuration (hidden size up to 6k). Optimal hidden sizes vary between models, see also Fig. 7. Note that H-LRU and BD-LRU can achieve better or matching performance than both linear and non-linear baselines while requiring fewer FLOPs per recurrent step. Diagonal LRU reports the best results across both H-LRU m1 and BD-LRU m1, which are identical for $m = 1$. **D.** Best performance for different window sizes m (H-LRU) and block sizes m (BD-LRU).

2024b). Although several recent studies have examined block-diagonal architectures, they either focus on parameterizations of non-selective models (Biegun et al., 2024; Rusch & Rus, 2024), analyze only the stability of the state-transition matrix norm (Fan et al., 2023), or rely on architectures where this matrix is normalized by design (Yang et al., 2024b; Walker et al., 2025), without fully addressing the problem of joint normalization of selective state-transition matrix and selective input gates, which has been previously shown critical for sequence modeling in diagonal LRNNs (Gu & Dao, 2023; De et al., 2024).

Building on this line of work, we explore how to improve expressivity of LRNNs through structured selective state mixing, without compromising their computational efficiency. We note that the goal of this paper is not to introduce a new SOTA, but to isolate and study the impact of structured state mixing on the capabilities of LRNNs. Our contributions are as follows:

- We introduce two architectures with structured selective state mixing: (i) Higher-order Linear Recurrent Units (H-LRU) which generalize recurrences to arbitrary order, mixing multiple past states, and (ii) Block-Diagonal LRUs (BD-LRU), which enable dense intra-block channel mixing.
- We present empirical results on synthetic sequence modeling and language modeling tasks showing that state mixing is a critical driver of performance in LRNNs.
- We propose a joint selective gate normalization that ensures stable scaling with respect to block and window sizes while preserving input-dependence.
- We provide a parallel-scan implementation that maintains high throughput for moderate block sizes, preserving the efficiency that motivates linear recurrences.

2. Higher-order and block diagonal linear recurrent networks.

To bridge the expressivity-efficiency gap, several structured LRNN architectures have been proposed, including those utilizing diagonal-plus-low-rank matrices (Yang et al., 2024a; Peng et al., 2025) and their products (Siems et al., 2025). An alternative approach involves designing higher-order recurrences (which are functionally equivalent to block-diagonal architectures), such as the second-order recurrence proposed by (Rusch & Rus, 2024). In this work, we introduce a novel class of block-diagonal architectures that, in contrast to prior methods, explicitly addresses the joint normalization of the selective state-transition matrix and the selective input gates.

Higher-order Recurrence In contrast to prior work, we generalize linear recurrences to arbitrary orders, introducing the Higher-order Linear Recurrent Units (H-LRUs) as follows:

$$\mathbf{h}_t = \sum_{i=1}^{m} \mathbf{a}_{i,t} \odot \mathbf{h}_{t-i} + \mathbf{a}_{0,t} \odot \mathbf{v}_t. \qquad \text{(H-LRU)}$$

This parametrizes the state evolution by an m-th order recurrence, where m denotes the window size of the difference

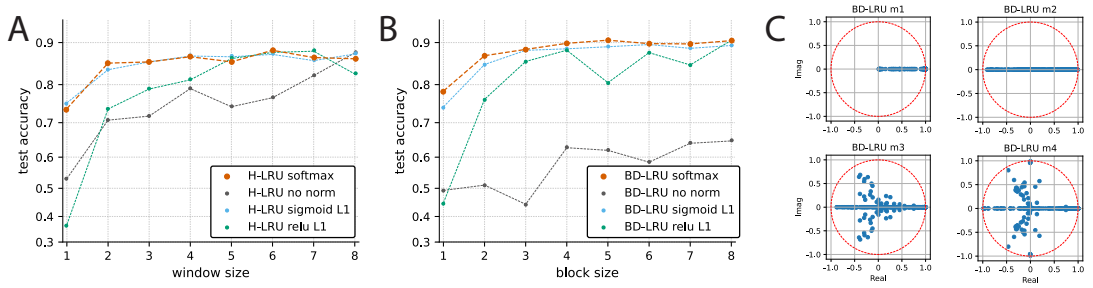

Figure 2. Scaling of performance with window/block size on the compression task for L1 normalization with different parameterizations. Results are shown for different window/block sizes m of the higher-order LRU (H-LRU) and block diagonal LRU (BD-LRU). **A**. Comparison between H-LRUs. **B**. Comparison between BD-LRUs. **C**. Eigenvalue spectra of the transition matrices learned by BD-LRU on the $S5$ dataset. BD-LRU with softmax exhibits negative eigenvalues starting from $m = 2$ and complex eigenvalues from $m = 3$. Other configurations are reported in Appendix A.

equation. Such models are a standard tool in time series statistics for forecasting (ARMA processes, see e.g. Hamilton (2020)) and are *canonical* in systems theory, since they lead to minimal realization (i.e., with the smallest memory size) of linear dynamical systems (Glad & Ljung, 2018).

To see the block-diagonal structure of H-LRU and the connection with controllable canonical forms for transition matrices in systems theory, it is sufficient to denote by h^k_{t-1} the k-th coordinate ($k \in \{1, 2, \ldots, N\}$) of $\mathbf{h}_t$ and by $a^k_{i,t}$ the k-th coordinate of $\mathbf{a}_{i,t}$. Then, with $\times$ denoting the standard matrix multiplication,

$$\begin{aligned}
\mathbf{h}^k_t &= \mathbf{A}^k_t \times \mathbf{h}^k_{t-1} + \mathbf{a}^k_{0,t} \odot \mathbf{v}^k_t, \\
\mathbf{A}^k_t &= \begin{bmatrix} a^k_{1,t} & \cdots & a^k_{m-1,t} & a^k_{m,t} \\ 1 & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & 1 & 0 \end{bmatrix}, \\
\mathbf{h}^k_{t-1} &= \begin{bmatrix} h^k_{t-1} \\ \vdots \\ h^k_{t-m} \end{bmatrix}, \ \mathbf{a}^k_{0,t} = \begin{bmatrix} a^k_{0,t} \\ \vdots \\ 0 \end{bmatrix}, \ \mathbf{v}^k_t = \begin{bmatrix} v^k_t \\ \vdots \\ 0 \end{bmatrix},
\end{aligned} \tag{1}$$

where $\mathbf{A}^k$ is a structured companion-like matrix which allows richer dynamic modes (e.g. oscillatory modes). Though eigenvalues for $\mathbf{A}^k_t$ are not available in closed form[1], stability for the system above can be guaranteed and is crucial for performance, as we will discuss in the next section.

Block Diagonal Recurrence. Independently, we also introduce a second kind of recurrence with structured state mixing, the block diagonal linear recurrent unit (BD-LRU). In contrast to the structured temporal state mixing implemented inside H-LRU blocks, BD-LRU implements dense channel mixing inside all blocks for all vectors and matrices by setting

$$\begin{aligned}
\mathbf{h}^k_t &= \mathbf{A}^k \times \mathbf{h}^k_{t-1} + \mathbf{a}^k_{0,t} \odot \mathbf{v}^k_t \qquad \text{(BD-LRU)} \\
\mathbf{A}^k_t &= \begin{bmatrix} a^k_{1,1,t} & \cdots & a^k_{1,m-1,t} & a^k_{1,m,t} \\ a^k_{2,1,t} & \cdots & a^k_{2,m-1,t} & a^k_{2,m,t} \\ \vdots & \ddots & \vdots & \vdots \\ a^k_{m,1,t} & \cdots & a^k_{m,m-1,t} & a^k_{m,m,t} \end{bmatrix}, \\
\mathbf{h}^k_{t-1} &= \begin{bmatrix} h^k_{1,t-1} \\ \vdots \\ h^k_{m,t-1} \end{bmatrix}, \ \mathbf{a}^k_{0,t} = \begin{bmatrix} a^k_{1,0,t} \\ \vdots \\ a^k_{m,0,t} \end{bmatrix}, \ \mathbf{v}^k_t = \begin{bmatrix} v^k_{1,t} \\ \vdots \\ v^k_{m,t} \end{bmatrix}.
\end{aligned}$$

To preserve the consistency of notation as in H-LRU, the block size m of BD-LRU corresponds to the size of a square matrix $\mathbf{A}^k$ and $k \in [1, N]$ corresponds to the block index of this matrix. The hidden size of BD-LRU is equal to the extended block diagonal representation of the H-LRU architecture. But in contrast to H-LRU, all vectors $\mathbf{a}^k_0, \mathbf{h}^k_t, \mathbf{v}^k_t \in \mathbb{R}^m$ and all matrices $\mathbf{A}^k \in \mathbb{R}^{m \times m}$ in BD-LRU are dense and there is no dependence on the several previous hidden states that is characteristic of the H-LRU architecture. Importantly, the structure of BD-LRU does not allow for the same eigenvalue analysis as is possible for H-LRU. Yet, as we show in the next section, we can guarantee its dynamical stability using a normalization technique similar to that of H-LRU.

Gate Selectivity. Prior literature demonstrates that input-dependent gate selectivity is crucial for diagonal models to achieve competitive performance (Gu & Dao, 2023). To equip our models with such selectivity, we introduce input-dependent gates for both H-LRU ($a'_{j,t} = \mathit{Linear}_j(\mathbf{x}_t)$) and BD-LRU ($a'_{i,j,t} = \mathit{Linear}_{i,j}(\mathbf{x}_t)$). Fig. 1 provides a schematic illustration of the proposed selective gating mechanisms in block-diagonal form, showing both the state gates that form the state-transition matrix and the input gates applied to external inputs. An ablation study investigating non-selective model variants is presented in Appendix F.

[1]Solve the equation $\chi_{\mathbf{A}^k}(\lambda) = \det(\lambda I - \mathbf{A}^k) = \lambda^m - a^k_{1,t}\lambda^{m-1} - a^k_{2,t}\lambda^{m-2} - \cdots - a^k_{m-1,t}\lambda - a^k_{m,t} = 0$.

Models	S_3 (10k samples)	S_3 (250)	S_4 (50k)	S_4 (3k)	S_5 (100k)	Overall
LSTM	**1.000**	0.316	**1.000**	0.368	**1.000**	0.737
Mamba2	0.657	0.281	0.429	0.119	0.261	0.349
H-LRU m5	0.622	0.318	0.448	0.194	0.377	0.392
BD-LRU m1	0.563	0.381	0.343	0.218	0.215	0.344
BD-LRU m2	**1.000**	0.488	0.699	0.356	0.338	0.576
BD-LRU m3	**1.000**	**1.000**	**1.000**	0.426	0.484	0.782
BD-LRU m4	**1.000**	**1.000**	**1.000**	**1.000**	0.882	0.976
BD-LRU m5	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**

Table 1. Model performance on permutation composition tasks. We note that BD-LRU performance improves with block size m, demonstrating strong sample efficiency by solving the tasks even given limited training data. See Tables 3 and 4 for extended results with statistics.

3. Normalization

Constraining the eigenvalues of the state-transition matrix is a standard technique to prevent vanishing and exploding gradients in RNNs (Pascanu et al., 2013; Arjovsky et al., 2016; Helfrich et al., 2018) and recent linear recurrent units (Orvieto et al., 2023) and state space models (Gu et al., 2021) apply similar constraints. However, as detailed in Orvieto et al. (2023), stability in a dynamical systems sense (i.e., requiring that the eigenvalues of the hidden-to-hidden transition be less than one in absolute value) does not necessarily guarantee a properly normalized forward pass in this case. If the model heavily relies on past states, the accumulated hidden states can still explode unless the incoming input is scaled down proportionally. Therefore, to maintain a bounded forward pass, both the input and forget gates must be controlled.

