PAYING ATTENTION TO HYBRID ATTENTION: UNTANGLING THE ISSUES WITH CONVERSION METHODS

Anonymous authors

000

001

002003004

006

008 009

010 011

012

013

014

015

016

017

018

019

021

024

025

026027028

029

031

033

034

037

038

040

041

042

043

044

046

047

048

049

050 051

052

Paper under double-blind review

ABSTRACT

Transformers' quadratic computational complexity limits their scalability despite remarkable performance. While linear attention reduces this to linear complexity, pre-training such models from scratch remains, in most cases, prohibitively expensive. Recent post-training linearisation methods convert pre-trained Transformers to linear models efficiently, often using hybrid approaches that combine linear attention with sliding-window softmax. We identify a critical flaw: existing hybrid methods inadvertently bypass the linear component, relying almost entirely on SWA. Component-level diagnostics reveal this previously undetected behaviour stems from overlooked evaluation practices on common-sense benchmarks. We propose three solutions to ensure balanced component usage: (i) inference-time hybridisation of linear-only conversions with sliding-window softmax; (ii) Hedge-CATs, combining attention-weight transfer with targeted LoRA fine-tuning; and (iii) Scheduled Sliding-window Dropout (SSD), which stochastically suppresses the softmax branch during training to prevent component collapse. Our methods maintain computational efficiency while recovering most base model performance and ensuring genuine linear attention adoption, restoring the validity of performance attributions in hybrid conversions.

1 Introduction

Transformers (Vaswani et al., 2017) have delivered state-of-the-art results across language, vision, and multimodal tasks, yet their quadratic attention cost in sequence length remains a central bottleneck for long-context inference and training. Linear attention (LA) (Katharopoulos et al., 2020), offers a compelling alternative by replacing the softmax kernel with linear feature maps that enable associative, streaming updates of a fixed-size recurrent state (Choromanski et al., 2021; Duman Keles et al., 2023; Banerjee et al., 2020; Peng et al., 2021; Qin et al., 2022). In principle, this reduces the asymptotic complexity of both memory and compute compared to softmax attention. In practice, however, fully pre-training LA models is costly (Liu et al., 2020), and performance often lags behind quadratic baselines trained with similar budgets due to limitations in representational complexity (Zhang et al., 2024b; Mercat et al., 2024).

A growing body of work focuses on circumventing the high cost of pre-training linear models via post-training linearisation (Zhang et al., 2024b;a; Lan et al., 2025; Mercat et al., 2024): converting a pre-trained quadratic Transformer into a fully linear or hybrid linear-softmax model. This approach amortises most of the cost into the pre-trained base model and performs a light "swap + adaptation" stage, in which softmax kernel is replaced with a learnable linear kernel, with the model then put through additional pre-training and/or supervised fine-tuning to recover performance. Such methods have been reported to require on the order of 0.02% (or less) of the data used to train the base model to recover performance. Existing conversion methods typically differ in (i) the LA formulation, (ii) whether and how the original weights are fine-tuned, and (iii) whether sliding-window softmax attention (SWA) (Beltagy et al., 2020; Zaheer et al., 2020) is retained alongside the linear path, yielding hybrid architectures.

Hybrid conversions are attractive as they pragmatically combine the representational capacity of SWA with the memory and computational efficiency of LA. This makes them powerful for long-context tasks. Yet, as we show, these reported gains can mask a critical failure mode: the model may lean almost entirely on the SWA path while effectively ignoring the linear component. This creates

misleading performance attribution: the hybrid is credited for "using LA" when, in reality, the model simply biases and adapts itself towards SWA entirely during the adaptation stage. The field lacks standard diagnostics to quantify each component's contribution, so such imbalances can remain hidden behind aggregate metrics.

Contributions (1) We identify and characterise a systematic issue in current hybrid attention conversions whereby models learn to ignore their LA component and over-rely on their SWA one, leading to misleading attribution of hybrid performance. (2) We provide component-level diagnostic that make this imbalance visible and reproducible across popular pre-trained base models on standard common-sense benchmarks. (3) We introduce three practical remedies: a zero-shot inference-time hybrid; HedgeCATs, which combines HedgeHog-style attention-weight transfer with brief LoRA fine-tuning; and Scheduled Sliding-window Dropout (SSD) to prevent component imbalance during training. We show that our proposed strategies recover most base-model performance while ensuring genuine use of the LA pathway, restoring attributional validity without sacrificing computational efficiency.

2 BACKGROUND & RELATED WORKS

2.1 LINEAR ATTENTION

Let $\mathbf{X} \in \mathbb{R}^{T \times d_{\text{model}}}$ be a sequence of length T with projections $\mathbf{Q} = \mathbf{X} \mathbf{W}_Q$, $\mathbf{K} = \mathbf{X} \mathbf{W}_K$, $\mathbf{V} = \mathbf{X} \mathbf{W}_V$, where $\mathbf{Q}, \mathbf{K} \in \mathbb{R}^{T \times d}$ and $\mathbf{V} \in \mathbb{R}^{T \times d_v}$. Standard softmax attention (Vaswani et al., 2017) computes

$$\mathbf{O} = \operatorname{softmax} \left(\frac{\mathbf{Q} \mathbf{K}^{\top}}{\sqrt{d}} \right) \mathbf{V}$$
 (1)

incurring $O(T^2)$ time and memory for the $T \times T$ similarity matrix. Using a kernel $\kappa(\mathbf{q}, \mathbf{k})$ such that $\kappa(\mathbf{q}, \mathbf{k}) \approx \phi(\mathbf{q})^{\top} \phi(\mathbf{k})$ for a non-negative feature map $\phi : \mathbb{R}^d \to \mathbb{R}^{d'}$ as proposed by Katharopoulos et al. (2020), we can avoid forming pairwise similarities by introducing global summaries:

$$\mathbf{o}_t = \frac{\sum_{i=1}^T \kappa(\mathbf{q}_t, \mathbf{k}_i) \mathbf{v}_i}{\sum_{i=1}^T \kappa(\mathbf{q}_t, \mathbf{k}_i)} = \frac{\phi(\mathbf{q}_t)^\top \mathbf{S}}{\phi(\mathbf{q}_t)^\top \mathbf{z}}, \quad \mathbf{S} = \sum_{i=1}^T \phi(\mathbf{k}_i) \mathbf{v}_i^\top, \quad \mathbf{z} = \sum_{i=1}^T \phi(\mathbf{k}_i).$$
(2)

These summaries can be accumulated in a single pass, so no $T \times T$ score matrix is materialised. This yields O(Tdd') time and O(dd') memory. Proposed ϕ in the literature include non-negative element-wise maps (such as $1 + \mathrm{ELU}(\mathbf{x})$ (Katharopoulos et al., 2020) and ReLU (Kasai et al., 2021), random feature maps (Choromanski et al., 2021), and exponential function approximations via low-order Taylor expansions (Duman Keles et al., 2023; Banerjee et al., 2020).

