arXiv:2611.24809v1 [cs.CL] 18 Nov 2026

On-Policy Morphing into Hybrid Attention Models

Disen Lan [1 2 *] **Xuanda Wang** [2] **Jianbin Zheng** [2] **Huafeng Kuang** [2] **Wei Tang** [2] **Yu Cheng** [3]

Abstract

Hybrid attention models combine full and linear attention to achieve a better quality–efficiency trade-off. Converting pretrained Transformers into such models requires identifying which layers should retain full attention under a prescribed budget. We identify two *off-policy* mismatches in FlashMorph's layer selection: (i) continuous branch interpolation and teacher-forced hidden states differ from discrete hybrid computation; and (ii) frozen model weights fail to account for parameter adaptation during subsequent recovery. We introduce **FlashMorph2**, an *on-policy* layer-selection method that (I) evaluates sampled discrete hybrid architectures end to end using their own hidden states, and (II) jointly optimizes layer-selection scores and model weights under the recovery objective to reflect post-adaptation performance. Experiments across Qwen3 model scales and linear-attention variants show that FlashMorph2 consistently outperforms existing layer-selection methods, improving average scores over FlashMorph by $+2.42$ points on RULER, $+1.08$ on LongBench, and $+0.76$ on recall-intensive tasks, while completing layer selection in only 0.72–3.57 GPU hours across 0.6B–8B models, demonstrating high efficiency and scalability.

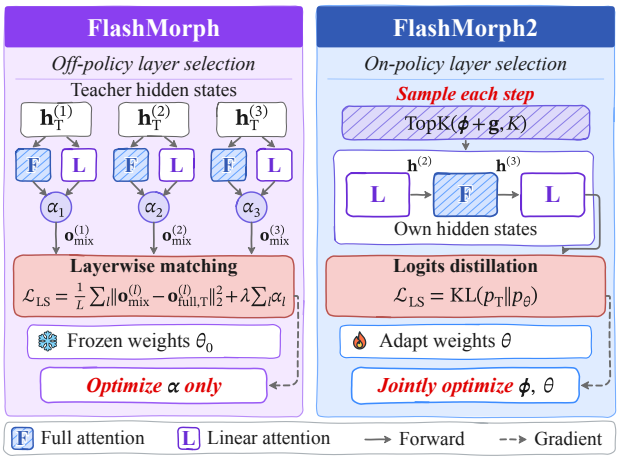


Figure 1. **Comparison between FlashMorph and FlashMorph2.** FlashMorph uses *off-policy* teacher-forced continuous gates with frozen model weights, whereas FlashMorph2 samples *on-policy* discrete hybrid architectures and jointly optimizes layer-selection scores and model weights.

1. Introduction

Hybrid attention models retain full attention in only a few layers and use linear attention elsewhere, preserving long-context modeling capabilities and precise retrieval (Arora et al., 2024a; Jelassi et al., 2024; Li et al., 2026a), while substantially reducing computation and KV-cache memory usage (Katharopoulos et al., 2020; Sun et al., 2025b;a). Owing to this favorable trade-off between performance and efficiency, hybrid attention has been adopted in a growing number of production-scale large language models (LLMs) (Li et al., 2025; Team et al., 2025; Cao et al., 2026; Team, 2026; Team et al., 2026). To avoid the prohibitive cost of pretraining from scratch, existing approaches reuse the weights of pretrained LLMs and replace a subset of full attention layers with linear attention, achieving Transformers-to-hybrid conversion. Model quality is then recovered through distillation and continued training (Kasai et al., 2021; Bick et al., 2024; Lan et al., 2025; Goldstein et al., 2025; Wang et al., 2024; Bick et al., 2025; Zhang et al., 2025). A central challenge in this conversion process is determining which full attention layers to retain under a given budget. As the only layers in the resulting hybrid model that can directly access the entire context, their placement largely determines how much of the original model's capability can be recovered (Li et al., 2026a; Chen et al., 2026; Lan et al., 2026).

Exhaustively evaluating all possible layer combinations is prohibitively expensive, motivating existing work to explore a variety of layer-selection strategies. Beyond fixed uniform interleaving patterns, PostNAS (Gu et al., 2026) performs supernet-based architecture search, KL-LS (Li et al., 2026a) ranks layers by the KL reduction from per-layer restoration and distillation, and HALO (Chen et al., 2026) uses replacement-based sensitivity scoring. More recently, FlashMorph (Lan et al., 2026) jointly optimizes layer-wise gates

[*]Work done at ByteDance Seed [1]Fudan University [2]ByteDance Seed [3]Nanyang Technological University. Correspondence to: Yu Cheng <yu.cheng@ntu.edu.sg>.

Preprint.

to learn a global hybrid configuration. Each layer interpolates between aligned full- and linear-attention branches through a continuous gate, which is then discretized under a given full-attention budget. FlashMorph preserves strong long-context recall and general benchmark performance while substantially reducing layer-selection cost.

Limitations. FlashMorph's layer-selection process still suffers from two forms of *off-policy* mismatch. **(L1)** ***Off-policy computation***: The forward computation used during layer selection differs from that of the candidate discrete model. Continuous mixing allows the full- and linear-attention branches to complement each other, while each layer operates on teacher-forced hidden states. At deployment, only one branch is retained per layer and hidden states are generated by the hybrid model itself. Consequently, complementarities observed under continuous mixing may not transfer to the final discrete architecture, and the resulting layer-importance estimates may fail to capture error accumulation and cross-layer compensation. **(L2)** ***Off-policy weights***: The weights used for layer selection are inherited from the original model and remain frozen during gate optimization, without adapting to the hybrid architecture induced by the selection policy. Layer selection therefore cannot account for the cross-layer co-adaptation of model weights during subsequent recovery, which may alter the layer-importance ranking observed before recovery.

Our Method. We introduce **FlashMorph2**, which improves upon FlashMorph by bringing the *on-policy* principle into its layer-selection process. To address **(L1)**, FlashMorph2 samples a discrete hybrid attention architecture that satisfies the prescribed full-attention budget and performs an end-to-end forward pass through the sampled architecture, allowing each layer to operate on hidden states produced by the model itself. To address **(L2)**, FlashMorph2 jointly trains model weights across these sampled hybrid attention architectures using the same optimization objective as in subsequent recovery training, allowing layer selection to account for performance after weight adaptation.

Results. Across Qwen3-0.6B and Qwen3-1.7B backbones and three linear-attention variants, FlashMorph2 consistently outperforms all layer-selection baselines, improving average scores over FlashMorph by $+2.42$ points on RULER, $+1.08$ on LongBench, and $+0.76$ on recall-intensive tasks while preserving language modeling and commonsense reasoning capabilities. Moreover, FlashMorph2 completes layer selection in only 0.72–3.57 GPU hours across 0.6B–8B models, running 8.4–$20.8\times$ faster than HALO and 560–$854\times$ faster than KL-LS.

Contributions. Our key contributions are as follows:

- We identify two forms of *off-policy* mismatch in FlashMorph's layer selection: (I) continuous interpolation and teacher-forced hidden states differ from deployed discrete hybrid computation, and (II) scores computed with frozen weights fail to reflect the optimal layer ranking after recovery training.
- We introduce FlashMorph2, an *on-policy* layer selection method for Transformer-to-hybrid conversion that performs end-to-end forward passes through sampled hybrid architectures under a full-attention budget, allowing selection to account for self-generated hidden states and weights adapted under the recovery objective.
- We evaluate FlashMorph2 across multiple Qwen3 model scales and linear-attention variants, demonstrating consistent improvements on long-context and recall-intensive benchmarks while maintaining high layer-selection efficiency and scalability.

2. Related Work

Layer Selection for Transformer-to-hybrid Conversion. Transformer-to-hybrid conversion requires selecting which full-attention layers to retain under a prescribed budget. A simple solution is fixed uniform interleaving (Lan et al., 2025; Goldstein et al., 2025), but it ignores layer-wise heterogeneity in pretrained models. PostNAS (Gu et al., 2026) trains a once-for-all supernet and uses task-guided beam search to determine full-attention layer placement, incurring substantial training and search costs; KL-LS (Li et al., 2026a) ranks layers by post-distillation KL loss after single-layer restoration; and HALO (Chen et al., 2026) ranks full-attention layers by the downstream performance changes induced by their individual replacement with linear-attention layers. These layerwise scoring methods incur substantial selection overhead from repeated single-layer evaluations and fail to account for interactions among layers in the resulting hybrid architecture. FlashMorph (Lan et al., 2026) jointly optimizes layerwise gate values, yielding more effective hybrid configurations at substantially lower selection cost. However, these methods remain *off-policy* due to mismatches in forward computation or model weights between the selection-time and recovered hybrid models. FlashMorph2 improves upon FlashMorph with *on-policy* layer selection that performs end-to-end forward passes through sampled discrete hybrid architectures while accounting for the effects of weight adaptation during recovery.

On-Policy Learning. Teacher forcing relies on ground-truth prefixes and is therefore *off-policy* with respect to the model's own generation policy; the resulting train-inference distribution mismatch can lead to exposure bias and error accumulation (Bengio et al., 2015). To mitigate this distribution shift, *on-policy* learning emphasizes obtaining supervision on states induced by the learner's own policy (Ross et al., 2011; Lin et al., 2020). Following this principle, *on-policy* distillation (OPD) trains students under teacher supervision on sequences generated by the students them-

selves (Lu & Lab, 2025; Li et al., 2026b; Ma et al., 2026). FlashMorph2 extends this idea to layer selection by sampling discrete hybrid architectures and estimating layer importance using self-generated hidden states and weights adapted under the recovery objective.

3. Preliminaries

Notation. Consider a pretrained Transformer (Vaswani et al., 2017) with L layers, indexed by $[L] = \{1, \dots, L\}$, and a token sequence $\mathbf{x}_{1:T} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$ of length T. Let $\mathbf{h}^{(l)}$ denote the input hidden states to layer l, and let $A^{(l)}_{\text{full}}$ and $A^{(l)}_{\text{lin}}$ denote the corresponding full-attention and linear-attention modules, respectively. At the position t, we define the query, key and value vectors (with layer and head indices omitted) as $\mathbf{q}_t = \mathbf{h}_t\mathbf{W}_Q$, $\mathbf{k}_t = \mathbf{h}_t\mathbf{W}_K$, and $\mathbf{v}_t = \mathbf{h}_t\mathbf{W}_V$, respectively, where $\mathbf{q}_t$, $\mathbf{k}_t$, $\mathbf{v}_t \in \mathbb{R}^{1\times d}$ and $\mathbf{W}_O$ denotes the output projection matrix.

Full attention. Under causal masking, full attention computes a weighted sum of value vectors in the current prefix, with weights determined by query-key similarities:

$$\mathbf{o}_{\text{full},t} = \frac{\sum_{i=1}^{t} \exp\left(\mathbf{q}_t\mathbf{k}_i^\top/\sqrt{d}\right)\mathbf{v}_i}{\sum_{j=1}^{t} \exp\left(\mathbf{q}_t\mathbf{k}_j^\top/\sqrt{d}\right)}\mathbf{W}_O. \tag{1}$$

This formulation incurs $\mathcal{O}(T^2)$ computational cost over the sequence and requires $\mathcal{O}(T)$ KV-cache memory to store past keys and values during autoregressive decoding.

Linear attention. Linear attention mitigates these inefficiencies by summarizing the sequence history in a fixed-size recurrent state. A broad class of linear-attention variants admits the following recurrent formulation:

$$\mathbf{o}_{\text{lin},t} = \mathbf{q}_t\mathbf{S}_t\mathbf{W}_O, \quad \mathbf{S}_t = \mathbf{S}_{t-1} + \mathbf{k}_t^\top\mathbf{v}_t, \tag{2}$$

where $\mathbf{S}_t \in \mathbb{R}^{d\times d}$ and $\mathbf{S}_0 = 0$. Different linear attention variants extend the recurrence with fixed state decay (Qin et al., 2023; 2024), data-dependent forget gates (Yang et al., 2023; Zhang et al., 2024; Du et al., 2026), or delta-rule updates (Yang et al., 2024; 2025b; Hu et al., 2025). These recurrences incur $\mathcal{O}(T)$ computational cost over the sequence and require $\mathcal{O}(1)$ cache memory during autoregressive decoding.