In purely diagonal models ($m = 1$), this is typically achieved by linking the gates (as seen in S4 (Gu et al., 2020) and Mamba (Gu & Dao, 2023)) or by defining the recurrence as a direct interpolation: $\mathbf{h}_t = \mathbf{a}_t \odot \mathbf{h}_{t-1} + (1 - \mathbf{a}_t) \odot \mathbf{v}_t$, analogous to the design of GRUs (Cho et al., 2014) and recent linear models (Feng et al., 2024). In this work, we derive the following proposition that establishes the conditions required to generalize this normalization principle to our higher-order and block-diagonal architectures ($m \geq 1$). The proof is provided in Appendix E.

Proposition 1 *Consider either the H-LRU or the BD-LRU architectures, written in matrix form as shown in Eq. 1 and in the BD-LRU equation. If for any $k \in [1, N]$, the k-th recurrent non-diagonal block $\mathbf{h}_t^k = \mathbf{A}_t^k \times \mathbf{h}_{t-1}^k + \mathbf{a}_{0,t}^k \odot \mathbf{v}_t^k$ is such that the matrix $\mathcal{A}_t^k := [\mathbf{A}_t^k, \mathbf{a}_{0,t}^k] \in \mathbb{R}^{m\times(m+1)}$ has the property that $\sum_{j=1}^{m+1} |(\mathcal{A}_t^k)_{i,j}| \leq 1$ for every row $i \in [1, m]$, then the recurrence is stable from a dynamical systems perspective and the forward pass is normalized, meaning that $\|\mathbf{h}_T\|_\infty \leq \max_{t\in[0,T]} \|\mathbf{v}_t\|_\infty$.*

The proposition above suggests that to achieve a normalized forward pass, L1-normalization should be applied to raw selective gates. For H-LRU, it is sufficient to normalize over all $m + 1$ coefficients of the m-th order recurrence, while for BD-LRU, we apply a row-wise normalization over the hidden state gates and the input gate. Let us therefore denote as a's the raw gates (linear functions of the input) before normalization. We set

$$\begin{aligned} \textbf{H-LRU: } a_{j,t} &= \frac{f(a'_{j,t})}{\sum_{l=0}^{m} |f(a'_{l,t})|}; \\ \textbf{BD-LRU: } a_{i,j,t} &= \frac{f(a'_{i,j,t})}{\sum_{l=0}^{m} |f(a'_{i,l,t})|}, \end{aligned} \tag{2}$$

where $f(\cdot)$ is a gate parametrization function; the block index is omitted for clarity. Note that this normalization only affects the elements inside on-diagonal blocks and has no impact on off-diagonal blocks (consisting of zero matrices). Note that according to the Gershgorin circle theorem, our proposed normalization scheme explicitly couples the eigenvalues of the state-transition matrix with the input gate. Specifically, for the BD-LRU, a row-wise application of the theorem yields the bound $|\lambda_{i,t}| \leq \max_i \left(\sum_{l=1}^{m} |a_{i,l,t}|\right) = \max_i(1 - |a_{i,0,t}|)$. Conversely, for the H-LRU, a column-wise application results in $|\lambda_{i,t}| \leq |a_{m,t}| = 1 - \sum_{l=0}^{m-1} |a_{l,t}|$. This results in a joint normalization for input and state gates that allows selective block-diagonal LRNNs to balance attention to hidden states and inputs in a similar way as in first-order models (Orvieto et al., 2023; De et al., 2024). This is in contrast to previous studies on selective block-diagonal LRNNs that only addressed the stability of the state-transition matrix (Fan et al., 2023).

Although the introduced normalization guarantees the stability of the recurrence, it has been shown that gradient-based learning is also highly sensitive to the specific choice of parametrization (Zucchet & Orvieto, 2024). In contrast to the normalization used in non-selective block-diagonal LRNNs that rely on structured parameterizations such as discretization schemes (Rusch & Rus, 2024; Walker et al., 2025), joint parametrization of the state-transition matrices and input gate (Biegun et al., 2024), and exponential reparametrization (Orvieto et al., 2023), our proposed nor-

Models	In-Context Recall	Selective Copy	Compression	Overall
LSTM	**1.000 (1.000 ± 0.000)**	**1.000 (0.998 ± 0.002)**	0.749 (0.697 ± 0.092)	0.916
Mamba2	**1.000 (0.988 ± 0.012)**	0.812 (0.786 ± 0.032)	0.721 (0.582 ± 0.131)	0.844
Deltanet[-1,1]	**1.000 (1.000 ± 0.000)**	0.902 (0.870 ± 0.040)	0.765 (0.665 ± 0.110)	0.889
Deltaproduct$_4$[-1,1]	**1.000 (1.000 ± 0.000)**	**1.000 (0.998 ± 0.001)**	0.735 (0.655 ± 0.073)	0.911
BD-LRU m1	0.779 (0.721 ± 0.036)	0.835 (0.815 ± 0.016)	0.725 (0.580 ± 0.119)	0.780
BD-LRU m2	**1.000 (1.000 ± 0.000)**	0.964 (0.954 ± 0.009)	0.760 (0.645 ± 0.129)	0.908
BD-LRU m3	**1.000 (1.000 ± 0.000)**	0.984 (0.975 ± 0.006)	0.764 (0.642 ± 0.126)	0.916
BD-LRU m4	**1.000 (1.000 ± 0.000)**	0.985 (0.976 ± 0.006)	**0.786 (0.748 ± 0.047)**	**0.924**
BD-LRU m8	**1.000 (1.000 ± 0.000)**	0.996 (0.946 ± 0.032)	0.748 (0.685 ± 0.079)	0.915
H-LRU m1	0.783 (0.721 ± 0.038)	0.847 (0.813 ± 0.024)	0.758 (0.650 ± 0.133)	0.796
H-LRU m2	0.998 (0.993 ± 0.004)	0.853 (0.827 ± 0.023)	0.772 (0.692 ± 0.061)	0.874
H-LRU m3	**1.000 (1.000 ± 0.000)**	0.853 (0.830 ± 0.018)	0.773 (0.673 ± 0.136)	0.875
H-LRU m4	**1.000 (1.000 ± 0.000)**	0.845 (0.823 ± 0.019)	0.774 (0.736 ± 0.038)	0.873
H-LRU m8	**1.000 (0.976 ± 0.024)**	0.807 (0.776 ± 0.022)	0.766 (0.694 ± 0.060)	0.858

Table 2. Performance on the in-context recall, selective copying and compression tasks. The presented results are the average of best test accuracies across four configurations of the corresponding synthetic dataset with different vocabulary sizes, sequence lengths and number of training examples (mean ± std in brackets across 5 seeds). Results are shown for different window (H-LRU) and block sizes (BD-LRU) m. The overall performance of our models consistently improves with window/block size up to approximately 3–5, after which the gains saturate. All models are single-layer configurations with a maximum overall hidden dimension of 6144. See Table 5 for extended results.

malization is more general as it can be applied to a variety of both non-selective and selective parameterizations. This allowed us to independently evaluate several variants of gate parameterizations that are defined by the function f in Eq. 2. As can be seen in Fig. 2, *our normalization strategy greatly improves performance* of both H-LRUs and BD-LRUs.

4. Experiments on state-tracking permutation tasks

An important property of dense recurrent networks is that one layer of such model can easily solve inherently sequential tasks that require state-tracking. In theory, linear diagonal networks can also solve any of these tasks, but only if we assume an infinite depth approximation. In practice, it has been shown that they cannot effectively approximate the evolution of recurrent state with a bounded number of layers (Merrill et al., 2024). Further, it was proposed that there is a parallelism-expressivity trade-off, in which efficient parallelization comes at the expense of decreased expressivity (Merrill & Sabharwal, 2023).

To evaluate the ability of a model to perform state-tracking we use a synthetic dataset based on the symmetric group S_n, the group of all permutations over n elements (Merrill et al., 2024). Each instance in the dataset corresponds to a specific permutation sampled from S_n, and the model is tasked with learning the mapping that defines the permutation purely from input-output examples within a sequence. We evaluated model performance on a series of increasingly complex permutation learning tasks derived from the symmetric groups S_2 through S_5.

Complex permutations are challenging. All tested recurrent architectures (H-LRU, BD-LRU, LSTM, Deltanet, Deltaproduct) were able to perfectly solve the S_2 task, which represents a uniquely simple permutation group as it is also a commutative cyclic group. However, as the group order increases over S_3 to S_5, the non-commutative structure of the permutation tasks increasingly posed challenges for the models, see Table 1, Table 3, and Table 4. Performance of the H-LRU was found to decrease progressively with increasing group size, indicating a limited capacity for modeling compositional permutations. Increasing the order of recurrence m did not seem to provide any benefits for the performance. We conclude that a strict temporal inductive bias on the structure of the transition matrix prevents H-LRU from solving this task. Moreover, we found that H-LRU is unable to solve our permutation tasks despite having access to negative and complex eigenvalues (see Fig. 2 and Appendix A for our eigenvalue analysis). This indicates that the presence of such eigenvalues is insufficient for these tasks and highlights that the structure of state mixing plays a more critical role.

BD-LRUs efficiently learn complex permutations. In contrast, BD-LRU with moderate block sizes was able to successfully solve all permutation tasks for all group sizes, matching the performance of LSTM and outperforming all other recurrent architectures tested. Notably, BD-LRU with block size 5 solved the S_5 task using as few as 200K parameters, matching the parameter efficiency of more computationally demanding non-linear LSTM model. Furthermore, we found that BD-LRUs are also sample-efficient in learning permutations, outperforming even LSTM in the regime of limited training data. We note that in our low training token regime Deltaproduct$_4$ fails to learn the S_5 dataset. However, when the number of training samples approaches the token counts used in the study (Siems et al., 2025), it is capable of solving S_5 task, showing that low-rank matrices

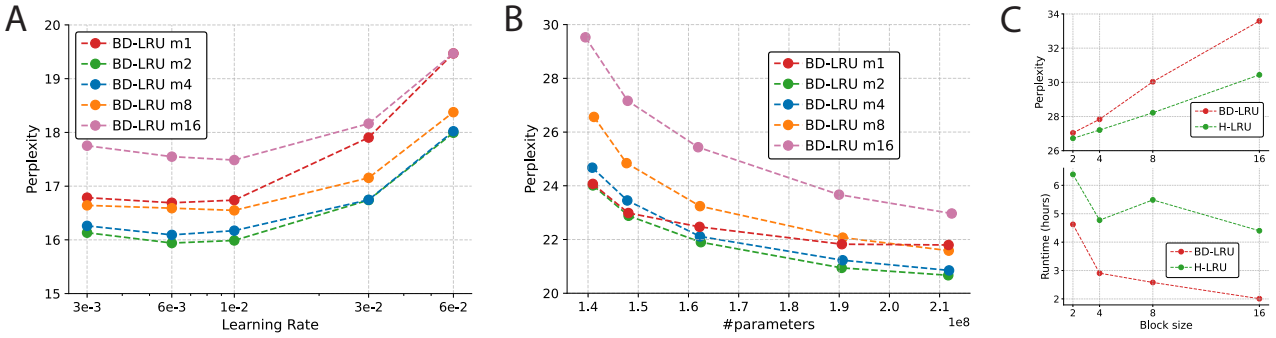

Figure 3. Language modeling results on FineWeb. **A.** Best achieved perplexity for 210M parameter BD-LRU models (trained on 10B tokens) across varying learning rates. **B.** Performance scaling of BD-LRU models on 2.5B tokens with varying hidden dimensions. Results indicate that moderate block sizes provide a superior inductive bias. **C.** Runtime and perplexity comparison for 140M parameter models. While H-LRUs are parameter-efficient, matching the parameter budget of a BD-LRU requires increasing the H-LRU hidden dimension by a factor of m, making them substantially more costly to scale.

are less sample-efficient than BD-LRU. Our findings align well with our predictions that dense blocks of BD-LRU are well-suited for implementing permutations between hidden states. The consistent improvement with larger block sizes on permutation tasks of increasing complexity highlights the advantage of the inductive bias of BD-LRU architecture.