2.2 Hybrid Attention

Hybrid attention mechanisms combine softmax attention with LA through various architectural approaches. Some methods interleave full softmax attention layers with LA layers (Lieber et al., 2024; Dong et al., 2024), while others employ SWA combined with LA, either in interleaved layers (Ren et al., 2024) or integrated within the same Transformer block (Beltagy et al., 2020; Zhang et al., 2024a; Lan et al., 2025; Irie et al., 2025; Munkhdalai et al., 2024). The adoption of SWA offers the key advantage of preserving linear complexity throughout the entire model. Following other hybrid conversion approaches (Zhang et al., 2024a; Lan et al., 2025; Irie et al., 2025), our work focuses on combining SWA with LA within Transformer blocks to approximate full softmax attention. These methods typically employ a scaled linear combination of the two attention outputs, utilising learned (Munkhdalai et al., 2024; Zhang et al., 2024a), fixed (Lan et al., 2025), or data-dependent scaling factors (Behrouz et al., 2024; Irie et al., 2025). This can be summarised by scalar or vector mixing terms a, b, such that the hybrid attention output is given by:

$$ATTN(Q, K, V, a, b) = a \odot SWA(Q, K, V) + b \odot LA(Q, K, V)$$
(3)

2.3 LINEARISING PRE-TRAINED TRANSFORMERS

Although numerous linear (Schlag et al., 2021; Gu & Dao, 2023; Yang et al., 2023; 2024; Peng et al., 2025) and hybrid (Beltagy et al., 2020; Zhu et al., 2021; Lieber et al., 2024; Yang et al., 2024; Behrouz et al., 2024) Transformers have been developed, the majority are trained from scratch. The prohibitive cost of pre-training constrains most of these approaches to small model sizes (typically < 1B parameters) and makes them expensive to reproduce or scale when checkpoints are unavailable, directly undermining the central promise of linear Transformers: computational efficiency. This limitation renders post-training conversion methods particularly appealing, as they offer the potential for high performance recovery in large-scale Transformers at a fraction of pre-training costs (typically < 1%).

While LA is theoretically equivalent to softmax-based self-attention under feature map ϕ and kernel $\phi(q_t)^\top \phi(k_i) = \exp\left((q_t^\top k_i) \cdot D^{-1/2}\right)$ (Katharopoulos et al., 2020), such a (finite-dimensional) feature map does not currently exist. As such, converting a pre-trained Transformer to use LA requires some adjustments and fine-tuning to make up for the change in attention weights.

Proposed kernels are designed to ensure positive attention weights via non-linear activation functions (Katharopoulos et al., 2020; Kasai et al., 2021; Mercat et al., 2024; Zhang et al., 2024b;a). However, these non-negative activation functions run the risk of suppressing any negative signals and may unnecessarily constrain the learned mappings. To this end, Zhang et al. (2024b) and Zhang et al. (2024a) concatenate the negative mapping to the output along the head dimension, applying their respective non-linear activation function σ to each separately:

$$\phi(\mathbf{x}) = [\sigma(\mathbf{W}_{\phi}^{\top}\mathbf{x} + \mathbf{b}) \oplus \sigma(-\mathbf{W}_{\phi}^{\top}\mathbf{x} - \mathbf{b})]$$
(4)

Zhang et al. (2024b) showed that the softmax function's unique spikiness and monotonicity with respect to the Query-Key dot-product are hard to match when using previously proposed candidates for ϕ . They therefore learn an exponential feature map (Eq. 4 with $\sigma = \exp(\cdot)$), and train it via an attention-weights transfer objective that minimises cross-entropy between softmax attention weights and linear weights. This "weights-to-weights" stage is followed by a fine-tuning stage of the original model weights which they claim makes up for any approximation errors. LoLCATs (Zhang et al., 2024a) then later explores general conversion methods using LoRA (Hu et al., 2022) fine-tuning as well as hybrid attention methods. In parallel, SUPRA (Mercat et al., 2024) follows T2R (Kasai et al., 2021) in adopting a ReLU-activated feature map with standard language-model fine-tuning rather than an explicit attention-transfer loss. This keeps training simple since no new parameters are introduced.

Other conversion methods (Mao, 2022; Lan et al., 2025) focus on converting pre-trained Transformers to gated linear/recurrent blocks (GLA-style). Both these methods repurpose the attention block into a gated linear update; Liger (Lan et al., 2025) differs in explicitly retaining a local softmax branch. A summary of different linear kernels and transfer objectives for various conversion methods can be found in Table 5.

2.4 Training Interventions to Tackle Component Collapse

As explored in this paper, hybrid attention conversion models can learn to ignore the linear path and rely solely on SWA. Related work in other settings tackles analogous "path collapse" with structured dropout (Srivastava et al., 2014): dropping whole substructures during training so models learn to balance all their trained components. In Transformers, DropHead (Zhou et al., 2020) targets multi-head attention (Cordonnier et al., 2021) directly by stochastically dropping whole heads with a scheduled rate to prevent a few heads from monopolising computation. Other works explore similar ideas but applied to different subcomponents such as Transformer layers (Fan et al., 2020), experts (Chen et al., 2023) in Mixture-of-Experts models, and even incoming keys (Li et al., 2023).

3 IDENTIFYING THE ISSUES WITHIN HYBRID CONVERSION METHODS

In this section, we outline some issues we have found to occur in conversion methods which make use of a hybrid attention-based training objective. We start by re-implementing and ablating the LoLCATs framework, which is considered to be the state-of-the-art (SOTA) method for converting pre-trained

Transformers to use hybrid attention, as well as repeating such analysis in their own codebase and checkpoints. We complete these findings with a component-wise investigation of similar methods, ablating key components with those used in similar SOTA LA-only methods, one at a time, in order to identify the ones responsible for the issues observed in hybrid methods.

3.1 EXPERIMENTAL SET-UP

Feature Map Φ We adopt a learned feature map (Mercat et al., 2024; Zhang et al., 2024b;a), with $W_{\phi} \in \mathbb{R}^{h_d \times \frac{h_d}{2}}$ and a softmax activation function (ie. Eq 4 with $\sigma = \operatorname{softmax}(\cdot)$), and apply the RoPE embeddings to queries and keys prior to applying ϕ , as motivated by Zhang et al. (2024a).

Combining Sliding-Window and Linear Attention For our experiments, we implement Eq. 2.2 with a simple choice of a = g, b = 1 - g, with fixed $g = \frac{1}{2}1$, as used in Lan et al. (2025). We avoid learned or dynamic mixing terms to ensure that LA and SWA are used equally in the hybrid attention outputs. The LA component only operates on tokens outside the sliding-window, which has size 64.

Models and Datasets In our experiments, we convert and evaluate up to three popular, pre-trained Transformers: Mistral-7B-Instruct-v0.1, Llama3-8B-Instruct, and Llama3.1-8B-Instruct¹. Attention transfer and fine-tuning are carried out on truncated samples of 1024 tokens from the FineWeb-Edu dataset (Penedo et al., 2024). For evaluation, we report performance on popular LM-Eval tasks, all zero-shot: PIQA, ARC-Easy, ARC-Challenge (acc norm), HellaSwag (acc norm), WinoGrande, and MMLU. In all results, any average performance shown is calculated across all six tasks.

Training & LoRA Parameters We follow Zhang et al. (2024a)'s exact fine-tuning settings, applying LoRA to the query, key, value, and output projection matrices (W_q , W_k , W_v , W_o), with rank r=8, $\alpha=16$. We use an AdamW (Loshchilov & Hutter, 2017) optimiser with learning rate 1e-4 (1e-2 during attention transfer), and a reduce-on-plateau scheduler. We train for 1 epoch for attention transfer, and 1-5 epochs during fine-tuning, using an effective (accumulated) batch size of 64 and approximately 25M tokens per epoch.