Hybrid attention. Hybrid attention combines full and linear attention across layers to preserve long-context retrieval capabilities while improving inference efficiency. Let $\mathcal{I} \subseteq [L]$ denote the set of layers retaining full attention; the remaining layers $[L] \setminus \mathcal{I}$ use linear attention. We encode this hybrid configuration as a binary mask $\mathbf{m} \in \{0,1\}^L$, where $m_l = \mathbf{1}[l \in \mathcal{I}]$. The sequence-mixing output at layer l is then given by

$$\begin{aligned} \mathbf{o}^{(l)} &= m_l\mathbf{o}^{(l)}_{\text{full}} + (1-m_l)\mathbf{o}^{(l)}_{\text{lin}}, \\ \mathbf{o}^{(l)}_{\text{full}} &= A^{(l)}_{\text{full}}(\mathbf{h}^{(l)}), \quad \mathbf{o}^{(l)}_{\text{lin}} = A^{(l)}_{\text{lin}}(\mathbf{h}^{(l)}). \end{aligned} \tag{3}$$

Given a budget of K full-attention layers, the feasible hybrid configuration set is $\mathcal{C}_K = \{\mathcal{I} \subseteq [L] : |\mathcal{I}| = K\}$, equivalently characterized by the mask constraint $\mathbf{1}^\top\mathbf{m} = K$. In Transformer-to-hybrid conversion, layer selection aims to identify the configuration in $\mathcal{C}_K$ that achieves the best downstream performance after recovery training.

3.1. FlashMorph Pipeline

FlashMorph (Lan et al., 2026) builds on the staged distillation pipeline of RADLADS (Goldstein et al., 2025) and HALO (Chen et al., 2026) for Transformer-to-hybrid conversion. Treating weight transfer as initialization, we organize the subsequent pipeline into four stages: layerwise alignment (Stage 1), layer selection (Stage 2), logits distillation (Stage 3), and long-context finetuning (Stage 4).

Stage 1: Layerwise alignment. FlashMorph first constructs morphable layers by equipping each full-attention layer with a linear-attention branch initialized from the corresponding pretrained projection weights. The original model serves as a frozen teacher, and only the linear-attention branches are trained. Let $\mathbf{h}^{(l)}_{\text{T}}$ denote the input to the teacher's attention module in layer l. The layerwise alignment objective is:

$$\mathcal{L}_{\text{align}} = \frac{1}{L}\sum_{l=1}^{L}\left\|A^{(l)}_{\text{lin}}(\mathbf{h}^{(l)}_{\text{T}}) - A^{(l)}_{\text{full}}(\mathbf{h}^{(l)}_{\text{T}})\right\|_2^2 \tag{4}$$

Layerwise alignment trains linear-attention branches to approximate teacher attention outputs, providing alternatives to full-attention layers for subsequent layer selection and a favorable initialization for end-to-end distillation.

Stage 2: Layer selection. FlashMorph freezes the model parameters θ_0 obtained after layerwise alignment and introduces a trainable continuous gate $\alpha_l \in [0,1]$ at each layer. Both branches take teacher hidden states as input, and their outputs are interpolated as follows:

$$\mathbf{o}^{(l)}_{\text{mix}} = \alpha_l\mathbf{o}^{(l)}_{\text{full,T}} + (1-\alpha_l)\mathbf{o}^{(l)}_{\text{lin,T}}. \tag{5}$$

During layer selection, only the gates are optimized by minimizing the following objective:

$$\mathcal{L}_{\text{LS}} = \frac{1}{L}\sum_{l=1}^{L}\left\|\mathbf{o}^{(l)}_{\text{mix}} - \mathbf{o}^{(l)}_{\text{full,T}}\right\|_2^2 + \lambda\sum_{l=1}^{L}\alpha^{(l)}, \tag{6}$$

where λ is a hyperparameter controlling the strength of linearization regularization. After optimization, FlashMorph retains full attention in the layers indexed by $\mathcal{I} =$

$\mathrm{TopK}(\{\alpha^{(l)}\}_{l=1}^{L}, K) \in \mathcal{C}_K$ and uses linear attention elsewhere. The gates are then discarded, yielding a discrete hybrid attention model for subsequent recovery training.

Stage 3: Logits distillation. The selected hybrid attention model $\mathcal{M}(\mathcal{I}; \theta_0)$ is then distilled end to end from the frozen original Transformer, with all student parameters optimized to minimize Kullback-Leibler (KL) divergence:

$$\mathcal{L}_{\mathrm{KD}} = \sum D_{\mathrm{KL}}\Big(p_{\mathrm{T}}(\cdot \mid x_{\le t}) \,\Big\|\, p_\theta(\cdot \mid x_{\le t})\Big), \quad (7)$$

where p_{T} and p_θ denote the output logits of the teacher and hybrid student model, and the KL divergence is computed over the full vocabulary.

Stage 4: Long-context finetuning. Following logits distillation, the hybrid model is finetuned on longer sequences using the standard next-token prediction objective:

$$\mathcal{L}_{\mathrm{FT}} = -\sum \log p_\theta\big(x_{t+1} \mid x_{\le t}\big). \quad (8)$$

3.2. Rethinking FlashMorph

Layer selection objective. FlashMorph formulates layer selection as a subset selection problem that maximizes hybrid model performance under a fixed full-attention budget (Lan et al., 2026). Since the quality of a selected layers also depends on subsequent recovery training, we evaluate selected architectures based on their post-recovery performance:

$$\mathcal{I}^\star = \arg\max_{\mathcal{I} \in \mathcal{C}_K} \mathrm{Score}\left(\mathcal{M}(\mathcal{I}; \theta^\star)\right), \quad (9)$$

where $\theta^\star$ denotes the model parameters of the selected hybrid architecture $\mathcal{I}$ obtained after initialization from θ_0 and subsequent recovery training, and $\mathrm{Score}(\cdot)$ measures post-recovery task performance. Individually recovering and evaluating all $\binom{L}{K}$ hybrid architectures is prohibitively expensive. FlashMorph therefore uses Eq. 6 as a proxy for layer selection before recovery. However, this proxy has two *off-policy* limitations that prevent it from fully reflecting the post-recovery performance of the resulting hybrid model:

L1: *Off-policy computation*. During FlashMorph's layer selection, continuous gates mix full- and linear-attention outputs computed from teacher hidden states (Eqs.5 and 6). The deployed model instead uses a discrete hybrid architecture under the prescribed full-attention budget, retaining one branch per layer and propagating its own hidden states (Eq.3). Thus, rankings learned under continuous mixing can misrepresent layer importance in the discrete hybrid architecture, while teacher-forced layerwise evaluation fails to capture how hidden-state deviations induced by attention replacement propagate through subsequent hybrid layers.

L2: *Off-policy weights*. Eq. 6 performs layer selection by optimizing continuous gates while keeping model weights frozen at their pre-recovery values θ_0, whereas Eq.9 evaluates discrete hybrid architectures using the recovered parameters $\theta^\star$. Since the frozen weights have not been adapted to the candidate hybrid architectures, subsequent recovery may alter the layer-importance ranking, potentially rendering the initial selection suboptimal.

These two forms of *off-policy* limitation highlight the need for layer selection to jointly account for end-to-end forward computation and weight adaptation on discrete hybrid architectures satisfying the prescribed full-attention budget. Building on this insight, Sec. 4 introduces our proposed FlashMorph2's *on-policy* layer-selection method.

4. Methodology

Drawing on *on-policy* learning, FlashMorph2 refines FlashMorph's optimization-based layer selection. Specifically, as illustrated in Fig. 2, FlashMorph2 samples discrete hybrid configurations satisfying the prescribed full-attention budget, performs end-to-end forward computation with self-propagated hidden states rather than teacher-forced hidden states, and jointly optimizes the layer-selection parameters and model weights using the same training objective as in the subsequent recovery stage to guide the layer selection.

4.1. On-Policy Hybrid Computation

We assign each layer a learnable layer-selection score ϕ_l, collectively forming the vector $\boldsymbol{\phi} \in \mathbb{R}^L$. To explore different hybrid configurations during layer-selection optimization, we add independent standard Gumbel noise to the layer-selection scores and select the K layers with the highest perturbed scores:

$$\mathcal{I} = \mathrm{TopK}\left(\boldsymbol{\phi} + \mathbf{g},\, K\right), \quad g_l \overset{\text{i.i.d.}}{\sim} \mathrm{Gumbel}(0,1). \quad (10)$$

We employ Gumbel-Top-K sampling trick without replacement (Kool et al., 2019), selecting exactly K distinct full-attention layers per sample, with higher layer-selection scores corresponding to higher selection probabilities. We denote the induced distribution over hybrid architectures in $\mathcal{C}_K$ by $q_{\boldsymbol{\phi}}$, with the corresponding binary mask defined as $m_l = \mathbf{1}[l \in \mathcal{I}]$. All layer-selection scores are initialized to the same constant, ensuring no prior preference for any layer during initial sampling.

Given a sampled mask, the model performs an *on-policy* forward pass through the corresponding discrete hybrid architecture (Eq. 3). Each layer operates on hidden states propagated through the sampled architecture and uses the output of its selected attention branch, allowing earlier branch choices to affect subsequent layer inputs and ultimately the final predictions. Layer-selection evaluation thus captures error accumulation and cross-layer compensation induced by attention replacement through forward computation consistent with that of the deployed hybrid architecture.

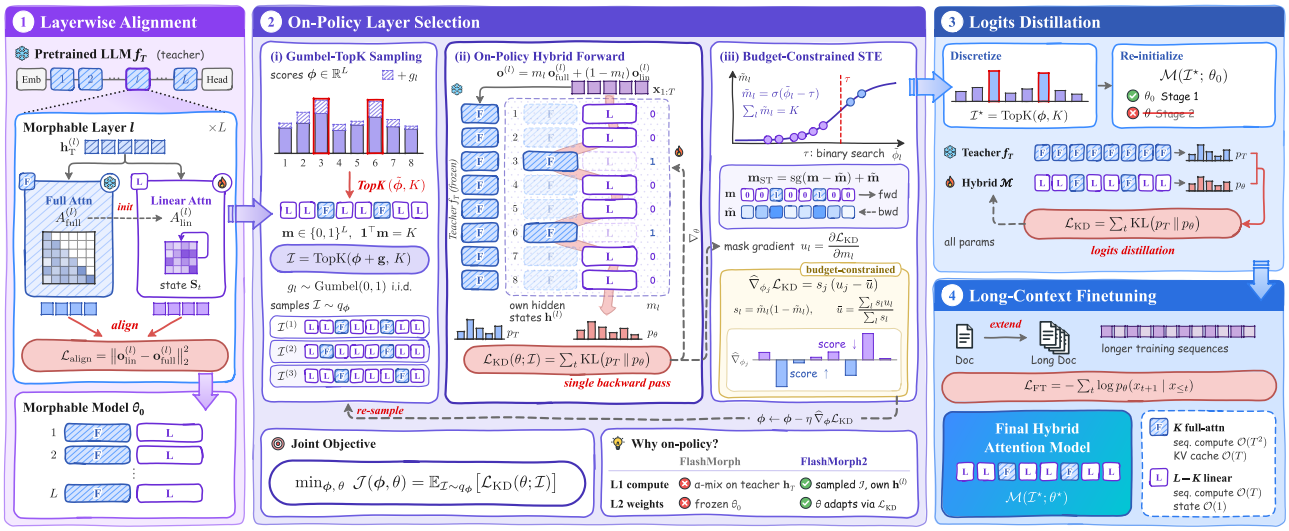


Figure 2. **Overview of FlashMorph2.** Following layerwise alignment, FlashMorph2 constructs a hybrid attention model through *on-policy* layer selection, jointly optimizing layer scores and model weights over sampled discrete hybrid architectures. The resulting hybrid attention model is further trained using logits distillation and long-context finetuning to recover model quality.