5. Experiments on token manipulation tasks

While sequence modeling is often evaluated through large-scale training, recent work demonstrates that critical capabilities can be effectively assessed via targeted synthetic benchmarks (Arora et al., 2023; Poli et al., 2024). Leveraging the established equivalence between lossless compression and generalization (Shannon, 1948; Delétang et al., 2023; Gu, 2025), we first evaluate temporal information integration using the auto-encoding compression task. Furthermore, to assess the dynamic adaptation required for general sequence modeling, we employ selective copying and associative recall (Poli et al., 2024; Olsson et al., 2022; Waleffe et al., 2024).

Normalization allows scaling with block size. In experiments on synthetic tasks, we show that our normalization strategies are crucial for performance, see Fig. 2. We tested several variants of the function f for L1 normalization in Eq. 2: exponentiated gate $exp(\cdot)$ (softmax normalization), sigmoidal gates $\sigma(\cdot)$, ReLU gates $relu(\cdot)$. As a baseline, we also tested all models without normalization.

We found that both softmax and sigmoidal L1 normalizations allowed the models to effectively scale with window and block size. Without normalization and with the ReLU normalization, both H-LRU and BD-LRU improve at lower rate with window size. With softmax or sigmoidal L1 normalizations, the improvement with window size was especially pronounced between a window/block size of 1 and 2. Our eigenvalue analysis (see Appendix A) indicates that this gain corresponds to the emergence of negative eigenvalues, consistent with the findings of (Grazzi et al., 2024). We also observe that further improvements in performance are associated with a broader range of complex eigenvalues, which are enabled starting from the block size 3. These results also align well with previous studies on the beneficial role of oscillatory dynamics in recurrent networks (Rusch & Mishra, 2021; Effenberger et al., 2022; Dubinin & Effenberger, 2024; Rusch & Rus, 2024).

We noticed that for moderate block sizes ($m \in [2, 5]$), the softmax normalization performed comparable to or better than sigmoidal normalization, making this the default choice for all the remaining experiments. That also agrees with previous findings that exponentiation of the gates benefits gradient descent (Orvieto et al., 2023; Zhang et al., 2024).

Performance is limited by state mixing. Next, we performed experiments in which we investigated the difference between scaling the window size and the hidden size. In these experiments we found that for both H-LRUs and BD-LRUs, the scaling with hidden size could not compensate for a lack of expressivity. In other words, window/block size was found to be the key factor for performance, see Fig. 7. We also found that scaling of H-LRUs and BD-LRUs results in models that are competitive with LSTMs and achieve higher performance than other linear recurrent baselines, both diagonal ones such as Mamba2 and low-rank ones such as Deltanet and Deltaproduct, see Table 2. In line with the observed limitations of diagonal RNNs, we found that scaling the hidden size in a Mamba2 model also had limited effect on performance, see Fig. 7. Notably, we also found distinct scaling behaviors for the compression and our other tasks, aligning with previous results (Delétang et al., 2023). In the compression (auto-encoding) task, models with smaller block size outperformed larger counterparts, while performance on autoregressive tasks scaled positively

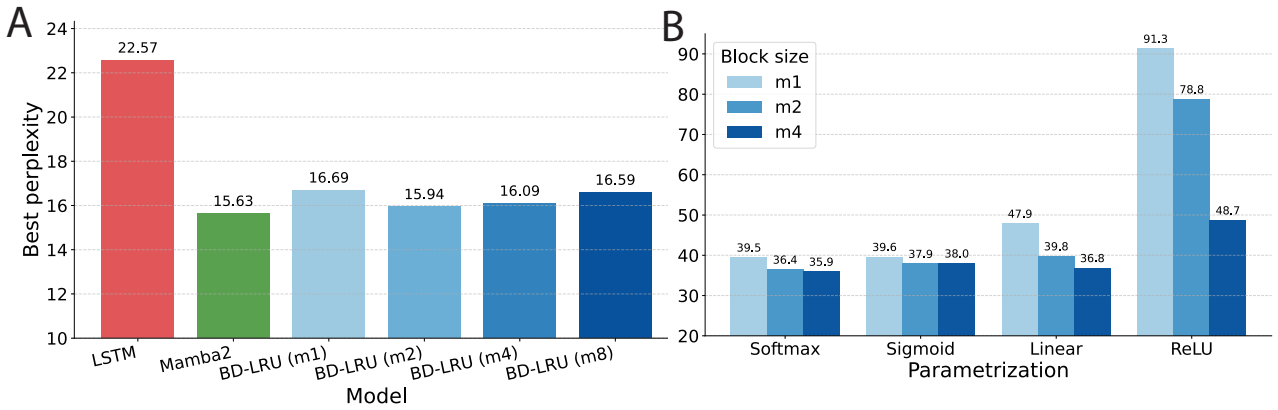

Figure 4. Language modeling results on FineWeb. **A.** Best achieved perplexity for ∼200M parameter models trained on 10B tokens. **B.** Performance scaling of BD-LRU models on 0.5B tokens with varying parametrization and block size.

with block size. Therefore, the decrease in aggregate performance for larger block sizes is substantially driven by the results on the compression task.

Our experiments show a direct trade-off between parameter efficiency and peak performance, as governed by the block and window sizes for BD-LRU and H-LRU, respectively. Models with smaller block/window sizes saturate at lower parameter counts, demonstrating high efficiency. In contrast, models with larger block/window sizes require a larger hidden dimension to match the performance of the smaller models, but can ultimately achieve a much higher performance. This indicates that richer state mixing increases a model's expressive power at the expense of parameter efficiency.

H-LRUs are parameter efficient. We also found that in the compression task which does not require complex token manipulation, H-LRU demonstrated the most parameter efficient scaling with hidden size, achieving accuracies not accessible to Mamba2 and LSTM of similar sizes (in terms of the number of trainable parameters), see Fig. 7. This aligns well with our predictions that the inductive bias introduced by extended temporal mixing results in hidden representations with better compression capabilities.

BD-LRUs are expressive across tasks. In contrast to the compression task, the selective copying task requires more extensive token manipulation. We found that the performance of BD-LRUs scales more favorably with hidden size than that of H-LRUs. Furthermore, BD-LRUs were able to outperform Mamba2 and Deltanet, achieving performance that is competitive with LSTMs and Deltaproduct. At the same time, BD-LRUs achieved the best performance in the compression task as well. Overall, the introduced normalization scheme allows BD-LRUs to efficiently utilize the expressivity of their dense block diagonal structure to approximate a variety of mixing patterns and to achieve the best overall results on our set of synthetic tasks.

6. Experiments on language modeling

Our language modeling experiments with BD-LRU and H-LRU further corroborate the findings from our synthetic task evaluations, see Fig. 3. When trained on the FineWeb dataset, BD-LRU with moderate block sizes achieves the lowest perplexity among parameter-matched BD-LRU variants (210M parameters trained for 10B tokens; see Fig. 3A). Architectures with block sizes of 2, 4 and 8 outperform diagonal networks, while models with block size of 16, despite being theoretically more expressive, underperform in practice.

By varying the hidden size of BD-LRU, we obtain models in the 140M–210M parameter range, see Fig. 3B. BD-LRU with moderate block sizes effectively scale with parameter count whereas diagonal models ($m = 1$) show early saturation with increasing hidden size. Overall, these results indicate that moderate block sizes provide a more effective inductive bias for scaling of language models, in line with our observations on synthetic tasks.

We also conducted language-modeling experiments with H-LRU using configurations matched in parameter count to their BD-LRU counterparts, see Fig. 3C for 140M parameters. Consistent with our synthetic benchmarks, H-LRU exhibits stronger parameter efficiency. However, to match the parameter budget of a BD-LRU, H-LRU requires increasing the hidden dimension by a factor of m, which in turn reduces throughput and increases memory consumption by approximately the same factor, see Sec. 7 and Fig. 3C. Therefore, although H-LRU is more parameter-efficient, it is more computationally demanding to scale compared to BD-LRU.

When controlling for parameter count, network depth, and hidden dimension, we find that BD-LRU achieves performance highly competitive with Mamba2, though Mamba2 maintains a marginal advantage (∼ 0.3 PPL; Fig. 4A). Notably, both BD-LRU and Mamba2 significantly outperform

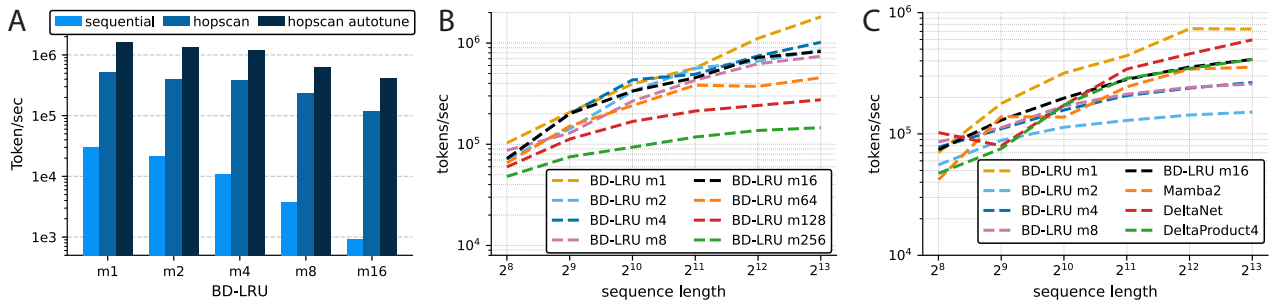

Figure 5. Model throughput on the selective copying task. (**A**) Comparison of sequential, higher-order parallel, and autotuned higher-order parallel implementations of BD-LRUs with 128 blocks and with a sequence length of 2048, illustrating the advantage of parallel scan implementation and the trade-off between expressivity and efficiency. BD-LRU is shown for illustration purposes only, but H-LRU employs the same parallel scan implementation. (**B**) Comparison for layers with a hidden size of 768 and accordingly adjusted number of blocks. Note that the trade-off between expressivity and efficiency increases over longer sequences. (**C**) Throughput comparison of parameter-matched layers (∼33M parameters). The number of blocks is adjusted to ensure consistent model sizes across architectures. BD-LRU achieves throughput competitive with other LRNN baselines. Notably, larger block sizes demonstrate higher practical efficiency despite increased theoretical complexity, due to superior utilization of GPU hardware operations.

the LSTM baseline, demonstrating that structured LRNNs can compensate for the lack of fully dense state mixing. Furthermore, our language modeling experiments also confirm that our proposed normalization scheme successfully enables performance to scale with increasing block sizes (Fig. 4B).

7. Implementation

The parallel scan algorithm in LRNNs allows them to efficiently process long sequences using constant memory and with logarithmic time complexity. Following the classic approach (Blelloch, 1990), we consider a recurrence of the form

$$\mathbf{h}_{i+1} = \begin{cases} \mathbf{b}_0, & \text{if} \quad i = 0 \\ (\mathbf{h}_i \bigotimes_v \mathbf{A}_i) \bigoplus \mathbf{b}_i, & \text{if } 0 \leq i < T \end{cases}, \tag{3}$$

where $\mathbf{h}_i, \mathbf{b}_i \in \mathbb{R}^H$, $\mathbf{A}_i \in \mathbb{R}^{H \times H}$ and associative operators: $\bigotimes_v$ is matrix-vector multiplication, $\bigotimes_M$ is matrix-matrix multiplication and $\bigoplus$ point-wise vector summation.