3.2 ABLATION STUDY OF LOLCATS HYBRID ATTENTION MODULES AT INFERENCE TIME

In this section, we explore the contribution of each attention module to performance for models trained using LoLCATs' proposed hybrid conversion methodology (Zhang et al., 2024a). We run four ablations at inference time: (i) SWA-only, disabling the LA module by forcing its outputs to zeros; (ii) LA-only where we disable SWA and retain only the LA component; (iii) attention sinks (Xiao et al., 2023) only, suppressing both SWA and LA and only passing the first 8 values through softmax attention; and (iv) no attention where we return an all-zeros attention output, removing any contribution from the attention mechanism. Additionally, we run the same ablations using the provided LoLCATs-trained checkpoints for Llama-3.1-8B using the LoLCATs codebase. Results are shown in Table 1.

Together, the ablations show that almost all the performance attribution sits within the SWA component. Using SWA-only in the resulting models either matches or slightly improves average accuracy for both Mistral-7B and the LoLCATs-trained Llama-3.1-8B 2 . By contrast, LA-only, attention sinks only, and no attention all collapse to roughly the same low performance. The results expose a clear SWA-LA imbalance in current hybrid attention conversion methods that leads to LA contributing little to downstream accuracy and even being detrimental. Similar findings are reported in Lan et al. (2025) (see their Table 6) where they find that SWA-only and GLA+SWA lead to very similar performance while GLA-only leads to a dramatic decrease in performance.

¹Note that we remove the "-Instruct" from the model names in our results for brevity but still refer to these models rather than the base models.

 $^{^2}$ https://huggingface.co/hazyresearch/lolcats-Llama-3.1-8b-distill https://huggingface.co/hazyresearch/lolcats-Llama-3.1-8b-ft-lora

Active Attn Modules	PIQA	ARC-E	ARC-C	HellaSwag	WG	MMLU	AVG	Rec. Perf
Mistral-7B	79.27	80.01	52.22	74.60	69.93	53.51	68.26	100.00
SWA + Linear	78.40	79.50	49.91	71.28	68.35	45.44	65.48	95.93
SWA only	78.56	79.46	49.83	70.57	68.35	46.56	65.56	96.04
Linear only	53.65	29.71	24.57	27.01	50.43	23.28	34.78	50.95
Attn sinks only	54.41	30.05	24.66	28.80	49.64	23.67	35.21	51.58
No Attention	53.92	25.63	24.66	25.99	50.67	25.51	34.40	50.39
Llama3-8B	78.13	81.69	56.66	75.94	71.67	63.85	71.32	100.00
SWA + Linear	78.35	80.47	54.35	71.62	71.67	48.84	67.55	94.71
SWA only	78.02	80.01	54.10	65.08	71.98	47.30	66.08	92.26
Linear only	52.07	25.97	25.77	26.31	48.22	22.95	33.55	47.04
Attn sinks only	56.31	29.00	22.61	28.83	49.41	22.95	34.85	48.86
No Attention	55.17	26.73	22.87	26.13	51.14	22.95	34.17	47.91
Llama3.1-8B (LoLCATs ckpt)	80.14	81.82	55.20	79.14	73.72	68.05	73.01	100.00
SWA + Linear	81.18	82.37	54.78	79.16	70.09	58.89	71.08	97
SWA only	81.56	82.37	55.29	79.76	74.11	55.63	71.45	98
Linear only	51.52	25.00	25.51	26.37	52.25	23.09	33.96	47
Attn sinks only	-	-	-	-	-	-	-	-
No Attention	54.62	26.68	24.40	25.88	48.86	23.12	33.93	46

Table 1: Measuring the effect of attention components on benchmark accuracy for models trained with LoLCATs hybrid conversion.

3.3 REVERTING BACK TO LINEAR ATTENTION-ONLY

In this section, we re-examine the HedgeHog (Zhang et al., 2024b) conversion method with LoL-CATs' (Zhang et al., 2024a) fine-tuning settings. HedgeHog has been shown to work with full-parameter fine-tuning and may be considered the SOTA conversion method for LA-only. Hence, we simply seek to determine whether a more conservative, LoRA-based fine-tune is also able to make use of LA-only in this setting, or whether this difference may be responsible for the gap in LA-only performance between HedgeHog and LoLCATs. More specifically, we evaluate models converted using HedgeHog's core methods, namely attention transfer on attention weights using soft-label cross-entropy and a square projections matrix $\mathbf{W}_{\phi} \in \mathbb{R}^{h_d \times h_d}$, but change the fine-tuning stage to use LoLCATs' LoRA settings. Our findings are illustrated in Figure 1, while per-task results are provided in Appendix A.4.1.

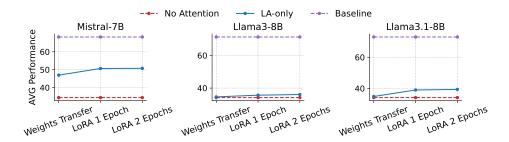


Figure 1: Measuring the performance of HedgeHog conversion at various points during the conversion process. LoRA refers to a LA-only fine-tuning.

We observe a clear contrast with the results in section 3.2, with the linear component in all models beating the no-attention ablation with no or little fine-tuning. However, we observe that the Llama models, and Llama3-8B especially, only see a minor improvement in performance over no use of attention at all.

3.4 Component-Wise Investigation

Our minimal implementation of the HedgeHog conversion method successfully makes use of LA, although with mixed performance recovery. On the other hand, LoLCATs claims to recover base model performance, but appears to do so entirely using SWA. In this section, we seek to identify

275 276

274

277

278

294

304

310 311

317 318

319 320

321

322

323

which parts of these two methods contribute to these two different results. To this end, we ablate various parts of these methods within the attention transfer stage and evaluate them with LA-only. Focusing on this stage of the conversion process allows us to determine exactly which components work best with LA without affecting any weights shared with SWA during fine-tuning. We focus our ablations on the Mistral-7B model as it saw the greatest difference in LA performance when changing the attention transfer and projections ϕ from LoLCATs' to HedgeHog's.

3.4.1 ATTENTION TRANSFER LEARNING

In this ablation, we keep the square HedgeHog feature map used in section 3.3 and simply ablate the Attention Transfer objective. We evaluate three modes: (1) The HedgeHog attention weights transfer using soft-label cross entropy loss between softmax attention weights and the Hedgehog attention weights (both quadratic cost) (2) The LoLCATs hybrid attention outputs transfer using MSE loss between full-context softmax attention outputs (quadratic cost) and the hybrid SWA + LA outputs (linear cost), and (3) the MSE between full softmax attention outputs (quadratic cost) and the full LA outputs (linear cost). We omit any LoRA fine-tuning in these evaluations to ensure we measure the direct impact of attention transfer only. Our findings are outlined in Table 2, while our results for all models are included in Appendix A.4.2.