4.2. On-Policy Weight Adaptation

To account for weight adaptation during subsequent recovery, we initialize the student model with the layer-wise aligned parameters θ_0 and jointly optimize the layer-selection parameters $\boldsymbol{\phi}$ and model weights θ over sampled hybrid architectures:

$$\min_{\boldsymbol{\phi},\,\theta} \; \mathcal{J}(\boldsymbol{\phi}, \theta) = \mathbb{E}_{\mathcal{I} \sim q_{\boldsymbol{\phi}}} \Big[\mathcal{L}_{\text{KD}}(\theta; \mathcal{I}) \Big], \tag{11}$$

where $\mathcal{L}_{\text{KD}}(\theta; \mathcal{I})$ denotes the logits distillation objective used in subsequent recovery training stage (Eq. 7), computed on the currently sampled hybrid architecture $\mathcal{I}$. All sampled hybrid architectures are constructed by selecting attention branches from a shared morphable model, directly reusing the parameters of the selected branches. During layer-selection optimization, we compute the distillation loss on sampled architectures and obtain gradients for both the layer-selection parameters $\boldsymbol{\phi}$ and model weights θ through a single backward pass. The layer-selection parameters are updated via gradient straight-through estimator (STE) (Bengio et al., 2013; Jang et al., 2016) (Sec. 4.3) to adjust the architecture-sampling distribution. For each sampled architecture, unselected attention branches receive no gradients. As the model weights continuously adapt to the sampled hybrid architectures, the layer-selection parameters are updated based on the corresponding optimization signals. This *on-policy* layer-selection method incorporates weight adaptation into architecture evaluation, mitigating the discrepancy between frozen-weight layer evaluation and post-recovery performance.

After optimization, we remove the sampling noise and select the final full-attention layers as $\mathcal{I}^\star = \text{TopK}(\boldsymbol{\phi}, K)$. To isolate the effect of layer selection, we retain only the selected architecture, reinitialize $\mathcal{M}(\mathcal{I}^\star; \theta_0)$ with the original parameters θ_0, and apply the same Stage 3-4 recovery training pipeline. The model weights θ updated during Stage 2 serve only to guide layer selection and are not used to initialize subsequent recovery training.

4.3. Budget-Constrained Straight-Through Estimator

Discrete Top-K selection is non-differentiable. To estimate gradients for the layer-selection parameters, we construct a continuous relaxation that preserves the same budget constraint (Wijk et al., 2026). For a sampled hybrid architecture, let $\tilde{\boldsymbol{\phi}} = \boldsymbol{\phi} + \mathbf{g}$ denote the perturbed layer-selection scores, and define:

$$\tilde{m}_l = \sigma\big(\tilde{\phi}_l - \tau\big), \quad \sum_{l=1}^{L} \sigma\big(\tilde{\phi}_l - \tau\big) = K, \tag{12}$$

where σ denotes the sigmoid function, and the threshold τ is determined via binary search. We use the binary mask $\mathbf{m}$ in the forward pass during layer selection stage and propagate gradients through the continuous relaxation $\tilde{\mathbf{m}}$ in the backward pass, following the gradient straight-through estimator (STE) (Bengio et al., 2013; Jang et al., 2016):

$$\mathbf{m}_{\text{ST}} = \text{sg}\left(\mathbf{m} - \tilde{\mathbf{m}}\right) + \tilde{\mathbf{m}}, \tag{13}$$

where $\text{sg}(\cdot)$ denotes the stop-gradient operator. Although the threshold is computed outside the automatic differentiation graph, its dependence on the selection scores is accounted for through implicit differentiation during backpropagation. Let $u_l = \partial \mathcal{L}_{\text{KD}} / \partial m_l$ denote the gradient of the loss with respect to the mask, and define $s_l = \tilde{m}_l(1 - \tilde{m}_l)$. We obtain the following biased gradient straight-through estimator for

Table 1. **Long-context Performance on RULER and LongBench.** Avg denotes the average score over 13 RULER or 14 LongBench tasks. For each backbone and linear-attention variant, the best and second-best scores among layer-selection methods are highlighted.

Attn.	Method	RULER NIAH	RULER VT	RULER Agg.	RULER QA	RULER **Avg**	LongBench SQA	LongBench MQA	LongBench Sum.	LongBench FS	LongBench Code	LongBench **Avg**
0.6B Backbone												
Full	Qwen3	94.38	95.32	55.78	33.56	79.15	21.22	20.86	19.39	60.78	33.26	30.95
Lightning	Uniform	61.12	1.56	51.64	29.99	50.29	19.90	16.46	17.92	48.58	33.28	26.80
	KL-LS	54.49	2.20	34.46	31.66	43.87	18.52	20.15	18.60	43.32	38.71	27.09
	HALO	59.51	1.20	50.48	32.61	49.50	19.05	18.24	17.11	44.72	34.65	26.19
	FlashMorph	68.56	10.84	55.03	35.33	56.93	23.36	19.16	19.55	46.43	36.73	28.50
	FlashMorph2	68.52	17.24	58.46	36.49	58.10	22.90	19.39	19.80	47.34	39.97	29.16
GLA	Uniform	62.16	4.44	48.80	27.13	50.27	19.77	15.61	16.82	46.88	29.56	25.45
	KL-LS	67.95	1.80	43.47	29.95	53.25	21.86	17.64	19.19	47.91	30.79	27.24
	HALO	65.84	2.08	47.37	29.98	52.58	20.03	19.23	17.37	46.83	29.00	26.31
	FlashMorph	68.99	5.52	51.27	32.56	55.78	23.64	18.99	18.40	43.72	29.64	26.68
	FlashMorph2	76.16	7.20	52.62	36.62	61.15	23.49	17.93	19.46	46.42	31.98	27.57
GDN	Uniform	64.09	3.76	50.23	33.26	52.57	21.19	16.40	17.89	49.09	35.19	27.43
	KL-LS	71.10	2.96	48.97	34.76	56.86	20.89	17.25	20.35	49.94	38.85	28.79
	HALO	60.20	20.28	44.81	33.44	50.64	20.78	16.24	18.48	46.47	36.97	27.13
	FlashMorph	72.48	15.60	56.54	37.65	60.29	23.37	18.91	19.68	48.78	37.94	29.15
	FlashMorph2	76.83	18.68	60.54	36.49	63.65	24.21	20.04	20.15	49.29	37.88	29.77
1.7B Backbone												
Full	Qwen3	99.51	95.72	77.47	46.44	87.66	26.23	13.12	18.53	66.02	41.14	32.43
Lightning	Uniform	71.36	29.16	51.97	40.73	60.41	21.19	18.24	20.27	57.69	37.59	30.53
	PostNAS	68.17	23.16	64.94	42.88	60.32	20.69	22.06	19.92	55.39	36.81	30.56
	KL-LS	56.95	0.24	44.01	37.22	47.56	20.07	18.09	21.11	56.05	42.42	30.77
	HALO	72.89	29.48	63.60	43.91	63.66	20.31	21.42	20.30	57.92	40.12	31.44
	FlashMorph	77.59	31.64	60.05	45.01	66.35	20.25	21.06	20.02	59.23	37.09	31.13
	FlashMorph2	79.96	41.24	62.91	47.74	69.40	22.71	23.56	21.30	59.72	42.53	33.35
GLA	Uniform	72.86	20.56	54.98	41.98	61.34	21.00	18.53	18.73	57.56	38.49	30.32
	PostNAS	74.95	19.76	51.59	45.92	62.64	24.23	21.76	17.21	56.74	38.77	31.24
	KL-LS	58.94	11.60	50.48	38.59	50.86	19.91	20.30	17.80	57.45	43.69	30.98
	HALO	64.45	25.84	62.20	37.39	56.97	22.38	17.85	15.93	56.19	40.59	29.87
	FlashMorph	80.43	30.52	56.48	42.83	67.12	21.93	19.16	15.77	58.89	43.82	31.06
	FlashMorph2	79.14	22.28	61.81	48.01	67.31	23.38	20.50	15.84	60.20	43.79	31.95
GDN	Uniform	75.56	63.32	59.41	45.42	67.50	19.53	22.27	20.48	59.23	35.41	31.10
	PostNAS	79.20	53.16	64.97	44.92	69.73	25.24	20.72	21.73	58.16	36.65	32.20
	KL-LS	79.71	68.12	65.44	47.91	71.73	24.57	19.74	21.41	59.62	41.06	32.73
	HALO	76.18	62.68	63.31	46.94	68.66	25.13	22.30	18.55	57.46	40.94	32.30
	FlashMorph	81.79	66.76	65.89	48.52	73.07	24.69	22.23	20.78	59.51	37.52	32.62
	FlashMorph2	84.83	58.80	66.28	48.96	74.46	25.72	20.53	21.39	61.16	43.48	33.81

the expected discrete objective:

$$\widehat{\nabla}_{\phi_j}\mathcal{L}_{\mathrm{KD}} = s_j\big(u_j - \bar{u}\big), \quad \bar{u} = \frac{\sum_{l=1}^{L} s_l u_l}{\sum_{l=1}^{L} s_l}. \tag{14}$$

Implicit differentiation of the threshold τ couples gradients with respect to layer-selection scores under the full-attention budget constraint. Computing u_l requires the gradient back-propagated from the end-to-end loss and the difference between full- and linear-attention outputs. Both branches are therefore evaluated, while the binary mask ensures that only the selected branch contributes to the forward output.

5. Experiments

5.1. Experiment Setup

Model and training. We use Qwen3-0.6B and Qwen3-1.7B (Yang et al., 2025a) as pretrained Transformer backbones and construct hybrid models with three linear-attention variants: Lightning Attention (Qin et al., 2024), Gated Linear Attention (GLA) (Yang et al., 2023), and Gated DeltaNet (GDN) (Yang et al., 2025b). We implement the linear-attention branches using the `flash-linear-attention` library (Yang & Zhang, 2024), incorporating QK normalization, GQA-to-MHA conversion, output normalization, and an output gate. Both full- and linear-attention layers adopt RoPE (Su et al., 2024). Unless otherwise specified, the hybrid ratio of linear-attention to full-attention layers is 3:1. We follow the four-stage Transformer-to-hybrid conversion pipeline described in Sec. 3.1, using training data randomly sampled from DCLM (Li et al., 2024). Details on model and training configurations are in Appendix A.

Evaluation. We evaluate FlashMorph2 on long-context capabilities using RULER (Hsieh et al., 2024) and 14

Table 2. **Performance on Recall-intensive Tasks.** Avg denotes the average score across all six recall-intensive tasks. For each backbone and linear-attention variant, the best and second-best scores among layer-selection methods are highlighted.

Attn.	Method	*0.6B Backbone*							*1.7B Backbone*						
		FDA	SWDE	SQD	NQ	TQA	DROP	**Avg**	FDA	SWDE	SQD	NQ	TQA	DROP	**Avg**
Full	Qwen3	76.95	66.52	46.35	30.92	57.41	28.27	51.07	70.78	71.11	51.21	34.75	63.39	30.81	53.67
Lightning	Uniform	54.90	61.93	45.58	25.34	57.76	27.17	45.45	49.36	62.83	50.97	27.91	64.34	31.53	47.82
	PostNAS	—	—	—	—	—	—	—	54.08	64.81	51.58	28.51	64.10	32.34	49.24
	KL-LS	49.36	57.61	46.92	22.68	58.35	27.93	43.81	48.64	60.13	50.94	29.17	63.15	32.20	47.37
	HALO	56.72	61.21	46.95	23.25	58.12	28.17	45.74	49.36	64.27	51.01	28.13	63.57	33.01	48.23
	FlashMorph	53.27	61.57	46.65	24.87	59.12	28.89	45.73	57.17	62.83	51.91	27.91	64.34	33.72	49.65
	FlashMorph2	59.35	61.57	47.92	25.94	56.87	28.80	46.74	60.16	64.27	52.65	30.00	64.45	32.06	50.60
GLA	Uniform	54.08	61.93	46.25	25.37	57.94	28.94	45.75	53.54	63.01	51.51	29.17	65.23	32.82	49.21
	PostNAS	—	—	—	—	—	—	—	57.62	66.07	51.98	29.90	65.17	32.82	50.59
	KL-LS	57.62	60.94	48.56	26.35	58.89	28.17	46.76	51.91	64.18	51.44	29.49	64.93	30.86	48.80
	HALO	60.25	60.31	47.65	24.77	58.00	29.80	46.80	49.73	59.50	50.80	27.75	64.75	31.62	47.36
	FlashMorph	62.98	62.20	46.75	26.13	57.88	27.93	47.31	61.98	61.75	53.35	30.66	65.46	33.92	51.19
	FlashMorph2	64.16	65.35	48.76	27.37	59.24	28.13	48.84	59.26	65.44	53.25	29.97	66.11	33.25	51.21
GDN	Uniform	55.08	62.38	45.81	25.94	57.88	28.80	45.98	52.72	62.38	52.38	29.36	65.76	32.10	49.12
	PostNAS	—	—	—	—	—	—	—	59.17	67.51	53.18	32.28	65.46	33.78	51.90
	KL-LS	64.25	62.29	47.92	26.42	58.77	27.55	47.87	61.89	64.36	53.08	32.15	65.28	32.06	51.47
	HALO	57.17	62.29	48.09	26.13	59.60	29.28	47.09	60.89	64.90	52.71	31.17	66.29	33.88	51.64
	FlashMorph	62.34	63.73	48.63	27.24	59.00	27.17	48.02	59.53	65.17	53.25	32.06	66.00	34.64	51.77
	FlashMorph2	63.43	64.18	47.39	27.78	58.77	29.47	48.50	61.71	66.88	53.62	31.58	66.94	33.21	52.32

tasks from LongBench (Bai et al., 2024). We further evaluate performance on 6 recall-intensive tasks using the `prefix-linear-attention` evaluation suite (Arora et al., 2024b). We also evaluate language modeling perplexity and zero-shot commonsense reasoning using `lm-evaluation-harness` (Gao et al., 2024). All evaluations are conducted using vLLM (Kwon et al., 2023) for accelerated inference. Detailed evaluation settings are also provided in the Appendix B.