Defining the following associative operator $\bullet$ and making substitution to sequence of pairs,

$$HOP = \begin{cases} \mathbf{c}_i = [\mathbf{A}_i, \mathbf{b}_i] \\ \mathbf{c}_i \bullet \mathbf{c}_j \equiv [\mathbf{c}_{i,A} \bigotimes_M \mathbf{c}_{j,A}, (\mathbf{c}_{i,b} \bigotimes_v \mathbf{c}_{j,A}) \bigoplus \mathbf{c}_{j,b}] \end{cases} \tag{4}$$

reduces recurrence 3 to a classic prefix sum and allows application of up and down sweeps of the Blelloch scan.

In many modern LRNNs, $\mathbf{A}_i$ is diagonal ($\mathbf{c}_{i,A} \bigotimes_M \mathbf{c}_{j,A} \sim H$), therefore parallel scan 4 enables efficient parallel processing by reducing the time complexity from HT to $H \log(T)$. However, in the more general case presented in Eq. 3, parallel scan changes the time complexity from H^2T to $H^3 \log(T)$. For large dense matrices $\mathbf{A}_i$ and/or short sequences, this change in complexity is not beneficial due to the high complexity of matrix-matrix multiplication ($\mathbf{c}_{i,A} \bigotimes_M \mathbf{c}_{j,A} \sim H^3$). However, if we exploit the block diagonal structure of the transition matrices in H-LRU and BD-LRU, we can reduce the time complexity of parallel scan from $H^3 \log(T)$ to $Nm^3 \log(T)$, where m is the block size and N is the number of blocks ($Nm = H$). Therefore, for moderate block sizes ($m \ll H$), our parallel scan significantly outperforms sequential implementations and remains competitive with purely diagonal scans, incurring only an $\mathcal{O}(m^2)$ computational overhead.

Parallel scan implementation enables competitive throughput. In experiments with single-layer models containing 128 blocks and trained on sequences of length 2048, when runtime is less influenced by GPU characteristics and more reflective of algorithmic complexity, we found that increasing block size reduces throughput, revealing the predicted trade-off between expressivity and efficiency, see Fig. 5A. For comparison, we also evaluated models with a fixed hidden size of 768 and adjusted the number of blocks accordingly, see Fig. 5B. We found that the expressivity–efficiency trade-off becomes more pronounced as sequence length increases. In particular, block sizes larger than 16 exhibit a substantial decline in throughput at longer sequence lengths.

We also tested larger parameter-matched layers (∼33M parameters), where the number of blocks is adjusted to ensure consistent model sizes across architectures, see Fig. 5C. We note that our most efficient implementation relies on compilation with maximal autotuning; thus, at such scale, the performance differences across block sizes primarily reflect kernel optimization in PyTorch and achieved GPU utilization. We found that certain block sizes align more favorably with GPU architectures. In particular, we found that moderately large block sizes ($m = 16$) demonstrate higher practical efficiency despite increased theoretical complexity,

due to superior utilization of GPU hardware operations.

Overall, we observed that our parallel scan implementation offers substantial improvements over sequential implementations, enables BD-LRUs and H-LRUs to achieve throughput comparable to that of linear baselines, and effectively scales in agreement with prior studies on SSMs (Walker et al., 2025).

8. Conclusion and outlook

We introduced H-LRU and BD-LRU as structured extensions of linear recurrent models that enhance temporal and channel-wise state mixing. Our results show that proper selective gate normalization is essential for scaling such models with window/block size, that H-LRU excels at parameter-efficient compression, while BD-LRU is overall the best-performing architecture on our benchmark of synthetic tasks. We also provide our parallel-scan implementation that maintains competitive efficiency of block-diagonal architectures. Overall, our empirical results indicate that the structure of state-mixing acts as an important driver for improved performance in LRNNs.

In our experiments, we observed clear task-dependent differences in how performance scales with block size. Simple tasks such as in-context recall, S3, and Parity are effectively solved with block size 2, nearly eliminating any expressivity–efficiency trade-off. More challenging autoregressive problems such as selective copying, S4, S5, and Regular Languages benefit substantially from larger block sizes. In contrast, the compression auto-encoding task exhibits a distinct scaling pattern: intermediate block sizes achieve the best results, while very large blocks degrade their performance.

We find that H-LRU is particularly effective on compression, likely due to its higher-order recurrence structure, whereas BD-LRU is highly parameter- and sample-efficient on permutation-heavy tasks, consistent with the advantages of dense intra-block mixing. Importantly, both architectures maintain strong throughput on long sequences, making moderate-to-large block sizes viable in practice.

We also observe the same scaling behavior with block size in our language modeling experiments, supporting the general nature of our findings. However, the strong state-tracking capabilities of the BD-LRU (and LSTM) did not translate to a proportional advantage over the Mamba2 model on the FineWeb dataset. This indicates that other architectural components of highly optimized Mamba2 models (e.g., causal convolutions) can compensate for this expressivity gap in autoregressive language modeling.

Overall, our results indicate that the optimal block or window size m is inherently task-dependent. In practice, we recommend beginning with moderate block/window sizes (with moderate hidden dimension) and adjusting upward or downward based on task complexity, sequence length, and modeling objective, thereby navigating the expressivity–efficiency trade-off. More broadly, the problem of selecting appropriate inductive biases remains an open research question in machine learning, and we hope that our findings contribute an additional perspective to this ongoing direction of research.

One potential limitation is that we do not evaluate our models on large-scale language modeling; however all our contributions are thoroughly supported without requiring such scale. We note that our primary objective is to study structured state mixing rather than to establish a new SOTA baseline on large-scale language modeling. We also acknowledge that, while our parallel-scan implementation is efficient, its throughput still lags behind heavily hardware-optimized baselines. Developing custom, hardware-aware kernels presents a clear objective for future engineering efforts.

Acknowledgments

We would like to thank Wolf Singer, Sajad Movahedi and Felix Sarnthein for the helpful discussions. Antonio Orvieto acknowledges financial support from the Hector Foundation and the AI2050 Early Career Fellowship from Schmidt Sciences.

9. Reproducibility, LLM usage and Impact statements

We provide anonymized code for our synthetic tasks experiments at `https://anonymous.4open.science/r/mad-lab-000E` and for our language modeling experiments at `https://anonymous.4open.science/r/plainLM-BD37`. To ensure full reproducibility, our released code repository includes: the exact experimental configurations, implementations of our models and baselines, the parallel-scan implementation, scripts for generating the synthetic data, and the data preparation pipeline for language modeling.

LLMs were used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research.

The primary goal of this work is to advance foundational research in machine learning. As such, the proposed architectures are not currently tied to any particular downstream applications that present direct or immediate societal impacts.

References

Ajroldi, N. plainlm: Language model pretraining in pytorch, 2024.

Arjovsky, M., Shah, A., and Bengio, Y. Unitary evolution recurrent neural networks. In *International Conference on Machine Learning*, pp. 1120–1128. PMLR, 2016.

Arora, S., Eyuboglu, S., Timalsina, A., Johnson, I., Poli, M., Zou, J., Rudra, A., and Ré, C. Zoology: Measuring and improving recall in efficient language models. *arXiv preprint arXiv:2312.04927*, 2023.

Biegun, K., Dolga, R., Cunningham, J., and Barber, D. Rotrnn: Modelling long sequences with rotations. *arXiv preprint arXiv:2407.07239*, 2024.

Blelloch, G. E. Prefix sums and their applications. 1990.

Chang, Y. and Bisk, Y. Language models need inductive biases to count inductively. *arXiv preprint arXiv:2405.20131*, 2024.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.

Chomsky, N. Three models for the description of language. *IRE Transactions on information theory*, 2(3):113–124, 1956.

Cirone, N. M., Orvieto, A., Walker, B., Salvi, C., and Lyons, T. Theoretical foundations of deep selective state-space models. *arXiv preprint arXiv:2402.19047*, 2024.

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

De, S., Smith, S. L., Fernando, A., Botev, A., Cristian-Muraru, G., Gu, A., Haroun, R., Berrada, L., Chen, Y., Srinivasan, S., et al. Griffin: Mixing gated linear recurrences with local attention for efficient language models. *arXiv preprint arXiv:2402.19427*, 2024.

Delétang, G., Ruoss, A., Grau-Moya, J., Genewein, T., Wenliang, L. K., Catt, E., Cundy, C., Hutter, M., Legg, S., Veness, J., et al. Neural networks and the chomsky hierarchy. *arXiv preprint arXiv:2207.02098*, 2022.

Delétang, G., Ruoss, A., Duquenne, P.-A., Catt, E., Genewein, T., Mattern, C., Grau-Moya, J., Wenliang, L. K., Aitchison, M., Orseau, L., et al. Language modeling is compression. *arXiv preprint arXiv:2309.10668*, 2023.

Dubinin, I. and Effenberger, F. Fading memory as inductive bias in residual recurrent networks. *Neural networks*, 173: 106179, 2024.

Effenberger, F., Carvalho, P., Dubinin, I., and Singer, W. A biology-inspired recurrent oscillator network for computations in high-dimensional state space. *BioRxiv*, 2022.

Fan, T.-H., Chi, T.-C., and Rudnicky, A. I. Advancing regular language reasoning in linear recurrent neural networks. *arXiv preprint arXiv:2309.07412*, 2023.

Feng, L., Tung, F., Ahmed, M. O., Bengio, Y., and Hajimirsadeghi, H. Were rnns all we needed? *arXiv preprint arXiv:2410.01201*, 2024.

Glad, T. and Ljung, L. *Control theory*. CRC press, 2018.

Grazzi, R., Siems, J., Zela, A., Franke, J. K., Hutter, F., and Pontil, M. Unlocking state-tracking in linear rnns through negative eigenvalues. *arXiv preprint arXiv:2411.12537*, 2024.

Gromov, A. Grokking modular arithmetic. *arXiv preprint arXiv:2301.02679*, 2023.

Gu, A. On the tradeoffs of state space models and transformers, 2025. URL `https://goombalab.github.io/blog/2025/tradeoffs/`.

Gu, A. and Dao, T. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.

Gu, A., Gulcehre, C., Paine, T., Hoffman, M., and Pascanu, R. Improving the gating mechanism of recurrent neural networks. In *International Conference on Machine Learning*, pp. 3800–3809. PMLR, 2020.

Gu, A., Goel, K., and Ré, C. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*, 2021.

Hamilton, J. D. *Time series analysis*. Princeton university press, 2020.

Helfrich, K., Willmott, D., and Ye, Q. Orthogonal recurrent neural networks with scaled cayley transform. In *International Conference on Machine Learning*, pp. 1969–1978. PMLR, 2018.

Hutter, M. *Universal artificial intelligence: Sequential decisions based on algorithmic probability*. Springer Science & Business Media, 2005.

Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and Amodei, D. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

Loshchilov, I. and Hutter, F. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.

Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

Merrill, W. and Sabharwal, A. The parallelism tradeoff: Limitations of log-precision transformers. *Transactions of the Association for Computational Linguistics*, 11:531–545, 2023.

Merrill, W., Petty, J., and Sabharwal, A. The illusion of state in state-space models. *arXiv preprint arXiv:2404.08819*, 2024.

Movahedi, S., Sarnthein, F., Cirone, N. M., and Orvieto, A. Fixed-point rnns: From diagonal to dense in a few iterations. *arXiv preprint arXiv:2503.10799*, 2025.

Olsson, C., Elhage, N., Nanda, N., Joseph, N., DasSarma, N., Henighan, T., Mann, B., Askell, A., Bai, Y., Chen, A., et al. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*, 2022.

Orvieto, A., Smith, S. L., Gu, A., Fernando, A., Gulcehre, C., Pascanu, R., and De, S. Resurrecting recurrent neural networks for long sequences. In *International Conference on Machine Learning*, pp. 26670–26698. PMLR, 2023.