Transfer Objective	PIQA	ARC-E	ARC-C	HellaSwag	WG	MMLU	AVG	Rec. Perf
Mistral-7B	79.27	80.01	52.22	74.60	69.93	53.51	68.26	100.00
Attention Weights	67.90	61.74	29.69	45.82	52.64	23.82	46.94	68.76
Attention Outputs	69.64	62.46	31.40	46.61	50.28	23.00	47.23	69.20
Hybrid Attention Out.	55.44	32.83	24.06	27.36	49.80	22.98	35.41	51.88
No Attention	53.92	25.63	24.66	25.99	50.67	25.51	34.40	50.39

Table 2: Comparing the performance of LA-only after the attention transfer phase across learning objectives used during this stage.

Our results suggest that the hybrid attention output objective is likely to be responsible for LoLCATs failure to make use of LA, as it barely beats no attention at all. On the other hand, using a full attention output objective seems to beat even the weights transfer objective in this setting. In fact, this objective showed a particular advantage with Llama 3-8B, although both Llama models still lag behind Mistral's recovered performance. We note that we have not observed any successful LA-only conversions of these models in the literature, as opposed to Mistral, suggesting that they may be particularly hard to convert. We leave the analysis as to why this might be to future work.

3.4.2 Φ DIMENSIONALITY

We now seek to measure the impact of the feature map size on conversion success. The original HedgeHog paper makes use of a square linear map $W_{\phi} \in \mathbb{R}^{h_d \times h_d}$ ($\phi : \mathbb{R}^{h_d} \to \mathbb{R}^{2h_d}$, given Eq. 4), with h_d : query-key head dimension, but LoLCATs reduce this to a rectangular linear map $W_{\phi} \in \mathbb{R}^{h_d \times \frac{h_d}{2}}$, for a significantly smaller $\phi : \mathbb{R}^{h_d} \to \mathbb{R}^{h_d}$. To this end, we repeat the experiments presented above in Table 2, with $W_{\phi} \in \mathbb{R}^{h_d \times \frac{h_d}{2}}$. Our findings are outlined in Table 3.

Transfer Objective	PIQA	ARC-E	ARC-C	HellaSwag	WG	MMLU	AVG	Rec. Perf
Mistral-7B	79.27	80.01	52.22	74.60	69.93	53.51	68.26	100.00
Attention Weights	64.85	54.84	26.11	39.48	51.07	23.17	43.25	63.37
Attention Outputs	63.71	53.83	27.30	38.50	49.33	22.90	42.60	62.40
Hybrid Attention Out.	53.97	30.81	22.44	27.42	49.88	22.99	34.59	50.67
No Attention	53.92	25.63	24.66	25.99	50.67	25.51	34.40	50.39

Table 3: Repeating ablations described in Table 2, but with half the W_{ϕ} size.

Our results broadly follow that of section 3.4.1, with performance decreasing across the board with this reduction of output features. More specifically, we see the hybrid attention objective still improves performance only marginally, while the attention weights and outputs objectives achieve a significant and similar improvement on no attention. We note that the weights transfer objective appears to have the upper hand with this smaller W_{ϕ} .

3.4.3 Φ ACTIVATION FUNCTION

Finally, HedgeHog implements an exponential activation function ($\sigma = \exp(\cdot)$). On the other hand, LoLCATs use the softmax, which is the layer-normalised equivalent ($\sigma = \operatorname{softmax}(\cdot)$). Here, we ablate the two, and compare the resulting performance along with T2R and SUPRA's ReLU activation, as well as no activation function at all. Our findings are outlined in Table 4.

Φ Activation Fn	PIQA	ARC-E	ARC-C	HellaSwag	WG	MMLU	AVG	Rec. Perf
Mistral-7B	79.27	80.01	52.22	74.60	69.93	53.51	68.26	100.00
Softmax	67.90	61.74	29.69	45.82	52.64	23.82	46.94	68.76
Exponential	66.32	61.45	29.69	44.67	50.99	23.52	46.11	67.55
ReLU	56.53	34.81	22.95	27.95	50.91	23.29	36.07	52.85
1+ ELU	54.68	29.12	23.81	25.98	51.46	25.64	35.12	51.45
None	51.74	25.51	29.61	25.97	51.07	24.95	34.81	51.00
No Attention	53.92	25.63	24.66	25.99	50.67	25.51	34.40	50.39

Table 4: Ablating the activation function used in HedgeHog projections (ϕ). All results show checkpoints for a single epoch of weights transfer with no fine-tuning, evaluated as LA-only.

Our results show the the exponential-based activation functions (exponential and softmax) far outperform the alternatives. In this regard, LoLCATs and HedgeHog appear to achieve comparable performance, with a slight edge in LoLCATs. Interestingly, 1+ELU does not appear to carry this same performance despite its close similarity to the exponential function by itself.

4 METHODS & RESULTS

Section 3 demonstrates some key issues in hybrid conversion methods, and isolates the components responsible for such shortcomings, namely the hybrid attention transfer objective and the dimensionality of ϕ . In this section, we build on top of these findings and propose, evaluate, and compare three different methods to fine-tune a model for hybrid attention while avoiding the dominance of SWA and decay of LA. Note that all experiments conducted in this section use the same experimental setup as in Subsection 3.1, unless explicitly stated otherwise.

4.1 Inference-Time Hybrid

Seeing as fine-tuning for hybrid attention appears to encourage the model to focus on the more expressive SWA, we first propose a "zero-shot" hybrid, introducing SWA in models which have either seen no fine-tuning or fine-tuning with LA-only. We implement two modes of hybrid attention: one where LA only sees tokens outside of SWA's sliding window, and one where LA sees all past tokens therefore overlapping with SWA's context ("Overlap" in our results). Our findings are illustrated in Figure 2, while per-task results are provided in Appendix A.4.3.

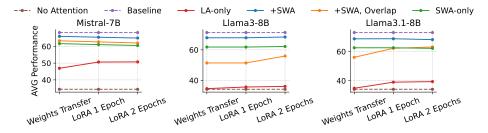


Figure 2: Adding SWA at inference time at various points of the conversion process.

We find that, while LA-only benefits from fine-tuning, the latter appears to slightly degrade the performance of SWA. As such, the best inference-time hybrid performance is achieved with no fine-tuning. Furthermore, it would appear as though overlapping LA's context with that of SWA also degrades performance. We theorise that this is most likely due to our fixed mixing term g=0.5 and

may be improved by a data-dependent mixing term which may dynamically choose how to weight the contributions of each component in the case of overlap, as employed by Irie et al. (2025).

4.2 HEDGECATS: HEDGEHOG TRANSFER + LOLCATS FINE-TUNING

Our second method, HedgeCATs, blends HedgeHog's attention weight transfer with LoLCATs' LoRA fine-tuning of hybrid attention. The first training stage performs HedgeHog-style transfer to learn a feature map ϕ that mimics full softmax attention, training with LA-only so the LA path stands up on its own.

Stage 2 applies LoRA fine-tuning while reintroducing SWA, aiming to recover base-model accuracy without letting the hybrid collapse back onto SWA. In practice, early stopping is key: Figure 3 shows that for Llama-3 8B, as fine-tuning proceeds the hybrid branch trends down, whereas SWA-only improves; for Mistral-7B, both hybrid and SWA-only degrade with more LoRA steps. These behaviours suggest a very short LoRA schedule gives the best trade-off between accuracy, attributional validity of hybrid attention, and training budget. Per-task performance can be seen in Appendix A.4.4.