Baselines. We compare FlashMorph2 against five layer selection baselines, including: uniform interleaving (Uniform), PostNAS (Gu et al., 2026), KL-LS (Li et al., 2026a), HALO (Chen et al., 2026), and FlashMorph (Lan et al., 2026). To isolate the effect of layer selection, all methods use the same full-attention budget, layerwise-aligned initialization, and subsequent recovery training settings.

5.2. Main Results

Long-context capabilities. Table 1 reports long-context performance under the same full-attention budget. FlashMorph2 achieves the highest average scores on both RULER and LongBench across all six combinations of backbone size and linear-attention variant. Averaged across these configurations, it raises the RULER score from 63.26 to 65.68 (+2.42 points) and the LongBench score from 29.86 to 30.94 (+1.08 points) compared with FlashMorph. These results demonstrate that FlashMorph2 identifies hybrid architectures with consistently superior long-context performance following an identical recovery training settings.

Recall-intensive tasks. Table 2 reports performance on six recall-intensive tasks. FlashMorph2 achieves the highest average score across tasks in all six combinations of backbone size and linear-attention variant. Averaged across these configurations, it improves the score from 48.94 to 49.70 (+0.76 points) compared with FlashMorph. These results indicate that the benefits of FlashMorph2 extend to real-world in-context retrieval tasks.

Language modeling and commonsense reasoning. Table 9 in the Appendix reports language modeling perplexity and zero-shot commonsense reasoning accuracy. FlashMorph2 remains competitive with the layer-selection baselines on these evaluations, suggesting that general language capabilities are preserved alongside the improvements in long-context and recall-intensive tasks.

5.3. Analysis

Layer-selection efficiency. Fig. 3 reports the computational cost of layer selection across different model scales. Compared with FlashMorph, FlashMorph2 requires 1.7–2.6× the FLOPs and 1.9–2.4× the GPU hours, reflecting the additional computation required for end-to-end evaluation of discrete hybrid architectures and concurrent model adaptation. This overhead is confined to a one-time selection

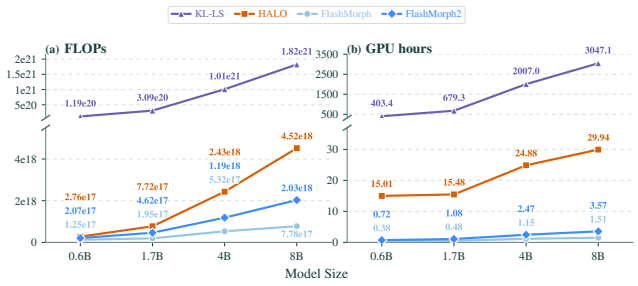


Figure 3. **Layer-selection Efficiency across Model Scales.** Comparison of (a) FLOPs and (b) GPU hours required for layer selection across Qwen3 models from 0.6B to 8B parameters.

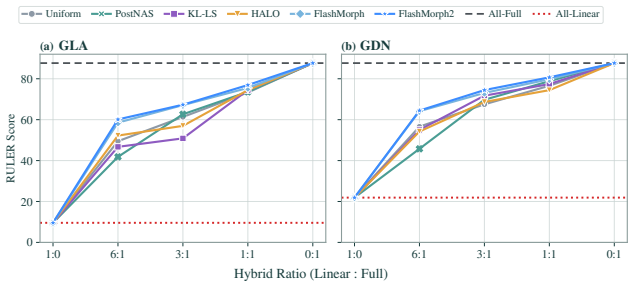


Figure 4. **RULER Performance across Hybrid Ratios.** Average RULER scores on the Qwen3-1.7B backbone across varying hybrid ratios for (a) GLA and (b) GDN variants.

stage, which accounts for approximately only $< 1\%$ of the overall conversion token budget, and does not affect deployment efficiency under the same full-attention budget. Across models ranging from 0.6B to 8B parameters, FlashMorph2 completes layer selection in only 0.72–3.57 GPU hours. It is also substantially more efficient than layerwise-scoring approaches. Relative to FlashMorph2, HALO requires 1.3–2.2× as many FLOPs and 8.4–20.8× as many GPU hours, while KL-LS requires 575–897× and 560–854× as many, respectively. Overall, despite its modest overhead over FlashMorph, FlashMorph2 remains highly efficient and scalable while being substantially more efficient than alternative layer-selection methods.

Hybrid ratios. Fig. 4 reports RULER performance on Qwen3-1.7B across different hybrid ratios for GLA and GDN. While full linearization (1:0) reduces the RULER score from 87.66 to 9.57 on GLA and 21.89 on GDN, retaining a small subset of full-attention layers substantially restores retrieval performance. Across all evaluated hybrid ratios (6:1, 3:1, and 1:1), FlashMorph2 consistently outperforms all baselines on both linear-attention variants. The advantage of optimization-based layer selection is particularly pronounced under the tightest full-attention budget (6:1), where FlashMorph2 achieves 60.12 on GLA and 64.45 on GDN, outperforming Uniform by +10.65 and +7.86 points while several competing methods fall below simple uniform interleaving. As the full-attention ratio increases to 3:1 and 1:1, performance across methods gradually converges toward the full-attention backbone, yet FlashMorph2 continues to improve upon FlashMorph in all settings, reaching 76.93 on GLA (+1.42) and 80.69 on GDN (+0.92) at 1:1. Overall, these results confirm that FlashMorph2 remains effective and robust across varying hybrid ratios.

Ablation study. Table 3 evaluates the effects of *on-policy* computation and *on-policy* weight adaptation on Qwen3-1.7B with GDN. Removing *on-policy* computation significantly reduces scores by 29.55 points on RULER and 5.71 points on LongBench. Inspecting the selected architectures reveals that this variant allocates all seven full-attention layers to the deepest layers $\{20, 21, 23, 24, 25, 26, 27\}$, leaving the first 20 layers entirely linear. This concentration reflects the mismatch between teacher-forced local scoring and end-to-end discrete hybrid computation, where early hidden-state deviations accumulate across layers. Removing *on-policy* weight adaptation instead lowers RULER by 4.88 points and LongBench by 1.38 points, replacing layers $\{6, 17, 19, 20\}$ selected by FlashMorph2 with $\{0, 1, 2, 26\}$, which indicates that weight adaptation substantially alters the preferred budget allocation. Removing both components reverts to FlashMorph, which still trails FlashMorph2 by 1.39 points on RULER and 1.19 points on LongBench. Overall, these results confirm that discrete hybrid computation and weight adaptation are complementary and must be jointly considered during layer selection.

Table 3. **Ablation Study of FlashMorph2.** Average RULER and LongBench scores on the Qwen3-1.7B backbone with GDN.

Method	RULER	LongBench
FlashMorph2	74.46	33.81
w/o On-Policy Computation	44.91	28.10
w/o On-Policy Weight	69.58	32.43
w/o Both (FlashMorph)	73.07	32.62

6. Conclusion

In this work, we present FlashMorph2, an *on-policy* layer-selection method for converting pretrained Transformers into hybrid attention models under a prescribed full-attention budget. To overcome the *off-policy* computation and weight mismatches in prior optimization-based selection, FlashMorph2 samples discrete budget-constrained hybrid architectures, performs end-to-end forward passes with self-propagated hidden states, and jointly optimizes layer-selection scores and model weights under the recovery training objective. Experiments across multiple Qwen3 model scales, linear-attention variants, and hybrid ratios show that FlashMorph2 consistently improves long-context and recall-intensive performance over existing baselines while maintaining high selection efficiency and scalability.

References

Arora, S., Yang, B., Eyuboglu, S., Narayan, A., Hojel, A., Trummer, I., and Ré, C. Language models enable simple systems for generating structured views of heterogeneous data lakes. *arXiv preprint arXiv:2304.09433*, 2023.

Arora, S., Eyuboglu, S., Zhang, M., Timalsina, A., Alberti, S., Zinsley, D., Zou, J., Rudra, A., and Ré, C. Simple linear attention language models balance the recall-throughput tradeoff. *arXiv preprint arXiv:2402.18668*, 2024a.

Arora, S., Timalsina, A., Singhal, A., Spector, B., Eyuboglu, S., Zhao, X., Rao, A., Rudra, A., and Ré, C. Just read twice: closing the recall gap for recurrent language models. *arXiv preprint arXiv:2407.05483*, 2024b.

Bai, Y., Lv, X., Zhang, J., Lyu, H., Tang, J., Huang, Z., Du, Z., Liu, X., Zeng, A., Hou, L., Dong, Y., Tang, J., and Li, J. LongBench: A bilingual, multitask benchmark for long context understanding. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3119–3137, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.172. URL https://aclanthology.org/2024.acl-long.172.

Bengio, S., Vinyals, O., Jaitly, N., and Shazeer, N. Scheduled sampling for sequence prediction with recurrent neural networks. *Advances in neural information processing systems*, 28, 2015.

Bengio, Y., Léonard, N., and Courville, A. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

Bick, A., Li, K. Y., Xing, E. P., Kolter, J. Z., and Gu, A. Transformers to ssms: Distilling quadratic knowledge to subquadratic models. *Advances in neural information processing systems*, 37:31788–31812, 2024.

Bick, A., Katsch, T., Sohoni, N., Desai, A., and Gu, A. Llamba: Scaling distilled recurrent models for efficient language processing. *arXiv preprint arXiv:2502.14458*, 2025.

Bisk, Y., Zellers, R., Gao, J., Choi, Y., et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 7432–7439, 2020.

Cao, R., Chen, M., Chen, J., Cui, Z., Feng, Y., Hui, B., Jing, Y., Li, K., Li, M., Lin, J., et al. Qwen3-coder-next technical report. *arXiv preprint arXiv:2603.00729*, 2026.

Chen, Y., Thai, Z. L., Zhou, Z., Zhang, Z., Shen, X., Wang, S., Xiao, C., Han, X., and Liu, Z. Hybrid linear attention done right: Efficient distillation and effective architectures for extremely long contexts. *arXiv preprint arXiv:2601.22156*, 2026.

Clark, C., Lee, K., Chang, M.-W., Kwiatkowski, T., Collins, M., and Toutanova, K. BoolQ: Exploring the surprising difficulty of natural yes/no questions. In Burstein, J., Doran, C., and Solorio, T. (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 2924–2936, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1300. URL https://aclanthology.org/N19-1300/.

Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.

Dasigi, P., Lo, K., Beltagy, I., Cohan, A., Smith, N. A., and Gardner, M. A dataset of information-seeking questions and answers anchored in research papers. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 4599–4610, 2021.

Du, J., Sun, W., Lan, D., Hu, J., Tao, Z., and Cheng, Y. Mom: Linear sequence modeling with mixture-of-memories. In *International Conference on Learning Representations*, volume 2026, pp. 106613–106631, 2026.

Dua, D., Wang, Y., Dasigi, P., Stanovsky, G., Singh, S., and Gardner, M. Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs, 2019.

Fabbri, A. R., Li, I., She, T., Li, S., and Radev, D. Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pp. 1074–1084, 2019.