Pascanu, R., Mikolov, T., and Bengio, Y. On the difficulty of training recurrent neural networks. *International conference on machine learning*, pp. 1310–1318, 2013.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

Penedo, G., Kydlíček, H., Lozhkov, A., Mitchell, M., Raffel, C. A., Von Werra, L., Wolf, T., et al. The fineweb datasets: Decanting the web for the finest text data at scale. *Advances in Neural Information Processing Systems*, 37: 30811–30849, 2024.

Peng, B., Zhang, R., Goldstein, D., Alcaide, E., Du, X., Hou, H., Lin, J., Liu, J., Lu, J., Merrill, W., et al. Rwkv-7" goose" with expressive dynamic state evolution. *arXiv preprint arXiv:2503.14456*, 2025.

Poli, M., Thomas, A. W., Nguyen, E., Ponnusamy, P., Deiseroth, B., Kersting, K., Suzuki, T., Hie, B., Ermon, S., Ré, C., et al. Mechanistic design and scaling of hybrid architectures. *arXiv preprint arXiv:2403.17844*, 2024.

Rusch, T. K. and Mishra, S. Coupled Oscillatory Recurrent Neural Network (coRNN): An accurate and (gradient) stable architecture for learning long time dependencies. *arXiv:2010.00951 [cs, stat]*, March 2021.

Rusch, T. K. and Rus, D. Oscillatory state-space models. *arXiv preprint arXiv:2410.03943*, 2024.

Shannon, C. E. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.

Siems, J., Carstensen, T., Zela, A., Hutter, F., Pontil, M., and Grazzi, R. Deltaproduct: Improving state-tracking in linear rnns via householder products. *arXiv preprint arXiv:2502.10297*, 2025.

Sun, Y., Li, X., Dalal, K., Xu, J., Vikram, A., Zhang, G., Dubois, Y., Chen, X., Wang, X., Koyejo, S., et al. Learning to (learn at test time): Rnns with expressive hidden states. *arXiv preprint arXiv:2407.04620*, 2024.

von Oswald, J., Scherrer, N., Kobayashi, S., Versari, L., Yang, S., Schlegel, M., Maile, K., Schimpf, Y., Sieberling, O., Meulemans, A., et al. Mesanet: Sequence modeling by locally optimal test-time training. *arXiv preprint arXiv:2506.05233*, 2025.

Waleffe, R., Byeon, W., Riach, D., Norick, B., Korthikanti, V., Dao, T., Gu, A., Hatamizadeh, A., Singh, S., Narayanan, D., et al. An empirical study of mamba-based language models. *arXiv preprint arXiv:2406.07887*, 2024.

Walker, B., Yang, L., Cirone, N. M., Salvi, C., and Lyons, T. Structured linear cdes: Maximally expressive and parallel-in-time sequence models. *arXiv preprint arXiv:2505.17761*, 2025.

Wang, S. and Li, Q. Stablessm: Alleviating the curse of memory in state-space models through stable reparameterization. *arXiv preprint arXiv:2311.14495*, 2023.

Yang, S., Kautz, J., and Hatamizadeh, A. Gated delta networks: Improving mamba2 with delta rule. *arXiv preprint arXiv:2412.06464*, 2024a.

Yang, S., Wang, B., Zhang, Y., Shen, Y., and Kim, Y. Parallelizing linear transformers with the delta rule over sequence length. *arXiv preprint arXiv:2406.06484*, 2024b.

Zhang, M., Bhatia, K., Kumbong, H., and Ré, C. The hedgehog & the porcupine: Expressive linear attentions with softmax mimicry. *arXiv preprint arXiv:2402.04347*, 2024.

Zucchet, N. and Orvieto, A. Recurrent neural networks: vanishing and exploding gradients are not the end of the story. *Advances in Neural Information Processing Systems*, 37:139402–139443, 2024.

A. Eigenvalue analysis

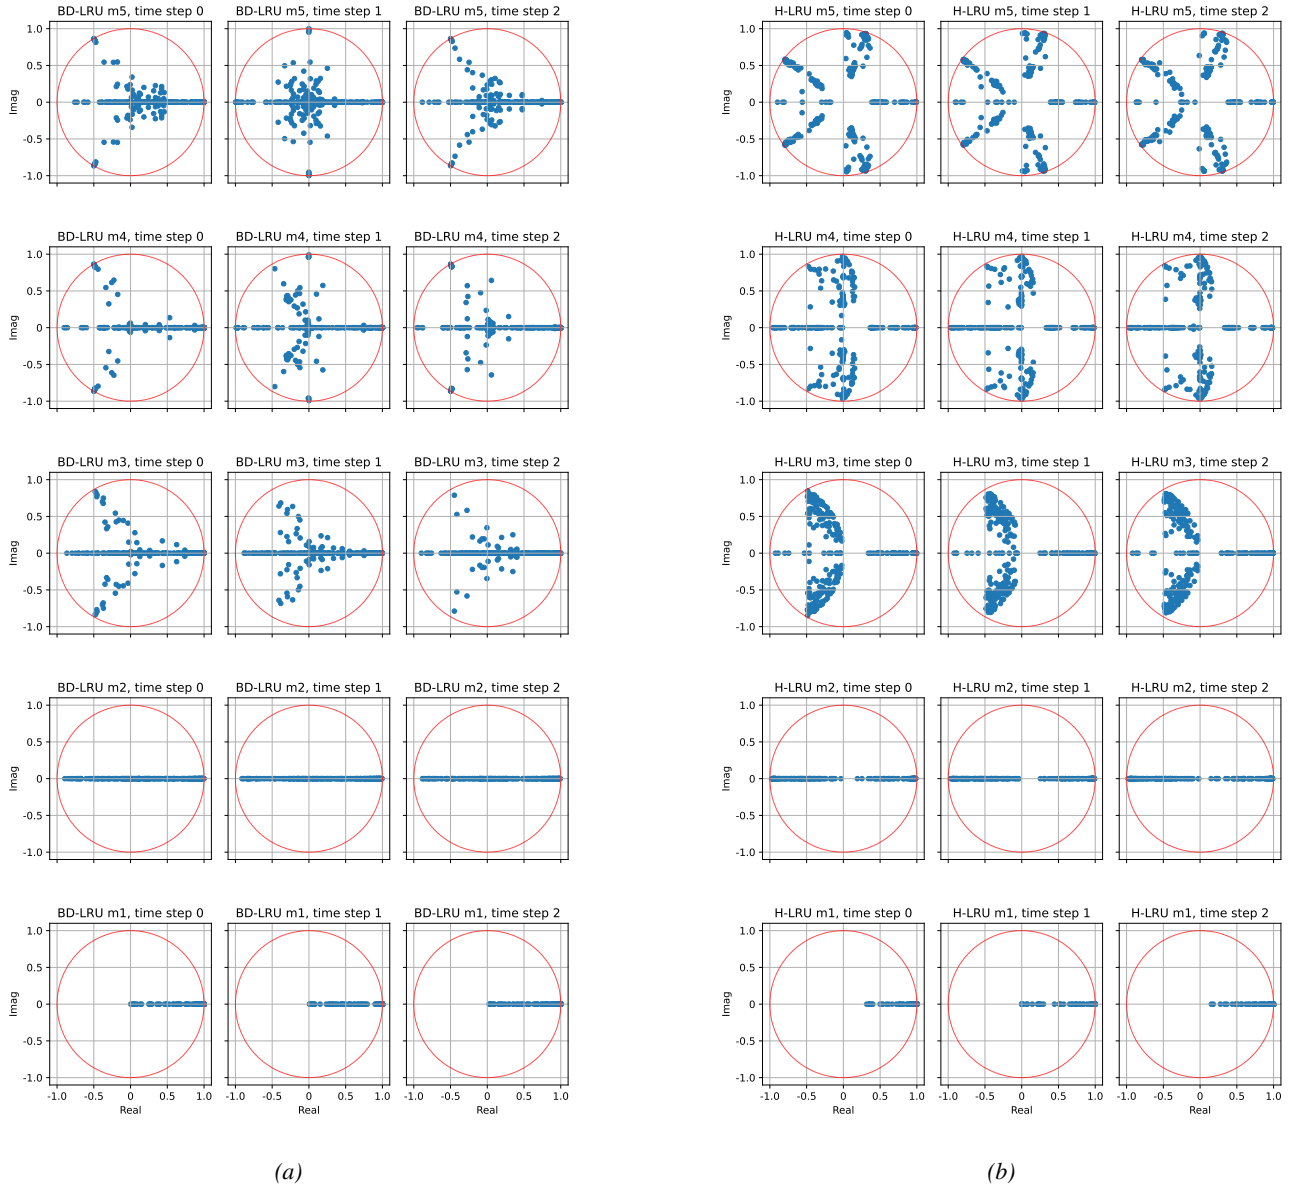

(a) *(b)*

Figure 6. Eigenvalues of transition matrices in H-LRU/BD-LRU on $S5$ dataset. (a) BD-LRU with softmax normalization. (b) H-LRU with softmax normalization. Each subplot corresponds to a specific time step (horizontal axis) and block size (vertical axis). Note that as block size increases, the number of available symmetries increases as well.

B. Extended tables and additional figures

Models	S_3 (10k samples)	S_4 (50k)	S_5 (100k)
LSTM	**1.000 (1.000 ± 0.000)**	**1.000 (1.000 ± 0.000)**	**1.000 (1.000 ± 0.000)**
Mamba2 c4 s2	0.640 (0.622 ± 0.020)	0.423 (0.402 ± 0.033)	0.268 (0.242 ± 0.028)
Mamba2 c4 s4	0.626 (0.609 ± 0.022)	0.432 (0.396 ± 0.035)	0.278 (0.237 ± 0.042)
Mamba2 c4 s8	0.657 (0.597 ± 0.055)	0.429 (0.400 ± 0.025)	0.261 (0.239 ± 0.033)
Deltanet[-1,1] nh8	**1.000 (0.752 ± 0.223)**	0.507 (0.456 ± 0.061)	0.140 (0.136 ± 0.004)
Deltaproduct$_4$[-1,1] nh8	**1.000 (1.000 ± 0.000)**	**1.000 (0.868 ± 0.132)**	0.140 (0.137 ± 0.004)
BD-LRU m1	0.563 (0.563 ± 0.000)	0.343 (0.323 ± 0.020)	0.215 (0.199 ± 0.018)
BD-LRU m2	**1.000 (1.000 ± 0.000)**	0.699 (0.651 ± 0.067)	0.338 (0.277 ± 0.057)
BD-LRU m3	**1.000 (1.000 ± 0.000)**	**1.000 (0.930 ± 0.060)**	0.484 (0.346 ± 0.148)
BD-LRU m4	**1.000 (1.000 ± 0.000)**	**1.000 (0.969 ± 0.031)**	0.882 (0.531 ± 0.308)
BD-LRU m5	**1.000 (1.000 ± 0.000)**	**1.000 (0.914 ± 0.086)**	**1.000 (0.652 ± 0.302)**
BD-LRU m6	**1.000 (1.000 ± 0.000)**	**1.000 (1.000 ± 0.000)**	**1.000 (0.923 ± 0.077)**
BD-LRU m8	**1.000 (1.000 ± 0.000)**	**1.000 (1.000 ± 0.000)**	**1.000 (0.821 ± 0.176)**
BD-LRU m16	**1.000 (1.000 ± 0.000)**	**1.000 (1.000 ± 0.000)**	**1.000 (1.000 ± 0.000)**
H-LRU m1	0.573 (0.559 ± 0.013)	0.347 (0.329 ± 0.016)	0.230 (0.194 ± 0.045)
H-LRU m2	0.597 (0.583 ± 0.012)	0.367 (0.354 ± 0.015)	0.260 (0.236 ± 0.031)
H-LRU m3	0.610 (0.589 ± 0.030)	0.399 (0.381 ± 0.016)	0.321 (0.301 ± 0.033)
H-LRU m4	0.624 (0.604 ± 0.018)	0.414 (0.390 ± 0.021)	0.337 (0.307 ± 0.031)
H-LRU m5	0.622 (0.609 ± 0.019)	0.448 (0.424 ± 0.025)	0.377 (0.337 ± 0.065)
H-LRU m6	0.630 (0.596 ± 0.053)	0.446 (0.431 ± 0.016)	0.386 (0.311 ± 0.065)
H-LRU m8	0.641 (0.612 ± 0.029)	0.494 (0.451 ± 0.037)	0.390 (0.351 ± 0.062)
H-LRU m16	0.659 (0.583 ± 0.066)	0.512 (0.459 ± 0.065)	0.393 (0.331 ± 0.063)

Table 3. Model performance on permutation composition tasks for different datasets of different sizes: S_3 (10k training samples), S_4 (50k training samples), S_5 (100k training samples). The accuracy values reflect the impact of window size (H-LRU) and block size (BD-LRU), both denoted by m. We note that BD-LRU performance improves with block size, demonstrating strong sample efficiency by solving the tasks even given limited training data. All models are single-layer configurations with a maximum overall hidden dimension of 6144. The presented results are the best test accuracies (mean ± std in brackets across 5 seeds). Deltaproduct and Deltanet are reported with 8 heads (nh) and [-1,1] denotes the range of eigenvalues. Mamba2 are reported with convolution size (c) and state space extension (s).