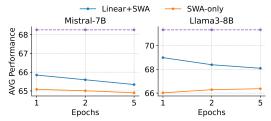


Figure 3: The average performance of HedgeCATstrained models for difference amounts of LoRA fine-tuning.

4.3 SSD: SCHEDULED SLIDING-WINDOW DROPOUT

Finally, to further guard against SWA dominance during training, we introduce Scheduled Sliding-window Dropout (SSD). SSD alters the SWA component in hybrid attention during LoRA fine-tuning according to a dropout and sliding-window size schedule. Our results across SSD fine-tuning regimes are summarised in Figure 4. We vary dropout and sliding-window schedules (either fixed or epochvarying) and evaluate at multiple fine-tuning epochs to characterise performance as a function of training time. Scheduled parameters are applied per epoch: at epoch k we use the k-th value in the schedule, and once the schedule is exhausted the final value is held fixed. This analysis illustrates how the SWA component of the hybrid attention mechanism evolves during fine-tuning. Figure 4a uses a dropout schedule $(0.9 \rightarrow 0.75 \rightarrow 0.5)$, i.e., the SWA branch is dropped 90% of the time in epoch 1, with a fixed sliding-window size of 32; Figure 4b fixes dropout at 0.5 and schedules the sliding-window size $(4 \rightarrow 8 \rightarrow 16 \rightarrow 32 \rightarrow 64)$; Figure 4c fixes both dropout and sliding-window size at 0.5 and 16, respectively. Per-task results for each experimented SSD setting are provided in Appendix A.4.5.

The results for SSD-trained models show consistent trends across Mistral-7B and Llama-3-8B. For Figure 4a, LA+SWA improves steadily with fine-tuning. This indicates that heavy early SWA dropout successfully pushes learning through the linear path before relaxing to 0.5. In Figure 4b, performance starts lower compared to the previous experiment (penalised by the short initial window size) and remains flat or slightly down for Mistral-7B, with only a mild late recovery for Llama3-8B; SWA-only also drifts down or flattens, suggesting limited benefit from widening the window alone in this setting. The fixed dropout and sliding-window model, Figure 4c, yields similar results to Figure 4a for Llama and slightly worse performance for Mistral. Across experiment settings, SWA-only sits well below Linear+SWA, confirming that our fine-tuning schedule leads to component-balanced hybrid attention.

5 DISCUSSION & CONCLUSION

Our analysis reveals a critical failure mode in hybrid attention conversion methods: models trained with hybrid objectives often bypass LA entirely, relying exclusively on SWA. In contrast, attention weights or full attention outputs transfer objectives successfully enable LA, though with varying performance recovery. We also confirm previous findings within our corrected framework, namely that exponential-family activations outperform alternatives, and larger feature map projections improve performance. These findings likely extend to other hybrid methods, like Liger, which report similar component imbalances. To address this issue, we proposed three solutions for converting a model to use hybrid attention: (1) Inference-time hybrid addition, which preserves LA without additional

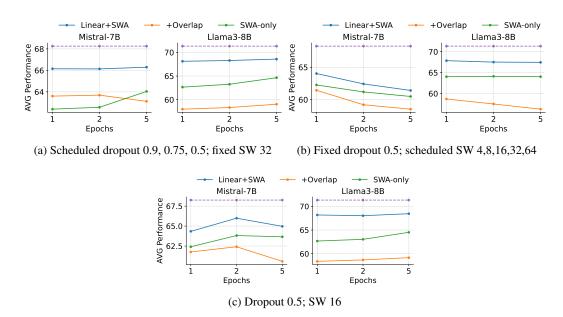


Figure 4: Performance comparison of dropout and sliding-window size schedules for a different number of fine-tuning epochs.

training; (2) HedgeCATs, combining successful weights transfer with careful early-stopping during hybrid fine-tuning; and (3) Scheduled Sliding-window Dropout (SSD), providing robust training that maintains component balance. Each offers different trade-offs between simplicity, individual component performance, and training cost, while all recovering over 95% of base model performance. Importantly, the SWA-driven LoLCATs-trained hybrids (Table 1) consistently underperform compared to all our proposed methods. This indicates that enforcing genuine LA utilisation improves overall accuracy as well as attributional validity.

Limitations and Future Work In this study, performance is limited by our simplified implementation of some of the components. For example, it should be noted that replacing our fixed mixing term g with learned and dynamic mixing mechanisms, as seen in other methods (see section 2) is likely to increase performance. Although we motivate such choices with this study's focus on clear performance attribution between LA and SWA, while minimising the model's ability to discard LA, future work should extend our resulting methods to maximise performance through further investigation of components such as the mixing term, normalisation methods, training datasets, LoRA settings, base models (instruction-tuned or not), etc.

LA is also believed to suffer from the same capacity issues observed in associative memory networks (Schlag et al., 2021). To this end, multiple gating mechanisms have been added to LA in order to manage information retention and retrieval accuracy across longer sequences (Schlag et al., 2021; Sun et al., 2023; Yang et al., 2023; 2024; Liu et al., 2024). Future work should examine how conversion methods may be applied to convert models to use such mechanisms. Additionally, LA conversion methods have the potential to improve performance of grouped KV retrieval methods (Xiao et al., 2024; Fountas et al., 2025), which reduce information dilution in long-context settings by only attending to subsets of past tokens. Such methods are limited by the memory requirements for storing long token histories. An interesting extension of this work would use LA to compress individual memory blocks, reducing memory requirements while limiting the capacity issues arising from LA being applied to entire long-context sequences.

Conclusion In conclusion, while hybrid conversions promise efficiency with minimal performance loss, without careful design they fail to genuinely adopt LA. By identifying this failure mode and proposing solutions that maintain component balance, we restore attributional validity, ensuring claimed architectural components actually contribute to model behaviour, which is essential for advancing efficient Transformer architectures.

ETHICS STATEMENT

This work focuses on foundational methods for converting large language models to use LA mechanisms. All experiments were conducted using publicly-available pre-trained models, and widely used training and benchmark datasets. No personally identifiable or sensitive data was used. The primary contribution is methodological, and thus does not introduce new societal risks beyond those already known for these models, such as potential biases or misuse. We encourage responsible and transparent use of these methods.

REPRODUCIBILITY STATEMENT

We include experimental details across sections 2-4, as well as Appendix A.2. Any details which aren't explicitly mentioned in this work are aligned with related works and clearly mentioned as such. Our experiments are reproducible within the LoLCATs public codebase with relatively minimal changes, and we are working on releasing our own public version of our codebase including all mentioned methods and ablations.

USE OF LARGE LANGUAGE MODELS

Our use of LLMs in the writing of this paper is limited to sparse and light improvements in wording within the main text.