Gao, L., Tow, J., Abbasi, B., Biderman, S., Black, S., DiPofi, A., Foster, C., Golding, L., Hsu, J., Le Noac'h, A., Li, H., McDonell, K., Muennighoff, N., Ociepa, C., Phang, J., Reynolds, L., Schoelkopf, H., Skowron, A., Sutawika, L., Tang, E., Thite, A., Wang, B., Wang, K., and Zou, A. The language model evaluation harness, 07 2024. URL https://zenodo.org/records/12608602.

Gliwa, B., Mochol, I., Biesek, M., and Wawer, A. SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization. In Wang, L., Cheung, J. C. K., Carenini, G., and Liu, F. (eds.), *Proceedings of*

the 2nd Workshop on New Frontiers in Summarization, pp. 70–79, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-5409. URL https://aclanthology.org/D19-5409/.

Goldstein, D., Alcaide, E., Lu, J., and Cheah, E. Radlads: Rapid attention distillation to linear attention decoders at scale. *arXiv preprint arXiv:2505.03005*, 2025.

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

Gu, Y., Hu, Q., Xi, H., Chen, J., Yang, S., Han, S., and Cai, H. Jet-nemotron: Efficient language model with post neural architecture search. *Advances in Neural Information Processing Systems*, 38:47191–47218, 2026.

Guo, D., Xu, C., Duan, N., Yin, J., and McAuley, J. Longcoder: A long-range pre-trained language model for code completion. In *International Conference on Machine Learning*, pp. 12098–12107. PMLR, 2023.

Ho, X., Nguyen, A.-K. D., Sugawara, S., and Aizawa, A. Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 6609–6625, 2020.

Hsieh, C.-P., Sun, S., Kriman, S., Acharya, S., Rekesh, D., Jia, F., Zhang, Y., and Ginsburg, B. Ruler: What's the real context size of your long-context language models? *arXiv preprint arXiv:2404.06654*, 2024.

Hu, J., Pan, Y., Du, J., Lan, D., Tang, X., Wen, Q., Liang, Y., and Sun, W. Improving bilinear RNN with closed-loop control. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=jlJaRXDzCE.

Hu, S., Tu, Y., Han, X., He, C., Cui, G., Long, X., Zheng, Z., Fang, Y., Huang, Y., Zhao, W., et al. Minicpm: Unveiling the potential of small language models with scalable training strategies. *arXiv preprint arXiv:2404.06395*, 2024.

Huang, L., Cao, S., Parulian, N., Ji, H., and Wang, L. Efficient attentions for long document summarization. In *Proceedings of the 2021 conference of the north American chapter of the association for computational linguistics: Human language technologies*, pp. 1419–1436, 2021.

Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.

Jelassi, S., Brandfonbrener, D., Kakade, S. M., and Malach, E. Repeat after me: Transformers are better than state space models at copying. *arXiv preprint arXiv:2402.01032*, 2024.

Joshi, M., Choi, E., Weld, D. S., and Zettlemoyer, L. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1601–1611, 2017.

Kasai, J., Peng, H., Zhang, Y., Yogatama, D., Ilharco, G., Pappas, N., Mao, Y., Chen, W., and Smith, N. A. Finetuning pretrained transformers into rnns. In *Proceedings of the 2021 conference on empirical methods in natural language processing*, pp. 10630–10643, 2021.

Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pp. 5156–5165. PMLR, 2020.

Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Kočiskỳ, T., Schwarz, J., Blunsom, P., Dyer, C., Hermann, K. M., Melis, G., and Grefenstette, E. The narrativeqa reading comprehension challenge. *Transactions of the Association for Computational Linguistics*, 6:317–328, 2018.

Kool, W., Van Hoof, H., and Welling, M. Stochastic beams and where to find them: The gumbel-top-k trick for sampling sequences without replacement. In *International conference on machine learning*, pp. 3499–3508. PMLR, 2019.

Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Devlin, J., Lee, K., et al. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466, 2019.

Kwon, W., Li, Z., Zhuang, S., Sheng, Y., Zheng, L., Yu, C. H., Gonzalez, J. E., Zhang, H., and Stoica, I. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.

Lan, D., Sun, W., Hu, J., Du, J., and Cheng, Y. Liger: Linearizing large language models to gated recurrent structures. *arXiv preprint arXiv:2503.01496*, 2025.

Lan, D., Zheng, J., Ren, Y., Xia, X., Wang, X., Xiao, X., Qiu, X., and Cheng, Y. Morphing into hybrid attention models. *arXiv preprint arXiv:2606.30562*, 2026.

Li, A., Gong, B., Yang, B., Shan, B., Liu, C., Zhu, C., Zhang, C., Guo, C., Chen, D., Li, D., et al. Minimax-01: Scaling foundation models with lightning attention. *arXiv preprint arXiv:2501.08313*, 2025.

Li, J., Fang, A., Smyrnis, G., Ivgi, M., Jordan, M., Gadre, S., Bansal, H., Guha, E., Keh, S., Arora, K., et al. Datacomp-lm: In search of the next generation of training sets for language models. *Advances in Neural Information Processing Systems*, 37:14200–14282, 2024.

Li, X. and Roth, D. Learning question classifiers. In *Coling 2002: The 19th international conference on computational linguistics*, 2002.

Li, Y., Yang, S., Tan, S., Mishra, M., Panda, R., Zhou, J., and Kim, Y. Distilling to hybrid attention models via kl-guided layer selection. In *International Conference on Learning Representations*, volume 2026, pp. 108531–108555, 2026a.

Li, Y., Zuo, Y., He, B., Zhang, J., Xiao, C., Qian, C., Yu, T., Gao, H.-a., Yang, W., Liu, Z., et al. Rethinking on-policy distillation of large language models: Phenomenology, mechanism, and recipe. *arXiv preprint arXiv:2604.13016*, 2026b.

Lin, A., Wohlwend, J., Chen, H., and Lei, T. Autoregressive knowledge distillation through imitation learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 6121–6133, 2020.

Liu, T., Xu, C., and McAuley, J. Repobench: Benchmarking repository-level code auto-completion systems, 2024. URL https://arxiv.org/abs/2306.03091.

Lockard, C., Shiralkar, P., and Dong, X. L. Openceres: When open information extraction meets the semi-structured web. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 3047–3056, 2019.

Lu, K. and Lab, T. M. On-policy distillation. *Thinking Machines Lab: Connectionism*, 2025. doi: 10.64434/tml.20251026. https://thinkingmachines.ai/blog/on-policy-distillation.

Ma, W., Wei, J., Zhao, L., Zhang, H., Xiao, B., Li, L., Yang, Q., Gao, B., Wang, Y., Li, R., et al. Mopd: Multi-teacher on-policy distillation for capability integration in llm post-training. *arXiv preprint arXiv:2606.30406*, 2026.

Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.

Paperno, D., Kruszewski, G., Lazaridou, A., Pham, N. Q., Bernardi, R., Pezzelle, S., Baroni, M., Boleda, G., and Fernández, R. The LAMBADA dataset: Word prediction requiring a broad discourse context. In Erk, K. and Smith, N. A. (eds.), *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1525–1534, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1144. URL https://aclanthology.org/P16-1144/.

Qin, Z., Li, D., Sun, W., Sun, W., Shen, X., Han, X., Wei, Y., Lv, B., Luo, X., Qiao, Y., et al. Transnormerllm: A faster and better large language model with improved transnormer. *arXiv preprint arXiv:2307.14995*, 2023.

Qin, Z., Sun, W., Li, D., Shen, X., Sun, W., and Zhong, Y. Lightning attention-2: A free lunch for handling unlimited sequence lengths in large language models. *arXiv preprint arXiv:2401.04658*, 2024.

Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pp. 2383–2392, 2016.

Rajpurkar, P., Jia, R., and Liang, P. Know what you don't know: Unanswerable questions for SQuAD. In Gurevych, I. and Miyao, Y. (eds.), *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 784–789, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-2124. URL https://aclanthology.org/P18-2124/.

Ross, S., Gordon, G., and Bagnell, D. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference Proceedings, 2011.

Sakaguchi, K., Bras, R. L., Bhagavatula, C., and Choi, Y. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.

Sap, M., Rashkin, H., Chen, D., Le Bras, R., and Choi, Y. Social IQa: Commonsense reasoning about social interactions. In Inui, K., Jiang, J., Ng, V., and Wan, X. (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 4463–4473, Hong

Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1454. URL https://aclanthology.org/D19-1454/.

Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.

Sun, W., Hu, J., Zhou, Y., Du, J., Lan, D., Wang, K., Zhu, T., Qu, X., Zhang, Y., Mo, X., Liu, D., Liang, Y., Chen, W., Li, G., and Cheng, Y. Speed always wins: A survey on efficient architectures for large language models. *arXiv preprint arXiv:2508.09834*, 2025a.

Sun, W., Lan, D., Zhu, T., Qu, X., and Cheng, Y. Linear-moe: Linear sequence modeling meets mixture-of-experts. *arXiv preprint arXiv:2503.05447*, 2025b.

Team, K., Zhang, Y., Lin, Z., Yao, X., Hu, J., Meng, F., Liu, C., Men, X., Yang, S., Li, Z., et al. Kimi linear: An expressive, efficient attention architecture. *arXiv preprint arXiv:2510.26692*, 2025.

Team, K., Bai, T., Bai, Y., Bao, Y., Cai, J., Cai, X., Cao, P., Cao, Y., Chai, Z., Charles, Y., et al. Kimi k3: Open frontier intelligence. *arXiv preprint arXiv:2607.24653*, 2026.

Team, Q. Qwen3.5: Towards native multimodal agents, February 2026. URL https://qwen.ai/blog?id=qwen3.5.

Trivedi, H., Balasubramanian, N., Khot, T., and Sabharwal, A. MuSiQue: Multihop questions via single-hop question composition. *Transactions of the Association for Computational Linguistics*, 2022.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Wang, J., Paliotta, D., May, A., Rush, A. M., and Dao, T. The mamba in the llama: Distilling and accelerating hybrid models. *Advances in Neural Information Processing Systems*, 37:62432–62457, 2024.

Wijk, K., Vinuesa, R., and Azizpour, H. Differentiable top-k: From one-hot to k-hot, 2026. URL https://openreview.net/forum?id=VMlajIH1oF.

Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025a.

Yang, S. and Zhang, Y. Fla: A triton-based library for hardware-efficient implementations of linear attention mechanism, January 2024. URL https://github.com/fla-org/flash-linear-attention.

Yang, S., Wang, B., Shen, Y., Panda, R., and Kim, Y. Gated linear attention transformers with hardware-efficient training. *arXiv preprint arXiv:2312.06635*, 2023.

Yang, S., Wang, B., Zhang, Y., Shen, Y., and Kim, Y. Parallelizing linear transformers with the delta rule over sequence length. *Advances in neural information processing systems*, 37:115491–115522, 2024.

Yang, S., Kautz, J., and Hatamizadeh, A. Gated delta networks: Improving mamba2 with delta rule. In *International Conference on Learning Representations*, volume 2025, pp. 29687–29707, 2025b.

Yang, Z., Qi, P., Zhang, S., Bengio, Y., Cohen, W., Salakhutdinov, R., and Manning, C. D. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pp. 2369–2380, 2018.

Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. Hellaswag: Can a machine really finish your sentence? In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pp. 4791–4800, 2019.

Zhang, M., Arora, S., Chalamala, R., Spector, B., Wu, A., Ramesh, K., Singhal, A., and Ré, C. Lolcats: On low-rank linearizing of large language models. In *International Conference on Learning Representations*, volume 2025, pp. 46200–46253, 2025.

Zhang, Y., Yang, S., Zhu, R., Zhang, Y., Cui, L., Wang, Y., Wang, B., Shi, F., Wang, B., Bi, W., et al. Gated slot attention for efficient linear-time sequence modeling. *Advances in Neural Information Processing Systems*, 37: 116870–116898, 2024.

Zhong, M., Yin, D., Yu, T., Zaidi, A., Mutuma, M., Jha, R., Hassan Awadallah, A., Celikyilmaz, A., Liu, Y., Qiu, X., and Radev, D. QMSum: A New Benchmark for Query-based Multi-domain Meeting Summarization. In *North American Association for Computational Linguistics (NAACL)*, 2021.