Models	S3 (250)	S4 (3k)
LSTM	0.316 (0.298 ± 0.016)	0.368 (0.297 ± 0.093)
Mamba2 c4 s2	0.284 (0.277 ± 0.006)	0.115 (0.113 ± 0.003)
Mamba2 c4 s4	0.284 (0.278 ± 0.006)	0.115 (0.105 ± 0.009)
Mamba2 c4 s8	0.281 (0.279 ± 0.002)	0.119 (0.115 ± 0.004)
Deltanet[-1,1] nh8	0.264 (0.259 ± 0.005)	0.138 (0.117 ± 0.018)
Deltaproduct_4[-1,1] nh8	0.279 (0.273 ± 0.008)	0.136 (0.136 ± 0.000)
BD-LRU m1	0.381 (0.347 ± 0.034)	0.218 (0.204 ± 0.012)
BD-LRU m2	0.488 (0.459 ± 0.032)	0.356 (0.265 ± 0.079)
BD-LRU m3	**1.000 (0.998 ± 0.002)**	0.426 (0.382 ± 0.038)
BD-LRU m4	**1.000 (0.792 ± 0.208)**	**1.000 (0.781 ± 0.219)**
BD-LRU m5	**1.000 (0.812 ± 0.188)**	**1.000 (0.582 ± 0.363)**
BD-LRU m6	**1.000 (1.000 ± 0.000)**	**1.000 (0.849 ± 0.151)**
BD-LRU m8	0.995 (0.962 ± 0.038)	**1.000 (0.737 ± 0.263)**
BD-LRU m16	**1.000 (0.959 ± 0.041)**	**1.000 (0.792 ± 0.208)**
H-LRU m1	0.363 (0.351 ± 0.014)	0.212 (0.203 ± 0.008)
H-LRU m2	0.311 (0.310 ± 0.001)	0.194 (0.182 ± 0.016)
H-LRU m3	0.316 (0.298 ± 0.016)	0.206 (0.179 ± 0.023)
H-LRU m4	0.309 (0.292 ± 0.015)	0.191 (0.175 ± 0.013)
H-LRU m5	0.318 (0.300 ± 0.015)	0.194 (0.179 ± 0.014)
H-LRU m6	0.283 (0.281 ± 0.003)	0.168 (0.165 ± 0.003)
H-LRU m8	0.281 (0.276 ± 0.004)	0.166 (0.164 ± 0.002)
H-LRU m16	0.260 (0.255 ± 0.004)	0.159 (0.157 ± 0.001)

Table 4. Extension of Table 3 for permutation composition tasks that have limited number of training samples: S_3 (250 training samples), S_4 (3k training samples). We note that BD-LRU performance improves with block size, demonstrating strong sample efficiency by solving the tasks even given limited training data.

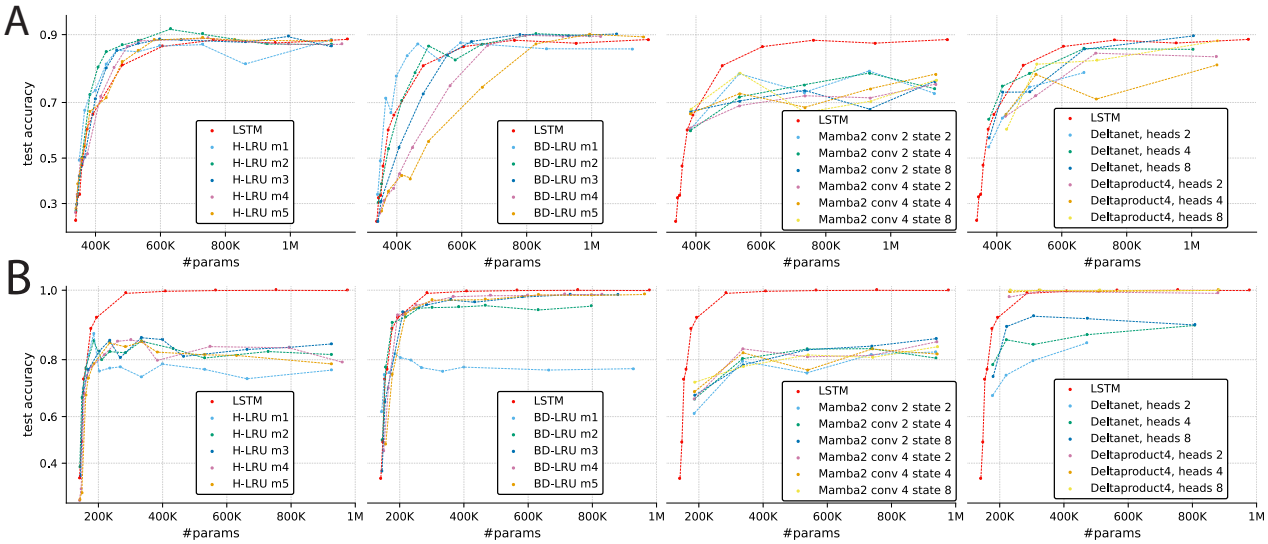

Figure 7. Performance of different single-layer models as a function of the hidden size (up to 6k overall hidden dimension) in the compression task (**A**) and the selective copying task (**B**). Results are shown for different window sizes (H-LRU) and block sizes (BD-LRU) m. We compare our networks with different configurations of Mamba2 (with two convolution kernel sizes (2,4) and several values of the state space expansion factor (2,4,8)). For comparison to low-rank models, we also include DeltaNet and DeltaProduct with 4 Householder transforms with different number of heads (2,4,8).

Models	In-Context Recall Tasks	Selective Copy Tasks	Compression Tasks
LSTM	**1.000 (1.000 ± 0.000)**	**1.000 (0.998 ± 0.002)**	0.749 (0.697 ± 0.092)
Mamba2 c2 s2	**1.000 (0.970 ± 0.030)**	0.791 (0.743 ± 0.051)	0.623 (0.508 ± 0.142)
Mamba2 c2 s4	**1.000 (0.988 ± 0.012)**	0.812 (0.786 ± 0.032)	0.721 (0.582 ± 0.131)
Mamba2 c2 s8	**1.000 (1.000 ± 0.000)**	0.830 (0.799 ± 0.032)	0.606 (0.535 ± 0.069)
Mamba2 c4 s2	**1.000 (1.000 ± 0.000)**	0.798 (0.753 ± 0.059)	0.623 (0.528 ± 0.108)
Mamba2 c4 s4	**1.000 (1.000 ± 0.000)**	0.800 (0.756 ± 0.060)	0.637 (0.548 ± 0.098)
Mamba2 c4 s8	**1.000 (1.000 ± 0.000)**	0.803 (0.742 ± 0.073)	0.634 (0.561 ± 0.065)
Deltanet[-1,1] nh2	**1.000 (1.000 ± 0.000)**	0.796 (0.744 ± 0.045)	0.709 (0.470 ± 0.180)
Deltanet[-1,1] nh4	**1.000 (1.000 ± 0.000)**	0.879 (0.806 ± 0.065)	0.718 (0.654 ± 0.068)
Deltanet[-1,1] nh8	**1.000 (1.000 ± 0.000)**	0.902 (0.870 ± 0.040)	0.765 (0.665 ± 0.110)
Deltaproduct_4[-1,1] nh2	**1.000 (1.000 ± 0.000)**	0.996 (0.985 ± 0.011)	0.753 (0.648 ± 0.090)
Deltaproduct_4[-1,1] nh4	**1.000 (1.000 ± 0.000)**	**1.000 (0.996 ± 0.004)**	0.720 (0.645 ± 0.096)
Deltaproduct_4[-1,1] nh8	**1.000 (1.000 ± 0.000)**	**1.000 (0.998 ± 0.001)**	0.735 (0.655 ± 0.073)
BD-LRU m1	0.779 (0.721 ± 0.036)	0.835 (0.815 ± 0.016)	0.725 (0.580 ± 0.119)
BD-LRU m2	**1.000 (1.000 ± 0.000)**	0.964 (0.954 ± 0.009)	0.760 (0.645 ± 0.129)
BD-LRU m3	**1.000 (1.000 ± 0.000)**	0.984 (0.975 ± 0.006)	0.764 (0.642 ± 0.126)
BD-LRU m4	**1.000 (1.000 ± 0.000)**	0.985 (0.976 ± 0.006)	**0.786 (0.748 ± 0.047)**
BD-LRU m5	**1.000 (1.000 ± 0.000)**	0.985 (0.978 ± 0.005)	0.782 (0.732 ± 0.056)
BD-LRU m6	**1.000 (1.000 ± 0.000)**	0.983 (0.977 ± 0.004)	0.775 (0.665 ± 0.101)
BD-LRU m8	**1.000 (1.000 ± 0.000)**	0.996 (0.946 ± 0.032)	0.748 (0.685 ± 0.079)
BD-LRU m16	**1.000 (1.000 ± 0.000)**	0.999 (0.991 ± 0.009)	0.725 (0.635 ± 0.121)
H-LRU m1	0.783 (0.721 ± 0.038)	0.847 (0.813 ± 0.024)	0.758 (0.650 ± 0.133)
H-LRU m2	0.998 (0.993 ± 0.004)	0.853 (0.827 ± 0.023)	0.772 (0.692 ± 0.061)
H-LRU m3	**1.000 (1.000 ± 0.000)**	0.853 (0.830 ± 0.018)	0.773 (0.673 ± 0.136)
H-LRU m4	**1.000 (1.000 ± 0.000)**	0.845 (0.823 ± 0.019)	0.774 (0.736 ± 0.038)
H-LRU m5	**1.000 (1.000 ± 0.000)**	0.835 (0.823 ± 0.009)	0.776 (0.698 ± 0.070)
H-LRU m6	**1.000 (0.954 ± 0.026)**	0.817 (0.805 ± 0.009)	0.775 (0.674 ± 0.129)
H-LRU m8	**1.000 (0.976 ± 0.024)**	0.807 (0.776 ± 0.022)	0.766 (0.694 ± 0.060)
H-LRU m16	**1.000 (0.996 ± 0.003)**	0.679 (0.611 ± 0.044)	0.708 (0.610 ± 0.114)

Table 5. Performance on the in-context recall, selective copying and compression tasks. The presented results are the average of best test accuracies (mean ± std in brackets across 5 seeds) across four configurations of the corresponding synthetic dataset with different vocabulary sizes, sequence lengths and number of training examples. Results are shown for different window (H-LRU) and block sizes (BD-LRU) m. Deltaproduct and Deltanet are reported for different number of heads (nh) and [-1,1] denotes the range of eigenvalues. Mamba2 are reported variants are reported with different convolution sizes (c) and state space extension (s). Note that overall performance of our models consistently improves with window/block size up to approximately 3–5, after which the gains saturate or exhibit slight degradation. All models are single-layer configurations with a maximum overall hidden dimension of 6144.