REFERENCES

- Kunal Banerjee, Vishak C., Rishi Raj Gupta, Kartik Vyas, Anushree H., and Biswajit Mishra. Exploring alternatives to softmax function. *ArXiv*, abs/2011.11538, 2020. URL https://api.semanticscholar.org/CorpusID:227127574.
- Ali Behrouz, Peilin Zhong, and Vahab Mirrokni. Titans: Learning to memorize at test time. *arXiv* preprint arXiv:2501.00663, 2024.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- Tianlong Chen, Zhenyu Zhang, Ajay Kumar Jaiswal, Shiwei Liu, and Zhangyang Wang. Sparse moe as the new dropout: Scaling dense and self-slimmable transformers. In *ICLR*, 2023. URL https://openreview.net/pdf?id=w1hwFUb_81.
- Krzysztof Marcin Choromanski, Valerii Likhosherstov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, David Benjamin Belanger, Lucy J Colwell, and Adrian Weller. Rethinking attention with performers. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=Ua6zuk0WRH.
- Jean-Baptiste Cordonnier, Andreas Loukas, and Martin Jaggi. Multi-head attention: Collaborate instead of concatenate, 2021. URL https://openreview.net/forum?id=bK-rJMKrOsm.
- Xin Dong, Yonggan Fu, Shizhe Diao, Wonmin Byeon, Zijia Chen, Ameya Sunil Mahabaleshwarkar, Shih-Yang Liu, Matthijs Van Keirsbilck, Min-Hung Chen, Yoshi Suhara, et al. Hymba: A hybridhead architecture for small language models. *arXiv preprint arXiv:2411.13676*, 2024.
- Feyza Duman Keles, Pruthuvi Mahesakya Wijewardena, and Chinmay Hegde. On the computational complexity of self-attention. In Shipra Agrawal and Francesco Orabona (eds.), *Proceedings of The 34th International Conference on Algorithmic Learning Theory*, volume 201 of *Proceedings of Machine Learning Research*, pp. 597–619. PMLR, 20 Feb–23 Feb 2023. URL https://proceedings.mlr.press/v201/duman-keles23a.html.
- Angela Fan, Edouard Grave, and Armand Joulin. Reducing transformer depth on demand with structured dropout. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=SylO2yStDr.

- Zafeirios Fountas, Martin Benfeghoul, Adnan Oomerjee, Fenia Christopoulou, Gerasimos Lampouras, Haitham Bou Ammar, and Jun Wang. Human-inspired episodic memory for infinite context LLMs. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=BI2int5SAC.
 - Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv* preprint arXiv:2312.00752, 2023.
 - Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.
 - Kazuki Irie, Morris Yau, and Samuel J Gershman. Blending complementary memory systems in hybrid quadratic-linear transformers. *arXiv preprint arXiv:2506.00744*, 2025.
 - Jungo Kasai, Hao Peng, Yizhe Zhang, Dani Yogatama, Gabriel Ilharco, Nikolaos Pappas, Yi Mao, Weizhu Chen, and Noah A. Smith. Finetuning pretrained transformers into rnns. ArXiv, abs/2103.13076, 2021. URL https://api.semanticscholar.org/CorpusID: 232335426.
 - Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pp. 5156–5165. PMLR, 2020.
 - Disen Lan, Weigao Sun, Jiaxi Hu, Jusen Du, and Yu Cheng. Liger: Linearizing large language models to gated recurrent structures. *arXiv* preprint arXiv:2503.01496, 2025.
 - Bonan Li, Yinhan Hu, Xuecheng Nie, Congying Han, Xiangjian Jiang, Tiande Guo, and Luoqi Liu. Dropkey for vision transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22700–22709, 2023.
 - Opher Lieber, Barak Lenz, Hofit Bata, Gal Cohen, Jhonathan Osin, Itay Dalmedigos, Erez Safahi, Shaked Meirom, Yonatan Belinkov, Shai Shalev-Shwartz, et al. Jamba: A hybrid transformer-mamba language model. *arXiv preprint arXiv:2403.19887*, 2024.
 - Bo Liu, Rui Wang, Lemeng Wu, Yihao Feng, Peter Stone, and Qiang Liu. Longhorn: State space models are amortized online learners. *arXiv preprint arXiv:2407.14207*, 2024.
 - Liyuan Liu, Xiaodong Liu, Jianfeng Gao, Weizhu Chen, and Jiawei Han. Understanding the difficulty of training transformers. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 5747–5763, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.463. URL https://aclanthology.org/2020.emnlp-main.463/.
 - Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint* arXiv:1711.05101, 2017.
 - Huanru Henry Mao. Fine-tuning pre-trained transformers into decaying fast weights. In *Conference on Empirical Methods in Natural Language Processing*, 2022. URL https://api.semanticscholar.org/CorpusID:252781097.
 - Jean Mercat, Igor Vasiljevic, Sedrick Scott Keh, Kushal Arora, Achal Dave, Adrien Gaidon, and Thomas Kollar. Linearizing large language models. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=soGxskHGox.
 - Tsendsuren Munkhdalai, Manaal Faruqui, and Siddharth Gopal. Leave no context behind: Efficient infinite context transformers with infini-attention. *arXiv preprint arXiv:2404.07143*, 101, 2024.
 - Guilherme Penedo, Hynek Kydlíček, Anton Lozhkov, Margaret Mitchell, Colin A Raffel, Leandro Von Werra, Thomas Wolf, 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.
 - Bo Peng, Ruichong Zhang, Daniel Goldstein, Eric Alcaide, Xingjian Du, Haowen Hou, Jiaju Lin, Jiaxing Liu, Janna Lu, William Merrill, et al. Rwkv-7" goose" with expressive dynamic state evolution. *arXiv preprint arXiv:2503.14456*, 2025.

- Hao Peng, Nikolaos Pappas, Dani Yogatama, Roy Schwartz, Noah Smith, and Lingpeng Kong.
 Random feature attention. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=QtTKTdVrFBB.
 - Zhen Qin, Weixuan Sun, Hui Deng, Dongxu Li, Yunshen Wei, Baohong Lv, Junjie Yan, Lingpeng Kong, and Yiran Zhong. cosformer: Rethinking softmax in attention. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=B18CQrx2Up4.
 - Liliang Ren, Yang Liu, Yadong Lu, Yelong Shen, Chen Liang, and Weizhu Chen. Samba: Simple hybrid state space models for efficient unlimited context language modeling. *arXiv preprint arXiv:2406.07522*, 2024.
 - Imanol Schlag, Kazuki Irie, and Jürgen Schmidhuber. Linear transformers are secretly fast weight programmers. In *International conference on machine learning*, pp. 9355–9366. PMLR, 2021.
 - Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
 - Yutao Sun, Li Dong, Shaohan Huang, Shuming Ma, Yuqing Xia, Jilong Xue, Jianyong Wang, and Furu Wei. Retentive network: A successor to transformer for large language models. *arXiv* preprint *arXiv*:2307.08621, 2023.
 - Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
 - Chaojun Xiao, Pengle Zhang, Xu Han, Guangxuan Xiao, Yankai Lin, Zhengyan Zhang, Zhiyuan Liu, and Maosong Sun. Infllm: Training-free long-context extrapolation for llms with an efficient context memory. *Advances in Neural Information Processing Systems*, 37:119638–119661, 2024.
 - Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*, 2023.
 - Songlin Yang, Bailin Wang, Yikang Shen, Rameswar Panda, and Yoon Kim. Gated linear attention transformers with hardware-efficient training. *arXiv preprint arXiv:2312.06635*, 2023.
 - Songlin Yang, Jan Kautz, and Ali Hatamizadeh. Gated delta networks: Improving mamba2 with delta rule. *arXiv preprint arXiv:2412.06464*, 2024.
 - Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. Big bird: Transformers for longer sequences. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), Advances in Neural Information Processing Systems, volume 33, pp. 17283–17297. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/c8512d142a2d849725f31a9a7a361ab9-Paper.pdf.
 - Michael Zhang, Simran Arora, Rahul Chalamala, Alan Wu, Benjamin Spector, Aaryan Singhal, Krithik Ramesh, and Christopher Ré. Lolcats: On low-rank linearizing of large language models. *arXiv preprint arXiv:2410.10254*, 2024a.
 - Michael Zhang, Kush Bhatia, Hermann Kumbong, and Christopher Ré. The hedgehog & the porcupine: Expressive linear attentions with softmax mimicry. *arXiv preprint arXiv:2402.04347*, 2024b.
 - Wangchunshu Zhou, Tao Ge, Furu Wei, Ming Zhou, and Ke Xu. Scheduled DropHead: A regularization method for transformer models. In Trevor Cohn, Yulan He, and Yang Liu (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 1971–1980, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.178. URL https://aclanthology.org/2020.findings-emnlp.178/.