A. Experimental Details

A.1. Model Configurations

Table 4 summarizes the model architectural specifications of the Qwen3 backbones (Yang et al., 2025a) and their corresponding hybrid attention models under the default 3:1 linear-to-full attention hybrid ratio. To convert grouped-query attention (GQA) in the pretrained full-attention layers into multi-head linear attention (MHA), we expand the pretrained key and value projection weights by repeating each KV head across its corresponding query group during initialization, while directly reusing the pretrained query and output projection weights. Each linear-attention branch is further equipped with RMSNorm on queries, keys, and branch outputs, an element-wise sigmoid output gate initialized to zero, and rotary position embeddings (RoPE) (Su et al., 2024). All newly introduced parameters in the linear-attention branches (e.g., decay/forget gates and output gates) are trained from scratch starting in Stage 1.

Table 6 presents the complete layer-importance rankings and selected full-attention layers produced by FlashMorph2 and baseline layer-selection methods across different backbones and linear-attention variants. Layer indices are zero-based, with boxed indices indicating the layers that retain full attention and all remaining layers replaced by linear attention.

Table 4. **Hybrid Attention Model Configurations across Different Scales.**

Attribute	0.6B	1.7B	8B	30B-A3B
#layers	28	28	36	48
#full-attention layers	7	7	9	12
#linear-attention layers	21	21	27	36
hidden dimension	1024	2048	4096	2048
dense FFN dimension	3072	6144	12288	—
#experts	—	—	—	128
#active experts	—	—	—	8
expert FFN dimension	—	—	—	768
head dimension	128	128	128	128
#query heads	16	16	32	32
#full-attention KV heads	8	8	8	4
#linear-attention KV heads	16	16	32	32

A.2. Training Configurations

Table 5 summarizes the training hyperparameters across all four stages of the Transformer-to-hybrid conversion pipeline: layerwise alignment (Stage 1), *on-policy* layer selection (Stage 2), logits distillation (Stage 3), and long-context finetuning (Stage 4). All stages are trained on randomly sampled subsets of DCLM (Li et al., 2024) using the Adam optimizer (Kingma & Ba, 2014) without weight decay. In Stage 1, only the linear-attention branches are updated while all pretrained backbone weights remain frozen. In Stage 2, we initialize all layer-selection scores ϕ to 0.5 and jointly optimize ϕ and the model weights θ under a Warmup–Stable–Decay (WSD) (Hu et al., 2024) learning rate schedule. After Stage 2, we retain the top-K full-attention layers, reinitialize the selected hybrid attention model with the Stage 1 weights θ_0, and train all model parameters end to end during Stage 3 and Stage 4. To ensure a fair comparison, all layer-selection methods share the same Stage 1 checkpoint and identical Stage 3–4 recovery training settings.

Table 5. **Training Configurations across the Four-Stage Transformer-to-hybrid Conversion Pipeline.**

Hyperparameter	Stage 1	Stage 2	Stage 3	Stage 4
training tokens	320M	20M	1B	1B
layer-selection score LR	—	$2e{-2} \rightarrow 2e{-3}$	—	—
model weight LR	$1e{-3} \rightarrow 1e{-5}$	$1e{-4} \rightarrow 1e{-5}$	$1e{-4} \rightarrow 1e{-5}$	$1e{-5}$
LR scheduler	cosine	WSD	cosine	constant
weight decay	0	0	0	0
sequence length	512	<4,096	512	16,384
batch size	32	32	96	128
warmup steps	50	50	50	50
decay steps	19,950	50	19,950	0
training steps	20,000	250	20,000	500

Table 6. **Complete Layer Importance Rankings across Backbones and Linear-Attention Variants.** Baseline rankings are adopted from Lan et al. (2026). Boxed zero-based indices denote retained full-attention layers (3:1 ratio), and red indices indicate selections differing from FlashMorph2.

Attn.	Method	Layer indices (most important → least important)
Qwen3-0.6B Backbone		
Lightning	**FlashMorph2**	[21, 11, 16, 6, 20, 19, 25,] 24, 2, 26, 8, 18, 0, 3, 17, 13, 22, 1, 9, 12, 4, 14, 5, 23, 27, 15, 10, 7
	FlashMorph	[1, 16, 21, 11, 19, 24, 0,] 25, 18, 2, 6, 8, 20, 3, 26, 13, 9, 23, 22, 10, 14, 17, 4, 15, 12, 27, 5, 7
	HALO	[10, 21, 9, 5, 11, 1, 13,] 16, 25, 12, 19, 24, 6, 18, 15, 8, 2, 26, 14, 23, 0, 27, 20, 7, 17, 22, 3, 4
	KL-LS	[21, 16, 19, 20, 18, 22, 25,] 24, 26, 17, 23, 14, 8, 12, 6, 13, 11, 15, 3, 2, 9, 4, 1, 5, 0, 10, 7, 27
GLA	**FlashMorph2**	[21, 16, 11, 20, 6, 19, 25,] 2, 24, 4, 3, 1, 26, 8, 18, 17, 13, 14, 0, 22, 12, 5, 9, 23, 15, 7, 10, 27
	FlashMorph	[21, 16, 11, 1, 2, 6, 25,] 19, 24, 18, 20, 8, 0, 26, 13, 22, 3, 23, 9, 10, 14, 17, 12, 15, 4, 7, 27, 5
	HALO	[8, 21, 13, 1, 24, 6, 25,] 19, 18, 11, 10, 16, 9, 12, 5, 15, 2, 26, 27, 23, 20, 4, 0, 14, 3, 22, 7, 17
	KL-LS	[21, 16, 20, 19, 18, 24, 8,] 22, 17, 6, 23, 25, 26, 14, 12, 11, 13, 2, 3, 15, 9, 4, 1, 0, 5, 10, 7, 27
GDN	**FlashMorph2**	[21, 11, 16, 6, 20, 2, 19,] 25, 24, 8, 13, 17, 26, 0, 3, 12, 18, 1, 14, 4, 9, 22, 23, 15, 7, 5, 10, 27
	FlashMorph	[1, 11, 21, 16, 19, 18, 24,] 6, 25, 8, 2, 20, 0, 13, 26, 3, 9, 23, 14, 22, 10, 12, 17, 15, 4, 7, 5, 27
	HALO	[10, 5, 21, 6, 19, 24, 18,] 11, 9, 13, 12, 25, 8, 2, 16, 26, 1, 4, 15, 0, 17, 7, 23, 14, 27, 3, 22, 20
	KL-LS	[21, 16, 19, 20, 18, 6, 22,] 24, 25, 11, 8, 14, 26, 12, 23, 17, 13, 3, 2, 9, 1, 4, 15, 5, 0, 10, 27, 7
Qwen3-1.7B Backbone		
Lightning	**FlashMorph2**	[21, 11, 16, 20, 19, 6, 17,] 2, 13, 0, 18, 8, 24, 25, 3, 26, 12, 1, 14, 9, 22, 15, 4, 10, 5, 23, 7, 27
	FlashMorph	[1, 16, 13, 11, 21, 3, 8,] 20, 6, 19, 18, 9, 14, 24, 2, 17, 0, 10, 15, 25, 23, 26, 22, 12, 4, 5, 7, 27
	HALO	[3, 14, 9, 2, 6, 16, 21,] 25, 24, 8, 23, 12, 11, 26, 27, 19, 18, 17, 7, 15, 4, 13, 1, 20, 10, 22, 0, 5
	KL-LS	[21, 16, 20, 19, 18, 22, 24,] 26, 25, 23, 17, 6, 14, 8, 12, 13, 11, 3, 15, 2, 9, 1, 4, 0, 5, 10, 7, 27
	PostNAS	[0, 21, 25, 19, 6, 11, 9,] 24, 12, 2, 26, 16, 17, 23, 18, 4, 7, 3, 14, 20, 1, 27, 10, 13, 8, 22, 15, 5
GLA	**FlashMorph2**	[21, 16, 11, 20, 19, 6, 2,] 17, 1, 13, 0, 8, 3, 25, 12, 14, 24, 18, 5, 4, 9, 26, 22, 7, 15, 23, 10, 27
	FlashMorph	[1, 11, 21, 3, 13, 14, 16,] 19, 9, 20, 6, 18, 2, 0, 24, 8, 10, 17, 25, 15, 23, 26, 12, 22, 4, 7, 27, 5
	HALO	[3, 14, 6, 8, 4, 25, 11,] 21, 16, 2, 24, 18, 26, 17, 19, 23, 1, 27, 12, 0, 15, 13, 7, 20, 9, 22, 10, 5
	KL-LS	[21, 16, 20, 19, 18, 24, 22,] 6, 8, 26, 11, 25, 23, 14, 12, 17, 3, 13, 2, 1, 9, 4, 15, 0, 5, 10, 27, 7
	PostNAS	[0, 21, 25, 19, 6, 11, 9,] 24, 12, 2, 26, 16, 17, 23, 18, 4, 7, 3, 14, 20, 1, 27, 10, 13, 8, 22, 15, 5
GDN	**FlashMorph2**	[21, 11, 16, 20, 6, 17, 19,] 2, 13, 8, 12, 14, 3, 0, 1, 24, 25, 18, 26, 9, 22, 4, 23, 15, 7, 10, 5, 27
	FlashMorph	[1, 11, 13, 21, 16, 14, 6,] 10, 20, 19, 2, 18, 8, 24, 0, 3, 9, 25, 15, 22, 17, 26, 23, 12, 4, 27, 7, 5
	HALO	[3, 2, 25, 21, 11, 14, 6,] 4, 12, 8, 16, 18, 24, 17, 19, 23, 7, 26, 27, 9, 1, 20, 22, 13, 0, 15, 5, 10
	KL-LS	[21, 16, 20, 19, 6, 11, 18,] 24, 22, 8, 25, 26, 12, 14, 23, 13, 2, 3, 17, 1, 9, 4, 15, 5, 10, 0, 27, 7
	PostNAS	[0, 21, 25, 19, 6, 11, 9,] 24, 12, 2, 26, 16, 17, 23, 18, 4, 7, 3, 14, 20, 1, 27, 10, 13, 8, 22, 15, 5
Qwen3-8B Backbone		
Lightning	**FlashMorph2**	[13, 29, 24, 20, 33, 15, 22, 17, 9,] 3, 19, 23, 0, 18, 16, 5, 32, 21, 1, 2, 34, 11, 4, 7, 8, 6, 26, 14, 12, 10, 31, 27, 28, 25, 30, 35
	FlashMorph	[1, 0, 7, 22, 29, 13, 3, 15, 24,] 20, 9, 19, 33, 34, 12, 2, 17, 8, 32, 16, 14, 23, 21, 35, 18, 5, 31, 11, 10, 4, 28, 26, 30, 6, 25, 27
	HALO	[7, 6, 0, 24, 8, 33, 1, 12, 34,] 22, 2, 15, 26, 9, 20, 16, 29, 31, 21, 4, 17, 30, 5, 35, 11, 32, 3, 25, 27, 14, 18, 19, 28, 13, 23, 10
Qwen3-30B-A3B Backbone		
Lightning	**FlashMorph2**	[42, 45, 0, 1, 7, 22, 6, 43, 39, 5, 11, 2,] 38, 27, 10, 26, 24, 34, 8, 13, 18, 30, 9, 15, 23, 12, 37, 17, 14, 3, 4,16, 25, 20, 46,41, 36, 21, 44, 31, 33, 19, 28, 29, 40, 35, 47, 32
	FlashMorph	[37, 2, 4, 3, 1, 38, 5, 0, 6, 42, 8, 26,] 13, 22, 18, 21, 11, 9, 14, 20, 15, 25, 36, 24, 10, 41, 12, 34, 47, 17, 23, 30, 7, 45, 39, 16, 19, 29, 35, 43, 33, 46, 27, 44, 40, 32, 28, 31
	HALO	[5, 45, 42, 8, 12, 43, 36, 19, 34, 41, 18, 22,] 46, 27, 37, 4, 31, 39, 13, 9, 35, 10, 24, 7, 33, 25, 47, 23, 28, 17, 21, 11, 14, 0, 30, 15, 20, 26, 40, 3, 1, 6, 32, 38, 2, 29, 16, 44

B. Evaluation

Long-context Capabilities. We evaluate long-context retrieval and understanding on RULER (Hsieh et al., 2024) and LongBench (Bai et al., 2024). For RULER, we evaluate across all 13 tasks grouped into four categories: Needle-in-a-Haystack (NIAH), including Single-NIAH (SG1–SG3), Multi-keys NIAH (MK1–MK3), Multi-queries NIAH (MQ), and Multi-values NIAH (MV); Variable Tracking (VT); Aggregation (Agg.), including Common Words Extraction (CWE) and Frequent Words Extraction (FWE); and Question Answering, including QA1 (Rajpurkar et al., 2016), and QA2 (Yang et al., 2018). For LongBench, we evaluate on 14 tasks spanning five categories: Single-Document QA, including NarrativeQA (NQA; Kočiskỳ et al. 2018), Qasper (QQA; Dasigi et al. 2021), and MultiFieldQA (MFQ; Bai et al. 2024); Multi-Document QA, including HotpotQA (HQA; Yang et al. 2018), 2WikiMultihopQA (2WM; Ho et al. 2020), and MuSiQue (Mus; Trivedi et al. 2022); Summarization, including GovReport (GvR; Huang et al. 2021), QMSum (QMS; Zhong et al. 2021), and MultiNews (MNs; Fabbri et al. 2019); Few-shot Learning (Few-shot), including TREC (TRC; Li & Roth 2002), TriviaQA (TQA; Joshi et al. 2017), and SAMSum (SSM; Gliwa et al. 2019); and Code Completion (Code), including Long Code Completion (LCC; Guo et al. 2023) and RepoBench-P (RBP; Liu et al. 2024).