C. Experiments

Motivation for synthetic tasks. The sequence modeling capabilities of modern neural architectures are often evaluated through large-scale experiments involving models with billions of parameters and trained on trillions of tokens (Kaplan et al., 2020; Waleffe et al., 2024). However, recent studies have shown that many of these capabilities can be assessed using smaller models trained on carefully designed synthetic datasets which target specific tasks that are crucial for general sequence modeling (Arora et al., 2023; Poli et al., 2024).

First, the well-established equivalence between lossless compression and probabilistic modeling suggests that models that compress well also generalize well (Shannon, 1948; Hutter, 2005). Indeed, recent work shows that there is a clear connection between language modeling and compression (Gu, 2025), although with some difference in scaling laws (Delétang et al., 2023). In light of this, we include in our evaluation a task that tests the efficiency of temporal information integration, the auto-encoding compression task from (Poli et al., 2024).

Next, general sequence modeling requires not only the ability to develop a fixed prediction algorithm, but also the capacity to adapt dynamically to changes within the input context. Such *in-context abilities* have been extensively studied and have been suggested to explain the success of the Transformer architecture (Olsson et al., 2022). To benchmark this basic capability, we choose the selective copying and associative recall tasks that have been shown to be good indicators of the in-context abilities of sequence models (Arora et al., 2023; Poli et al., 2024), as well as indicators of downstream capabilities (Waleffe et al., 2024).

Synthetic token manipulation tasks. We benchmarked our architectures using the Mechanistic Architecture Design (MAD) framework (Poli et al., 2024), a framework for efficient model evaluation and prototyping. The MAD protocol is motivated by the challenge of predicting how architectural choices impact performance at scale. The working hypothesis of MAD is that an architecture's macroscopic scaling behavior can be effectively predicted by its performance on a set of microscopic, mechanistic tasks.

The benchmark consists of a diverse suite of sequence modeling challenges designed to test core token manipulation capabilities. By evaluating models at a small, fixed computational scale, MAD produces a relative ranking of architectures that has been shown to be predictive of their compute-optimal performance in large-scale language modeling (Poli et al., 2024). This approach not only approximates scaling outcomes, but also provides valuable insights into the compositional skills and failure modes of a given design.

In particular, we utilize three tasks from the MAD framework:

- **Compression task**. Models are tasked to compress a random sequence of input tokens into a single aggregation token. Then, this aggregation token is passed through an encoder MLP, the output of which is used to reconstruct the original sequence via a decoder MLP. All models were tested using a standard encoder-decoder architecture (Embedding, Tested Model, MLP Encoder, MLP Decoder).

- **Selective copying task**. Models are tasked with copying tokens from one position of an input sequence to a later position of the sequence, while ignoring irrelevant noise tokens that are randomly inserted into the sequence. This task is designed to evaluate the ability of a model to perform selective temporal integration in the specific order of occurrence in the sequence. All models were tested using a standard decoder-only architecture (Embedding, Tested Model, MLP Decoder).

- **Associative recall task**. Models are presented with an input sequence of key-value pairs and tasked with retrieving all values from the input sequence associated with the presented keys. This task tests the ability of a model to adaptively retrieve information depending on the established in-context associations. All models were tested using a standard decoder-only architecture (Embedding, Tested Model, MLP Decoder).

In our experiments, each model was evaluated across four configurations: a baseline (vocabulary size: 16, sequence length: 64, training examples: 20,000) and three variations designed to probe specific failure modes. These variations all use the same base parameters, but independently (i) increase the vocabulary size to 32, (ii) extend the sequence length to 128, or (iii) reduce the training set to 10,000 examples to test vocabulary handling, long-range capabilities, and sample efficiency, respectively.

Synthetic permutation tasks. In our experiments, we employ synthetic datasets derived from the symmetric permutation groups S_n, which denotes the group of all possible permutations of n elements. These groups provide a natural hierarchy of complexity: S_2 contains only two permutations and is fully commutative, making it relatively simple to model. In contrast, groups with $n \geq 3$ (e.g., S_3, S_4, S_5) are non-commutative, and their size grows factorially with n, which rapidly increases the difficulty of learning the underlying structure. For instance, S_3, with six elements, is the smallest non-commutative group. Geometrically, S_3 can be interpreted as the group of symmetries of an equilateral triangle, including both rotations and reflections. The complexity increases substantially with S_4, which contains 24 elements and corresponds to the full symmetry group of a regular tetrahedron. S_4 introduces more intricate subgroup structures and non-trivial normal subgroups. Extending further, S_5 has 120 elements and is the first symmetric group that is not solvable.

We assess model performance on the synthetic permutation group task from (Merrill et al., 2024), which is designed to probe state-tracking and generalization to complex structures. Using their toolbox, we generated datasets for the symmetric groups S_3, S_4, and S_5 with a fixed sequence length of 16. To evaluate sample efficiency, we created five distinct data configurations: S_3 (10k and 250 examples), S_4 (50k and 3k examples), and S_5 (100k examples). The S_5 setting is particularly data-limited compared to the multi-million-example setups used in previous studies (Siems et al., 2025). All models were tested using a standard decoder-only architecture (Embedding, Tested Model, MLP Decoder), consistent with the MAD benchmark protocol.

Training details on synthetic tasks. All models were implemented in PyTorch (Paszke et al., 2019). For training, we follow the experimental settings of the MAD framework. All models are trained with the AdamW optimizer (Loshchilov & Hutter, 2017) with parameters $\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^{-8}$ and a cosine scheduler (Loshchilov & Hutter, 2016) (minimum LR: 0.00001), with the initial learning rate selected from 0.001, 0.0005, 0.0001. The final reported metric is the best test accuracy across all three learning rate configurations and five runs with distinct random seeds. For training, we used NVIDIA A100 and NVIDIA H100, while we used NVIDIA H100 for benchmarking the best throughput across models. The estimated total compute required to perform all experiments on synthetic tasks is 16 000 GPU hours.

Training details on language modeling. We conduct our experiments on the well-established FineWeb dataset (Penedo et al., 2024) using the PlainLM training setup (Ajroldi, 2024). All models are trained with the AdamW optimizer (Loshchilov & Hutter, 2017) with parameters $\beta_1 = 0.9, \beta_2 = 0.95, \epsilon = 10^{-8}$ and a cosine scheduler (Loshchilov & Hutter, 2016). All models are trained on a single NVIDIA H100 GPU, or a total of approximately 2,000 GPU hours.

Table 6. Summary of computational costs for hidden state updates.

Architecture	FLOPs per recurrent step	Implementation complexity
LSTM	$8H^2 + 25H$	$O(TH^2)$
H-LRU	$2Nm + 2N$	$O(Nm^3 log(T))$
BD-LRU	$2Nm^2 + 2N$	$O(Nm^3 log(T))$
Mamba2	$2NS$	$O(T(N^2 + NS))$
Deltanet	$N_h(4Nr + 4N)$	$O(TCN + TN^2)$
Deltaproduct4	$H_n N_h(4Nr + 4N)$	$O(H_n(TCN + TN^2))$

D. Computational complexity

This section provides a breakdown of the Floating Point Operations (FLOPs) required for hidden-to-hidden state transition in the recurrent architectures discussed. For this breakdown, we define the number of blocks for BD-LRU and H-LRU as N and H denotes overall hidden dimension, $m \cdot N = H$ as previously. The sequence length is denoted as T. For Mamba2, the state expansion factor is denoted by S. In Deltanet and Deltaproduct4, N_h denotes the number of heads, C denotes the number of chunks in the Deltanet implementation, H_n denotes the number of Householder transformations, and $r = 1$ denotes low rank. The calculations focus on the recurrence mechanism, omitting additional components like the input projections or gating, as they can be precomputed in advance. A multiply-add operation is counted as 2 FLOPs.

E. Proof of Proposition 1.

First, note that stability directly follows from the induced norm bound. We can reason blockwise: assuming $\sum_j |(\mathcal{A}_t^k)_{i,j}| \leq 1$ implies that the eigenvalues of state-transition matrix $|\lambda_{i,t}^k| \leq 1$. Therefore, the product of such matrices will result in dynamical stability.

Next, by block-diagonality, it is sufficient to show that for all $k \in [1, N]$, $\|\mathbf{h}_T^k\|_\infty \leq \max_{t\in[0,T]} \|\mathbf{v}_t^k\|_\infty$. Let $h_{i,t}^k$ be the i-th coordinate of the generic k-th block hidden state $\mathbf{h}_t^k$ at time t.

$$\begin{bmatrix} h_{1,t}^k \\ \vdots \\ h_{m,t}^k \end{bmatrix} = \begin{bmatrix} a_{1,1,t}^k & \cdots & a_{1,m-1,t}^k & a_{1,m,t}^k \\ a_{2,1,t}^k & \cdots & a_{2,m-1,t}^k & a_{2,m,t}^k \\ \vdots & \ddots & \vdots & \vdots \\ a_{m,1,t}^k & \cdots & a_{m,m-1,t}^k & a_{m,m,t}^k \end{bmatrix} \times \begin{bmatrix} h_{1,t-1}^k \\ \vdots \\ h_{m,t-1}^k \end{bmatrix} + \begin{bmatrix} a_{1,0,t}^k \\ \vdots \\ a_{m,0,t}^k \end{bmatrix} \odot \begin{bmatrix} v_{1,t}^k \\ \vdots \\ v_{m,t}^k \end{bmatrix}. \tag{5}$$

Hence,

$$h_{i,t}^k = \sum_{j=1}^m a_{i,j,t}^k h_{j,t-1}^k + a_{i,0,t}^k v_{i,t}^k. \tag{6}$$

It is then clear that by subadditivity of the absolute value,

$$|h_{i,t}^k| \leq \sum_{j=1}^m |a_{i,j,t}^k| \cdot |h_{j,t-1}^k| + |a_{i,0,t}^k| \cdot |v_{i,t}^k|. \tag{7}$$

Hence, by collecting the non-coefficient terms, we find a further upper bound

$$|h_{i,t}^k| \leq \left(\sum_{j=1}^m |a_{i,j,t}^k| + |a_{i,0,t}^k| \right) \cdot \max \left[|v_{i,t}^k|, \max_{j\in[1,m]} |h_{j,t-1}^k| \right]. \tag{8}$$

By hypothesis, $\sum_{j=1}^m |a_{i,j,t}^k| + |a_{i,0,t}^k| = \sum_j |(\mathcal{A}_t^k)_{i,j}| \leq 1$, and hence we conclude that

$$|h_{i,t}^k| \leq \max \left[|v_{i,t}^k|, \max_{j\in[1,m]} |h_{j,t-1}^k| \right]. \tag{9}$$

At this point, we can finalize the proof by induction. We want to show that $\|\mathbf{h}_T^k\|_\infty \leq \max_{t\in[0,T]} \|\mathbf{v}_t^k\|_\infty$. Let us start from $T = 1$. Since $h_{i,0}^k = 0$ for all $i \in [1, m]$, we have

$$h_{i,1}^k = a_{i,0,t}^k v_{i,1}^k, \tag{10}$$

hence, again because $\sum_j |(\mathcal{A}_0^k)_{i,j}| \leq 1$, $|h_{i,1}^k| \leq |v_{i,1}^k|$, we can conclude that $\|\mathbf{h}_1^k\|_\infty \leq \|\mathbf{v}_1^k\|_\infty$. Let us then assume by induction that $\|\mathbf{h}_{T-1}^k\|_\infty \leq \max_{t\in[0,T-1]} \|\mathbf{v}_t^k\|_\infty$. Recall that by Equation 9,

$$|h_{i,t}^k| \leq \max \left[|v_{i,t}^k|, \max_{j\in[1,m]} |h_{j,t-1}^k| \right] \tag{11}$$

$$= \max \left[|v_{i,t}^k|, \|\mathbf{h}_{t-1}^k\|_\infty \right]. \tag{12}$$

Hence,

$$\|\mathbf{h}_t^k\|_\infty = \max_{j\in[1,m]} |h_{j,t}^k| \tag{13}$$

$$\leq \max_{j\in[1,m]} \max\left[|v_{i,t}^k|, \|\mathbf{h}_{t-1}^k\|_\infty\right] \tag{14}$$

$$= \max\left[\max_{j\in[1,m]} |v_{i,t}^k|, \|\mathbf{h}_{t-1}^k\|_\infty\right] \tag{15}$$

$$\leq \max\left[\|\mathbf{v}_t^k\|_\infty, \max_{t\in[0,T-1]} \|\mathbf{v}_t^k\|_\infty\|\right] \tag{16}$$

$$= \max_{t\in[0,T]} \|\mathbf{v}_t^k\|_\infty, \tag{17}$$

where in the second-last line we used the induction hypothesis.