Chen Zhu, Wei Ping, Chaowei Xiao, Mohammad Shoeybi, Tom Goldstein, Anima Anandkumar, and Bryan Catanzaro. Long-short transformer: Efficient transformers for language and vision. *Advances in neural information processing systems*, 34:17723–17736, 2021.

A APPENDIX

A.1 NOTATION

Vector operations Take vectors x, y

• $x^{\top}y$: vector inner-product • xy^{\top} : vector outer-product

Operators

- ⊙: Hadamard product
- ⊕: Concatenation

A.2 IMPLEMENTATION DETAILS

Most of the experimental details are included across sections 2- 4. Any details which aren't explicitly mentioned either in theses sections or this one are aligned with related works and clearly mentioned as such.

A.2.1 WEIGHTS INITIALISATION

 Φ Linear projections, as in $\phi(\cdot)$, were initialised as the identity, following Zhang et al. (2024b), with added Gaussian-sampled noise ($\mu=0$ and $\sigma=0.1$) using a seed-based generator to ensure the same initialisation across all runs.

A.3 EXTENDED RELATED WORKS

Method	Attn Type	Feature map $\phi(x)$	Transfer objective
T2R/SUPRA	Fully linear	$\operatorname{ReLU}(W_{\phi}^{\top}x + b)$	Uptrain with standard causal LM cross- entropy (no attention-distillation objec- tive).
HedgeHog	Fully linear	$\exp(W_{\phi}^{\top}x+b)^*$	Minimises cross-entropy / KL between softmax attention weights and linear weights.
LoLCATs	Hybrid (LA + SWA)	$Softmax(W_{\phi}^{\top}x+b)^*$	Minimises MSE between softmax attention outputs and linear (hybrid) attention outputs.
Liger	Hybrid (GLA + SWA)	Softmax(x)	Causal LM cross-entropy with LoRA fine-tuning (no attention-distillation objective).

Table 5: Side-by-side comparison of linearising conversions for pre-trained Transformers. *Methods for which the negative mapping is concatenated to the final output of ϕ (see Eq. 4).

A.4 FURTHER RESULTS

A.4.1 LINEAR ATTENTION ONLY

Active Attn Modules	PIQA	ARC-E	ARC-C	HellaSwag	WG	MMLU	AVG	Rec. Perf
Mistral-7B	79.27	80.01	52.22	74.60	69.93	53.51	68.26	100.00
No Attention	53.92	25.63	24.66	25.99	50.67	25.51	34.40	50.39
LA Weights Transfer	67.90	61.74	29.69	45.82	52.64	23.82	46.94	68.76
+LoRA 1 Epoch	70.13	66.20	34.90	56.39	51.38	24.95	50.66	74.22
+LoRA 2 Epochs	70.29	66.29	34.39	56.41	52.17	24.97	50.75	74.36
Llama3-8B	78.13	81.69	56.66	75.94	71.67	63.85	71.32	100.00
No Attention	55.17	26.73	22.87	26.13	51.14	22.95	34.17	47.90
LA Weights Transfer	54.24	26.94	26.11	26.42	50.43	23.42	34.59	48.50
+LoRA 1 Epoch	54.90	29.46	25.68	32.06	49.01	22.79	35.65	49.98
+LoRA 2 Epochs	54.73	30.26	25.43	33.28	49.64	22.94	36.05	50.54
Llama3.1-8B	80.14	81.82	55.20	79.14	73.72	68.05	73.01	100.00
No Attention	53.59	26.89	25.94	26.19	49.09	22.95	34.11	46.72
LA Weights Transfer	54.41	28.16	25.34	27.50	48.46	25.04	34.82	47.69
+LoRA 1 Epoch	58.71	38.51	24.23	35.38	53.51	23.53	38.98	53.39
+LoRA 2 Epochs	58.87	39.69	23.89	37.41	52.49	23.78	39.36	53.90

Table 6: Measuring the performance of HedgeHog conversion at various points during the conversion process. LoRA refers to a LA-only finetuning.

A.4.2 ATTENTION TRANSFER LEARNING OBJECTIVE

Transfer Objective	PIQA	ARC-E	ARC-C	HellaSwag	WG	MMLU	AVG	Rec. Perf
Mistral-7B	79.27	80.01	52.22	74.60	69.93	53.51	68.26	100.00
Attention Weights	67.90	61.74	29.69	45.82	52.64	23.82	46.94	68.76
Attention Outputs	69.64	62.46	31.40	46.61	50.28	23.00	47.23	69.20
Hybrid Attention Out.	55.44	32.83	24.06	27.36	49.80	22.98	35.41	51.88
No Attention	53.92	25.63	24.66	25.99	50.67	25.51	34.40	50.39
Llama3-8B	78.13	81.69	56.66	75.94	71.67	63.85	71.32	100.00
Attention Weights	54.24	26.94	26.11	26.42	50.43	23.42	34.59	48.50
Attention Outputs	61.04	41.04	24.06	30.02	49.25	23.88	38.22	53.58
Hybrid Attention Out.	53.10	26.22	24.74	26.23	49.41	23.43	33.86	47.47
No Attention	55.17	26.73	22.87	26.13	51.14	22.95	34.17	47.90
Llama3.1-8B	80.14	81.82	55.20	79.14	73.72	68.05	73.01	100.00
Attention Weights	54.41	28.16	25.34	27.50	48.46	25.04	34.82	47.69
Attention Outputs	54.90	28.49	25.60	28.60	48.78	26.30	35.45	48.55
Hybrid Attention Out.	54.57	25.46	26.02	26.04	47.75	23.37	33.87	46.39
No Attention	53.59	26.89	25.94	26.19	49.09	22.95	34.11	46.72

Table 7: Comparing the performance of linear attention only after the attention transfer phase across learning objectives used during this stage.