Recall-intensive Tasks. We evaluate real-world recall-intensive tasks across six datasets: FDA (Arora et al., 2023) for PDF key-value retrieval, SWDE (Lockard et al., 2019) for structured HTML relation extraction, and four question-answering datasets including SQuAD (SQD; Rajpurkar et al. 2018), Natural Questions (NQ; Kwiatkowski et al. 2019), TriviaQA (TQA; Joshi et al. 2017), and DROP (Dua et al., 2019).

Language Modeling and Commonsense Reasoning. Following previous work (Gu & Dao, 2023; Yang et al., 2025b), we evaluate our method on language modeling perplexity and zero-shot commonsense reasoning benchmarks: WikiText (Wiki.; Merity et al. 2016), LAMBDA (LMB.; Paperno et al. 2016), ARC-easy (ARC-e) and ARC-challenge (ARC-c) (Clark et al., 2018), PIQA (Bisk et al., 2020), HellaSwag (Hella.; Zellers et al. 2019), WinoGrande (Wino.; Sakaguchi et al. 2021), SIQA (Sap et al., 2019), and BoolQ (Clark et al., 2019).

C. Detailed Benchmark Results

Tables 7, 8, and 9 present the detailed per-task results on RULER, LongBench, and language modeling and zero-shot commonsense reasoning benchmarks, respectively. In addition to the Qwen3-0.6B and Qwen3-1.7B configurations reported in the main text, we further scale our evaluation to larger dense and MoE backbones, including Qwen3-8B and Qwen3-30B-A3B with Lightning Attention. Across these larger model scales, FlashMorph2 continues to consistently outperform competing layer-selection baselines on both RULER and LongBench, while preserving strong language modeling perplexity and zero-shot commonsense reasoning accuracy.

Table 7. **Detailed Results on RULER.** Per-task scores and average across all 13 RULER tasks. For each backbone and linear-attention variant, the best and second-best scores among layer-selection methods are highlighted.

Attn.	Method	SG1	SG2	SG3	MK1	MK2	MK3	MQ	MV	VT	CWE	FWE	QA1	QA2	Avg
0.6B Backbone															
Full	Qwen3	100.00	100.00	99.60	99.20	99.00	63.60	97.70	95.90	95.32	52.96	58.60	41.32	25.80	79.15
Lightning	Uniform	99.80	99.80	95.60	65.80	15.40	0.40	53.25	58.95	1.56	46.62	56.67	36.58	23.40	50.29
	KL-LS	100.00	100.00	73.60	42.20	30.80	0.20	48.35	40.75	2.20	18.52	50.40	37.32	26.00	43.87
	HALO	83.80	100.00	97.40	62.00	13.20	1.00	60.05	58.60	1.20	48.04	52.93	38.62	26.60	49.50
	FlashMorph	100.00	100.00	94.20	79.00	43.40	1.40	64.50	66.00	10.84	54.80	55.27	42.25	28.40	56.93
	FlashMorph2	100.00	100.00	88.00	80.00	31.60	3.40	65.30	79.85	17.24	51.80	65.13	45.18	27.80	58.10
GLA	Uniform	100.00	100.00	88.40	74.80	26.20	1.00	52.95	53.90	4.44	50.48	47.13	32.87	21.40	50.27
	KL-LS	100.00	100.00	98.20	81.20	32.00	2.20	68.45	61.55	1.80	37.06	49.87	35.90	24.00	53.25
	HALO	100.00	100.00	97.80	77.60	15.40	0.80	63.85	71.30	2.08	58.54	36.20	36.97	23.00	52.58
	FlashMorph	99.80	100.00	97.60	88.00	18.60	2.80	71.65	73.50	5.52	53.06	49.47	39.92	25.20	55.78
	FlashMorph2	100.00	100.00	89.40	92.40	61.60	4.20	83.05	78.65	7.20	57.84	47.40	44.25	29.00	61.15
GDN	Uniform	99.60	99.80	94.60	70.80	30.20	1.40	56.30	60.00	3.76	48.38	52.07	40.32	26.20	52.57
	KL-LS	100.00	100.00	92.80	81.80	56.00	3.00	72.20	63.00	2.96	41.14	56.80	42.12	27.40	56.86
	HALO	100.00	100.00	85.00	69.20	25.80	2.00	54.40	45.20	20.28	35.36	54.26	39.67	27.20	50.64
	FlashMorph	100.00	100.00	91.20	89.00	52.00	3.20	75.15	69.30	15.60	55.08	58.00	46.10	29.20	60.29
	FlashMorph2	100.00	100.00	98.80	95.20	79.60	4.40	63.80	72.85	18.68	60.88	60.20	45.18	27.80	63.65
1.7B Backbone															
Full	Qwen3	100.00	100.00	99.80	99.80	98.20	98.40	99.90	99.95	95.72	88.62	66.33	51.88	41.00	87.66
Lightning	Uniform	100.00	100.00	87.80	73.60	67.60	2.00	75.75	64.10	29.16	42.20	61.73	46.45	35.00	60.41
	PostNAS	100.00	99.80	97.40	76.40	58.20	1.80	55.15	56.65	23.16	63.36	66.53	48.15	37.60	60.32
	KL-LS	100.00	100.00	70.80	49.60	42.00	0.60	44.55	48.05	0.24	42.02	46.00	42.05	32.40	47.56
	HALO	100.00	100.00	93.00	84.60	45.00	1.80	77.30	81.40	29.48	59.94	67.27	52.02	35.80	63.66
	FlashMorph	99.80	100.00	98.00	86.80	56.20	2.00	86.80	91.15	31.64	58.78	61.33	52.22	37.80	66.35
	FlashMorph2	100.00	100.00	96.80	80.40	87.00	7.80	81.85	85.80	41.24	58.62	67.20	56.08	39.40	69.40
GLA	Uniform	100.00	100.00	86.40	83.20	65.00	2.00	68.45	77.85	20.56	58.28	51.67	48.57	35.40	61.34
	PostNAS	99.60	100.00	99.20	92.00	48.40	3.60	70.20	86.60	19.76	56.04	47.13	52.25	39.60	62.64
	KL-LS	100.00	100.00	65.00	43.80	48.00	2.40	43.15	69.15	11.60	54.28	46.67	41.98	35.20	50.86
	HALO	100.00	99.60	95.60	68.40	25.80	0.80	58.00	67.40	25.84	63.52	60.87	45.17	29.60	56.97
	FlashMorph	100.00	99.60	95.60	87.00	80.60	5.00	83.20	92.45	30.52	57.50	55.47	49.65	36.00	67.12
	FlashMorph2	100.00	100.00	98.40	94.40	63.00	8.00	76.75	92.60	22.28	61.22	62.40	58.22	37.80	67.31
GDN	Uniform	100.00	100.00	97.20	75.80	64.40	3.20	76.45	87.45	63.32	51.70	67.13	51.65	39.20	67.50
	PostNAS	100.00	100.00	100.00	88.80	75.00	6.00	74.35	89.45	53.16	61.54	68.40	53.25	36.60	69.73
	KL-LS	100.00	100.00	99.00	89.40	60.00	9.60	87.55	92.10	68.12	61.40	69.47	55.22	40.60	71.73
	HALO	100.00	100.00	98.60	88.60	50.20	2.40	81.40	88.25	62.68	58.88	67.73	54.88	39.00	68.66
	FlashMorph	100.00	99.80	100.00	90.60	66.80	8.60	92.50	96.05	66.76	60.84	70.93	55.25	41.80	73.07
	FlashMorph2	100.00	100.00	98.60	90.00	91.80	12.80	90.30	95.15	58.80	63.90	68.67	55.52	42.40	74.46
8B Backbone															
Full	Qwen3	100.00	100.00	100.00	100.00	100.00	100.00	100.00	99.85	99.96	99.00	93.73	71.32	56.60	93.88
Lightning	Uniform	100.00	99.80	99.80	87.80	76.80	7.20	90.85	91.25	65.20	61.22	81.13	58.50	43.00	74.04
	HALO	100.00	100.00	96.40	56.40	22.40	1.80	86.40	90.55	55.32	67.86	55.07	51.35	38.20	63.21
	FlashMorph	100.00	100.00	99.80	90.20	80.80	17.80	95.95	97.75	51.00	58.66	70.53	60.98	43.40	74.38
	FlashMorph2	100.00	100.00	100.00	90.00	71.20	28.60	95.30	95.05	83.44	64.44	68.93	54.68	44.40	76.62
30B-A3B Backbone															
Full	Qwen3	100.00	100.00	100.00	100.00	100.00	100.00	91.80	31.80	99.96	59.28	67.00	58.58	39.80	80.63
Lightning	Uniform	99.60	100.00	97.00	58.80	16.20	2.20	85.55	87.45	32.76	57.46	61.47	43.05	40.00	60.12
	HALO	100.00	100.00	98.80	70.20	15.00	1.40	81.70	93.15	36.08	63.10	57.00	34.82	38.00	60.71
	FlashMorph	100.00	100.00	99.20	79.20	53.60	6.20	97.75	97.40	19.04	61.76	69.40	45.77	33.80	66.39
	FlashMorph2	100.00	100.00	90.80	80.20	54.80	7.20	96.35	96.60	44.96	65.78	71.47	45.73	33.20	68.24

Table 8. **Detailed Results on LongBench.** Per-task scores and average across 14 LongBench tasks spanning five categories. For each backbone and linear-attention variant, the best and second-best scores among layer-selection methods are highlighted.