F. Selectivity ablation

To isolate and quantify the contribution of selectivity, we conducted an ablation study. In this analysis, the input-dependent selective gates in both the H-LRU and BD-LRU architectures were replaced with data-invariant, learnable parameters.

As hypothesized, the non-selective variants exhibited a significant performance degradation compared to their selective counterparts on our synthetic benchmark. On tasks requiring dynamic token manipulation—such as in-context recall, selective copying, and permutation composition—the non-selective models failed to achieve meaningful performance. For these tasks, increasing the window or block size yielded no discernible improvement, confirming the necessity of selectivity.

However, the results on the compression task were more nuanced, see Fig. 8. We observed that our proposed $L1$ normalization scheme enabled the non-selective models to improve with larger block and window sizes, albeit at a lower rate than their selective analogs.

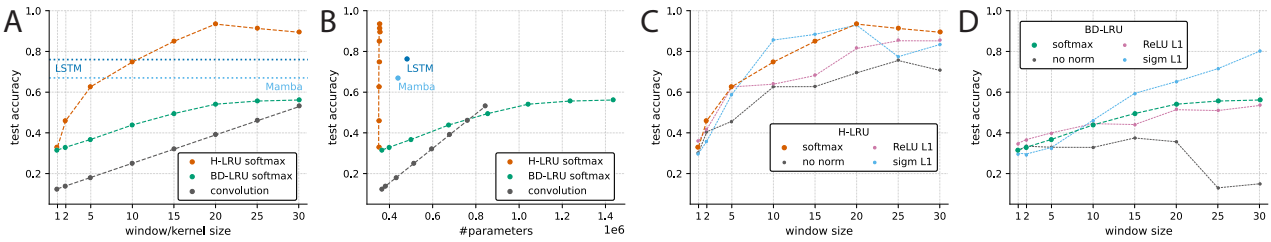

Figure 8. Scaling analysis of non-selective models on the compression task. **A**. Performance as a function of window size m of non-selective higher-order LRU (H-LRU) and block size m of block diagonal LRU (BD-LRU). For the convolutional baseline, the performance presented as a function of kernel size. **B**. The same results plotted against parameter count. Note that scaling with window size of non-selective H-LRU demonstrates extreme parameter efficiency, resulting in a nearly vertical trajectory on the plot. **C**. Comparison of scaling properties between different parameterizations for H-LRU. **D**. Comparison of scaling properties between different parameterizations for BD-LRU.

To highlight the advantages of recurrent architectures, we used a convolution layer as a baseline. This model is limited to explicit, local time mixing within its kernel, in contrast to the implicit and unbounded temporal integration provided by a hidden state. Our experiments showed that H-LRU decisively outperforms the convolution on the compression task. This demonstrates the critical role of recurrent state mixing for tasks requiring efficient long-range temporal reasoning. Furthermore, the non-selective H-LRU with large window sizes ($m > 15$) demonstrated strong performance, surpassing the LSTM and Mamba2 baselines and even approaching the performance of our selective models. This finding highlights the powerful inductive bias of the higher-order recurrence for parameter-efficient compression.

In contrast, the non-selective BD-LRU performed poorly on the compression task, only marginally surpassing the convolution baseline. Interestingly, for this non-selective variant, the sigmoidal $L1$ normalization outperformed softmax normalization, highlighting a difference in how these schemes interact with selective versus fixed parameterizations.

In addition, when we analyzed H-LRU with minimal point-wise selective gates that don't mix channel dimensions, we

observed a very moderate improvement in compression task. This indicates that not only selectivity itself but also density of selectivity in gates plays an important role in improving networks' expressivity.

While the overall performance of these non-selective models is modest, their parameter efficiency can become advantageous in resource-constrained settings. Given the strong compression results of the non-selective H-LRU, we hypothesize that such models could be optimized for use as highly efficient embedding layers, a direction we leave for future research.

G. Relation between the expressivity of LRUs and State Space Duality

Recently, it has been shown that there is a direct correspondence between state space models, the Transformer architecture and structured attention matrices (Dao & Gu, 2024). Following this approach, we can reformulate the general LRU as a general discrete time SSM

$$\begin{aligned} \mathbf{h}_t &= \mathbf{A}_t \times \mathbf{h}_{t-1} + \mathbf{B}_t \times \mathbf{v}_t \\ \mathbf{y}_t &= \mathbf{C}_t \times \mathbf{h}_t. \end{aligned} \tag{18}$$

Here, we consider the general case of SSMs in which mixing matrices $\mathbf{C}_t, \mathbf{A}_t, \mathbf{B}_t$ are dense matrices. We note that although state space models are commonly defined in continuous time, they have to be discretized for implementation, at which point they conform to the discrete form described by Eq. 18. In this study, we effectively ignored the role of $\mathbf{C}_t$, but it can be introduced without affecting the validity of our arguments.

Following the approach to reformulating state space models (SSMs) as attention mechanisms, the architecture given in Eq. 18 can be expressed in block matrix representation assuming a fixed sequence length T:

$$\begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_3 \\ \vdots \\ \mathbf{y}_T \end{bmatrix} = \begin{bmatrix} \mathbf{C}_1\mathbf{B}_1 & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{C}_2\mathbf{A}_1\mathbf{B}_1 & \mathbf{C}_2\mathbf{B}_2 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{C}_3\mathbf{A}_2\mathbf{A}_1\mathbf{B}_1 & \mathbf{C}_3\mathbf{A}_2\mathbf{B}_2 & \mathbf{C}_3\mathbf{B}_3 & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{C}_T \prod_{j=1}^{T} \mathbf{A}_j\mathbf{B}_1 & \mathbf{C}_T \prod_{j=2}^{T} \mathbf{A}_j\mathbf{B}_2 & \cdots & \cdots & \mathbf{C}_T\mathbf{B}_T \end{bmatrix} \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \mathbf{v}_3 \\ \vdots \\ \mathbf{v}_T \end{bmatrix}$$

If we abstract the details of SSMs matrices, we obtain the generalized attention formulation:

$$\begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} \overline{\mathbf{A}}_{1,1} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \overline{\mathbf{A}}_{2,1} & \overline{\mathbf{A}}_{2,2} & \mathbf{0} & \cdots & \mathbf{0} \\ \overline{\mathbf{A}}_{3,1} & \overline{\mathbf{A}}_{3,2} & \overline{\mathbf{A}}_{3,3} & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots \end{bmatrix} \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \mathbf{v}_3 \\ \vdots \end{bmatrix}. \tag{19}$$

Importantly, elements $\overline{\mathbf{A}}_{k,l}$ of the block attention matrix are also matrices in this representation. According to State Space Duality (Dao & Gu, 2024), both the attention in Transformers and diagonal SSMs result in diagonal matrices $\overline{\mathbf{A}}_{k,l}$. So, their architecture allows for efficient parallelization as it separates temporal mixing from channel mixing.

In contrast to diagonal SSMs and LRUs, both H-LRU and BD-LRU architectures result in block-diagonal matrices $\overline{\mathbf{A}}_{k,l}$, allowing richer but block-limited channel mixing inside the generalized block attention matrix in Eq. 19. Such channel mixing allows for the state mixing patterns that are not accessible to one layer of diagonal LRU or SSMs. Although the channel mixing in H-LRU is more expressive than the one in a diagonal LRU, it is still more restricted compared to BD-LRU (it is equivalent to mixing only in one row of a block-diagonal matrix), placing the expressivity of H-LRU between diagonal LRU and BD-LRU. Notably, if we extend SSMs with higher-order or block-diagonal structures, their expressivity would lag behind analogous LRUs due to the restrictions on mixing patterns imposed by the chosen discretization scheme. Overall, the generalized block attention formulation 19 reveals that diagonal, higher-order, block diagonal and dense variants of LRUs and SSMs form a hierarchy of architectures, each providing access to increasingly complex state mixing patterns which result in increased expressivity.

H. Chomsky Hierarchy Tasks

The Chomsky hierarchy formalizes increasing levels of expressiveness and computational complexity of formal languages into several hierarchical classes (Chomsky, 1956; Delétang et al., 2022). Here, we tested several tasks from this hierarchy: Parity, Cycle Navigation, Modular Arithmetic with and without brackets. The Parity task requires computing whether given

binary string is even or not. Cycle Navigation requires computing the end position given a sequence of movements on a cycle of length 5. Modular Arithmetic tasks require computing the result modulo 5 for a given sequence of numbers in $(0, 1, 2, 3, 4)$ and operations in $(+, -, \cdot)$, with or without brackets.

In our experiments, we observe that similar to S_3 task, the Parity task can be solved by BD-LRU with access to negative eigenvalues ($m \geq 2$). For the Cycle Navigation task we obtain similar results as for S_5 task. BD-LRU is able to solve it starting from $m = 5$. Therefore, the results on these two tasks from Chomsky Hierarchy support our previously found advantage of BD-LRUs on permutation tasks.

Modular arithmetic tasks present a challenge for highly parallel Transformer architecture, often require grokking and have pure generalization (Gromov, 2023). In contrast, it has been shown that the sequential nature of state mixing in RNNs has a strongly beneficial bias for arithmetic-like induction (Merrill & Sabharwal, 2023). However, both our linear variants and other modern LRNNs struggle with such arithmetic tasks (Siems et al., 2025), supporting the idea that the nonlinearity of state transitions is crucial in such tasks (Chang & Bisk, 2024). In our experiments, we found that BD-LRUs were able to solve Modular Arithmetic without brackets, while the version with brackets remained challenging, similar to other RNNs.

Models	cycle nav	mod arith no brack	mod arith w brack	parity
LSTM	**1.000**	0.976	0.663	**1.000**
BD-LRU m1	0.434	0.370	0.370	0.512
BD-LRU m2	0.425	0.493	0.417	**1.000**
BD-LRU m3	0.597	0.546	0.434	**1.000**
BD-LRU m4	0.608	0.459	0.435	**1.000**
BD-LRU m5	**1.000**	0.525	0.422	**1.000**
BD-LRU m6	**1.000**	0.433	0.440	**1.000**
BD-LRU m8	**1.000**	0.553	0.395	**1.000**
BD-LRU m16	**1.000**	**1.000**	0.448	**1.000**

Table 7. Accuracy on Chomsky hierarchy tasks.