Attn.	Method	Single-Doc QA			Multi-Doc QA			Summarization			Few-shot			Code		Avg
		NQA	QQA	MFQ	HQA	2WM	Mus	GvR	QMS	MNs	TRC	TQA	SSM	LCC	RBP	
0.6B Backbone																
Full	Qwen3	10.04	14.12	39.49	26.25	25.91	10.41	25.05	19.10	14.03	70.50	73.49	38.34	31.31	35.21	30.95
Lightning	Uniform	14.34	17.78	27.59	20.21	20.80	8.38	15.87	20.36	17.52	56.00	55.80	33.95	30.14	36.42	26.80
	KL-LS	11.50	18.76	25.30	24.76	26.37	9.32	17.15	20.16	18.50	48.50	49.03	32.44	36.22	41.20	27.09
	HALO	10.33	19.44	27.39	22.66	24.81	7.25	13.70	20.58	17.05	50.00	53.28	30.87	32.19	37.11	26.19
	FlashMorph	15.05	23.34	31.69	24.95	25.29	7.24	18.30	20.50	19.86	51.00	55.01	33.27	33.46	40.01	28.50
	FlashMorph2	14.05	23.09	31.56	23.72	25.44	9.02	19.46	20.89	19.06	52.50	56.17	33.36	36.71	43.23	29.16
GLA	Uniform	11.70	19.82	27.78	16.78	22.31	7.73	15.30	18.50	16.67	58.00	49.11	33.52	27.43	31.69	25.45
	KL-LS	14.09	24.84	26.66	19.46	24.18	9.28	20.02	18.98	18.58	52.50	56.56	34.66	30.25	31.32	27.24
	HALO	11.96	19.28	28.86	24.68	24.24	8.76	15.70	20.54	15.87	48.50	56.62	35.38	26.17	31.83	26.31
	FlashMorph	12.95	26.31	31.66	23.87	24.08	9.02	16.53	20.29	18.37	46.50	51.67	33.00	28.70	30.58	26.68
	FlashMorph2	15.23	23.66	31.59	21.31	24.31	8.18	18.92	20.45	19.02	49.50	55.75	34.01	30.65	33.31	27.57
GDN	Uniform	14.49	20.92	28.15	21.79	18.33	9.09	16.61	21.39	15.66	60.00	52.86	34.42	30.39	39.99	27.43
	KL-LS	13.95	21.14	27.59	23.16	20.13	8.45	19.84	21.22	20.00	63.00	53.66	33.16	35.42	42.28	28.79
	HALO	14.19	22.03	26.13	20.48	21.66	6.58	18.95	20.55	15.93	55.50	49.76	34.16	33.23	40.70	27.13
	FlashMorph	14.74	23.05	32.31	24.22	24.19	8.31	19.24	20.70	19.10	58.50	54.45	33.39	33.33	42.55	29.15
	FlashMorph2	15.13	23.42	34.07	24.38	25.17	10.57	19.03	22.11	19.31	62.00	51.17	34.71	33.34	42.41	29.77
1.7B Backbone																
Full	Qwen3	14.05	23.99	40.64	17.17	14.29	7.91	25.73	21.95	7.91	71.50	84.58	41.97	44.62	37.66	32.43
Lightning	Uniform	16.79	17.56	29.21	22.92	21.97	9.84	18.53	21.82	20.47	66.00	66.36	40.70	34.58	40.60	30.53
	PostNAS	15.19	15.52	31.36	29.72	25.67	10.78	18.48	22.03	19.25	61.50	65.41	39.27	33.25	40.37	30.56
	KL-LS	17.38	16.28	26.54	22.00	23.47	8.81	19.07	21.66	22.61	61.00	68.24	38.91	38.97	45.87	30.77
	HALO	18.29	15.43	27.20	29.23	23.89	11.15	18.36	22.07	20.47	61.50	72.16	40.09	37.67	42.58	31.44
	FlashMorph	16.43	15.07	29.25	29.05	20.66	13.48	17.01	22.55	20.51	63.00	75.49	39.19	36.43	37.75	31.13
	FlashMorph2	18.98	15.87	33.28	31.73	25.29	13.65	20.62	22.13	21.15	67.50	72.12	39.54	38.60	46.47	33.35
GLA	Uniform	14.59	18.63	29.78	26.08	20.53	8.99	15.71	19.62	20.86	68.50	66.54	37.65	39.12	37.87	30.32
	PostNAS	16.50	25.50	30.69	26.51	24.95	13.83	18.98	13.97	18.68	65.50	65.56	39.16	37.51	40.04	31.24
	KL-LS	15.80	19.36	24.58	25.69	24.08	11.12	19.58	13.84	19.99	63.00	69.03	40.32	43.77	43.61	30.98
	HALO	15.00	23.16	28.97	20.45	23.26	9.84	15.63	12.23	19.92	66.00	66.79	35.79	38.09	43.08	29.87
	FlashMorph	16.28	21.88	27.63	23.04	23.23	11.22	16.68	8.89	21.74	66.50	73.01	37.16	42.78	44.86	31.06
	FlashMorph2	16.91	23.84	29.39	26.80	24.49	10.22	18.91	13.17	15.44	68.00	72.55	40.05	42.21	45.37	31.95
GDN	Uniform	18.45	23.29	16.85	28.07	25.92	12.82	18.96	22.06	20.42	70.50	67.77	39.43	34.82	35.99	31.10
	PostNAS	18.63	25.57	31.52	26.05	21.63	14.48	21.68	21.93	21.57	68.50	65.95	40.02	34.50	38.80	32.20
	KL-LS	19.11	25.78	28.82	25.82	21.50	11.90	21.50	21.66	21.08	68.50	70.47	39.90	38.40	43.72	32.73
	HALO	19.66	24.59	31.15	30.99	22.87	13.03	20.08	18.68	16.90	65.50	66.68	40.19	39.07	42.80	32.30
	FlashMorph	18.02	27.09	28.95	30.89	22.78	13.02	20.06	22.16	20.13	64.50	76.95	37.08	34.65	40.39	32.62
	FlashMorph2	19.32	26.05	31.79	27.60	21.71	12.27	22.02	21.05	21.09	68.50	76.51	38.47	40.30	46.66	33.81
8B Backbone																
Full	Qwen3	24.29	43.66	48.42	41.25	33.59	14.02	29.05	19.43	22.98	71.00	89.82	44.30	67.55	61.76	43.65
Lightning	Uniform	20.91	31.88	40.29	38.04	27.73	17.10	22.85	23.52	20.81	65.00	69.51	41.17	44.64	43.51	36.21
	HALO	17.08	30.97	28.25	37.02	27.51	16.86	23.17	22.42	19.27	60.50	62.56	38.95	42.18	41.99	33.48
	FlashMorph	22.78	27.69	33.77	40.54	31.24	17.79	26.08	23.91	21.75	68.50	75.09	42.45	43.11	41.83	36.89
	FlashMorph2	27.13	30.63	40.56	39.10	35.67	20.11	27.30	24.20	22.91	70.50	77.82	43.68	49.05	42.34	39.36
30B-A3B Backbone																
Full	Qwen3	21.14	42.91	31.69	17.18	18.17	6.85	29.92	18.65	22.54	73.00	87.55	44.16	71.17	68.06	39.50
Lightning	Uniform	16.50	30.88	27.93	40.40	32.85	22.35	21.66	21.61	17.30	60.50	56.33	41.43	45.50	45.32	34.33
	HALO	17.87	21.02	26.58	37.43	29.30	15.89	26.74	21.73	13.91	47.50	68.24	39.01	47.06	50.06	33.02
	FlashMorph	20.07	32.30	29.90	37.81	28.58	15.20	25.20	21.06	13.43	65.50	69.76	40.02	50.33	50.53	35.69
	FlashMorph2	21.43	33.71	28.90	37.76	33.03	13.01	20.65	21.99	17.81	62.50	78.09	42.10	59.42	55.96	37.60

Table 9. **Performance on Language Modeling and Zero-Shot Commonsense Reasoning.** Avg denotes the average accuracy across all eight zero-shot commonsense reasoning tasks. For each backbone and linear-attention variant, the best and second-best results among layer-selection methods are highlighted.

Attn.	Method	Wiki. ppl ↓	LMB. ppl ↓	LMB. acc ↑	ARC-e acc ↑	ARC-c acc_n ↑	PIQA acc ↑	Hella. acc_n ↑	Wino. acc ↑	SIQA acc ↑	BoolQ acc ↑	Avg
0.6B Backbone												
Full	Qwen3	26.18	24.72	41.02	60.65	33.96	67.90	47.24	55.96	39.05	64.40	51.27
Lightning	Uniform	22.58	12.94	49.18	65.19	33.79	68.44	49.01	57.85	40.79	61.41	53.21
	KL-LS	31.43	12.44	50.01	64.81	34.64	68.28	48.75	58.25	41.25	60.73	53.34
	HALO	23.02	12.74	49.62	65.91	34.47	68.28	48.99	56.35	40.99	63.94	53.57
	FlashMorph	24.40	12.56	49.87	66.46	34.39	68.93	49.26	56.12	41.15	61.16	53.42
	FlashMorph2	20.20	12.52	50.32	65.99	34.50	68.88	49.15	57.46	41.20	61.74	53.66
GLA	Uniform	21.34	12.53	50.67	65.78	33.70	69.86	49.19	55.88	41.35	63.03	53.68
	KL-LS	29.25	11.93	51.19	65.95	36.01	69.26	49.36	55.80	41.61	65.11	54.29
	HALO	27.29	12.15	51.14	66.04	34.47	68.72	49.20	56.04	41.30	66.67	54.20
	FlashMorph	20.41	12.41	50.61	65.74	33.96	68.93	49.34	54.93	41.20	64.59	53.66
	FlashMorph2	20.14	12.37	50.15	64.94	34.39	69.37	49.52	56.27	41.50	63.46	53.70
GDN	Uniform	20.17	12.48	50.24	65.78	34.13	68.99	49.43	55.72	41.20	63.49	53.63
	KL-LS	19.83	12.16	50.65	65.74	35.32	68.88	49.64	56.67	41.81	64.25	54.12
	HALO	20.05	12.30	50.69	65.95	35.58	68.72	49.27	55.17	41.10	64.40	53.86
	FlashMorph	19.64	11.87	51.33	66.46	35.92	68.93	49.56	55.25	41.61	65.38	54.31
	FlashMorph2	19.72	12.03	51.12	66.08	34.47	68.99	49.80	55.25	41.66	65.78	54.14
1.7B Backbone												
Full	Qwen3	20.98	12.19	51.41	72.22	42.66	72.31	60.36	61.56	45.04	77.65	60.40
Lightning	Uniform	15.26	6.91	59.64	76.73	44.45	74.16	62.16	62.19	44.01	77.86	62.65
	PostNAS	14.79	6.85	59.65	77.53	44.54	74.32	62.33	63.93	44.98	77.22	63.06
	KL-LS	14.87	6.93	59.34	75.42	46.93	74.32	62.16	62.27	44.63	73.91	62.37
	HALO	14.91	7.05	58.72	76.47	44.97	74.81	62.03	61.01	43.86	77.46	62.42
	FlashMorph	14.84	6.98	59.07	76.35	45.22	74.27	62.30	64.09	44.27	77.52	62.89
	FlashMorph2	14.71	6.77	59.54	76.64	45.65	74.43	62.51	62.98	43.91	77.55	62.90
GLA	Uniform	14.91	6.72	61.17	77.61	45.82	74.27	62.70	62.83	44.88	77.37	63.33
	PostNAS	14.69	6.68	60.28	76.81	45.22	74.05	62.67	64.09	44.63	76.97	63.09
	KL-LS	14.70	6.63	60.61	77.23	47.35	74.37	62.75	62.43	45.09	75.57	63.18
	HALO	14.92	6.68	62.56	76.98	45.99	74.32	62.56	62.98	44.47	75.55	63.18
	FlashMorph	14.73	6.64	60.59	76.89	45.73	74.65	62.71	62.59	44.78	78.17	63.26
	FlashMorph2	14.60	6.64	60.47	76.77	46.93	74.48	62.71	64.09	45.29	76.91	63.46
GDN	Uniform	14.65	6.81	60.37	75.60	45.56	74.48	62.75	62.19	44.01	77.80	62.85
	PostNAS	14.45	6.74	59.98	77.15	46.33	74.16	62.54	64.33	44.42	77.68	63.32
	KL-LS	14.35	6.68	60.06	76.60	45.56	74.86	63.00	62.75	44.78	76.94	63.07
	HALO	14.46	6.77	59.67	76.60	45.99	74.54	62.69	63.93	44.52	77.43	63.17
	FlashMorph	14.40	6.78	59.83	76.85	45.05	74.16	62.84	63.14	43.76	78.10	62.97
	FlashMorph2	14.33	6.67	60.26	76.89	45.05	74.32	62.82	62.75	44.52	77.71	63.04
8B Backbone												
Full	Qwen3	12.20	4.60	65.09	83.38	56.48	76.61	74.88	67.48	51.69	86.61	70.28
Lightning	Uniform	11.15	4.27	68.95	81.99	55.46	78.45	74.22	71.59	49.08	82.20	70.24
	HALO	11.27	4.20	68.99	81.86	55.97	78.89	74.18	71.59	48.82	81.56	70.23
	FlashMorph	11.09	4.26	68.91	81.65	56.06	78.89	74.43	71.74	49.59	83.03	70.54
	FlashMorph2	10.87	4.15	70.35	82.03	55.72	78.24	74.70	73.01	49.13	82.11	70.66
30B-A3B Backbone												
Full	Qwen3	10.89	4.13	65.63	79.29	55.72	79.43	77.69	70.09	49.49	88.75	70.76
Lightning	Uniform	10.20	3.77	71.78	82.28	57.42	79.60	77.02	72.77	48.93	84.34	71.77
	HALO	10.11	3.67	72.40	82.95	58.45	79.49	76.82	73.80	48.46	83.79	72.02
	FlashMorph	10.28	3.77	71.84	82.91	57.59	79.65	76.78	74.03	48.31	83.61	71.84
	FlashMorph2	10.11	3.75	72.48	83.21	57.42	79.76	77.03	73.56	48.98	83.24	71.96