A.4.3 Adding SWA at Inference Time

Active Attn Modules	PIQA	ARC-E	ARC-C	HellaSwag	WG	MMLU	AVG	Rec. Perf
Mistral-7B	79.27	80.01	52.22	74.60	69.93	53.51	68.26	100.00
LA Weights Transfer	67.90	61.74	29.69	45.82	52.64	23.82	46.94	68.76
+SWA	79.11	79.76	51.62	70.57	69.93	45.13	66.02	96.72
+SWA, Overlap	77.48	78.79	50.77	67.13	61.96	43.88	63.34	92.79
SWA-only	77.97	79.92	51.02	51.19	69.93	39.92	61.66	90.33
+LoRA 1 Epoch	70.13	66.20	34.90	56.39	51.38	24.95	50.66	74.22
+SWA	78.73	78.83	51.96	70.44	70.24	42.66	65.48	95.93
+SWA, Overlap	77.42	78.20	49.06	67.38	63.30	40.84	62.70	91.86
SWA-only	77.91	78.75	51.54	48.63	70.24	39.41	61.08	89.49
+LoRA 2 Epochs	70.29	66.29	34.39	56.41	52.17	24.97	50.75	74.36
+SWA	78.62	78.70	51.11	69.75	69.61	42.14	64.99	95.21
+SWA, Overlap	77.26	77.78	48.04	66.22	62.98	39.99	62.05	90.90
SWA-only	77.80	78.49	50.77	47.69	69.61	38.91	60.55	88.70
Llama3-8B	78.13	81.69	56.66	75.94	71.67	63.85	71.32	100.00
LA Weights Transfer	54.24	26.94	26.11	26.42	50.43	23.42	34.59	48.50
+SWA	77.80	81.06	55.97	71.71	71.67	49.22	67.91	95.21
+SWA, Overlap	67.03	56.06	36.86	49.67	58.56	40.26	51.41	72.08
SWA-only	76.28	80.09	55.12	44.48	71.67	43.31	61.83	86.68
+LoRA 1 Epoch	54.90	29.46	25.68	32.06	49.01	22.79	35.65	49.98
+SWA	77.80	81.06	55.97	71.71	71.67	49.22	67.91	95.90
+SWA, Overlap	67.03	56.06	36.86	49.67	58.56	40.26	51.41	78.42
SWA-only	76.28	80.09	55.12	44.48	71.67	43.31	61.83	87.20
+LoRA 2 Epochs	54.73	30.26	25.43	33.28	49.64	22.94	36.05	50.54
+SWA	78.35	80.22	55.55	73.15	73.24	49.87	68.40	95.27
+SWA, Overlap	70.62	61.20	41.13	58.61	60.30	43.71	55.93	78.10
SWA-only	76.99	79.46	54.01	45.71	73.24	43.77	62.20	86.74
Llama3.1-8B	80.14	81.82	55.20	79.14	73.72	68.05	73.01	100.00
LA Weights Transfer	54.41	28.16	25.34	27.50	48.46	25.04	34.82	47.69
+SWA	79.87	80.98	54.78	72.81	73.72	50.11	68.71	94.11
+SWA, Overlap	71.38	64.35	40.44	51.66	62.51	45.90	56.04	76.75
SWA-only	78.51	80.35	53.24	46.08	73.72	43.51	62.57	85.70
+LoRA 1 Epoch	58.71	38.51	24.23	35.38	53.51	23.53	38.98	53.39
+SWA	80.09	81.44	53.58	72.77	74.03	50.76	68.78	94.20
+SWA, Overlap	76.88	75.93	48.89	61.16	62.59	47.79	62.21	85.20
SWA-only	78.94	80.64	51.96	46.24	74.03	43.64	62.58	85.71
+LoRA 2 Epochs	58.87	39.69	23.89	37.41	52.49	23.78	39.36	53.90
+SWA	79.65	80.72	53.58	72.07	73.40	49.56	68.16	93.36
+SWA, Overlap	77.48	77.82	50.51	61.46	63.38	47.98	63.11	86.43
SWA-only	78.45	79.92	52.39	45.27	73.40	43.58	62.17	85.15

Table 8: Adding SWA at inference time at various points of the conversion process.

A.4.4 HEDGECATS

Active Attn Module	PIQA	ARC-E	ARC-C	HellaSwag	WG	MMLU	AVG	Rec. Perf
Mistral-7B	79.27	80.01	52.22	74.60	69.93	53.51	68.26	100.00
1 epoch								
Linear+SWA	78.51	79.46	51.45	71.14	69.38	45.14	65.85	96.47
SWA-only	78.24	79.34	50.68	65.92	69.38	46.96	65.09	95.36
2 epochs								
Linear+SWA	78.24	79.59	51.02	70.73	68.75	45.24	65.60	96.10
SWA-only	78.07	79.71	50.17	66.30	68.75	47.07	65.01	95.25
5 epochs								
Linear+SWA	78.56	79.42	50.26	70.55	67.64	45.62	65.34	95.73
SWA-only	78.51	79.67	49.66	66.51	67.64	47.41	64.90	95.08
Llama3-8B	78.13	81.69	56.66	75.94	71.67	63.85	71.32	100.00
1 epoch								
Linear+SWA	79.38	81.78	55.97	72.28	73.32	51.18	68.99	96.72
SWA-only	78.89	81.61	55.63	59.67	73.32	47.02	66.02	92.57
2 epochs								
Linear+SWA	78.62	80.89	53.92	71.54	73.48	51.87	68.39	95.88
SWA-only	78.40	81.14	54.27	62.68	73.48	47.76	66.29	92.94
5 epochs								
Linear+SWA	78.73	80.05	52.99	71.49	72.38	52.90	68.09	95.47
SWA-only	78.78	80.39	53.33	65.57	72.38	47.76	66.37	93.05

Table 9: HedgeCATs performance for different activate attention modules at inference time and at different stages of LoRA finetuning.

A.4.5 SSD: TESTING SCHEDULED HYBRID CONVERSION

Active Attn Modules	PIQA	ARC-E	ARC-C	HellaSwag	WG	MMLU	AVG	Rec. Perf
Mistral-7B	79.27	80.01	52.22	74.60	69.93	53.51	68.26	100.00
1 epoch								
Linear+SWA	78.94	79.55	52.47	71.55	69.61	44.79	66.15	96.92
+ Overlap	77.86	78.58	50.34	68.46	63.61	42.79	63.61	93.19
SWA-only	78.24	79.59	52.39	53.77	69.61	40.74	62.39	91.40
2 epochs								
Linear+SWA	78.62	79.59	52.90	71.61	68.98	45.13	66.14	96.90
+ Overlap	77.64	78.49	50.43	68.55	64.25	42.80	63.69	93.31
SWA-only	77.80	79.76	52.39	55.45	68.98	41.02	62.57	91.66
5 epochs								
Linear+SWA	78.62	79.55	53.58	71.06	69.93	45.04	66.30	97.13
+ Overlap	77.48	77.90	50.68	68.09	62.35	42.24	63.12	92.48
SWA-only	78.40	79.25	53.24	61.45	69.93	41.97	64.04	93.82
Llama3-8B	78.13	81.69	56.66	75.94	71.67	63.85	71.32	100.00
1 epoch								
Linear+SWA	78.40	80.72	55.72	72.07	71.43	50.38	68.12	95.59
+ Overlap	70.89	63.80	41.81	63.53	59.98	47.98	58.00	81.32
SWA-only	77.26	79.88	54.61	48.04	71.43	44.70	62.65	87.84
2 epochs								
Linear+SWA	78.18	80.85	55.80	72.34	72.22	50.34	68.29	95.74
+ Overlap	71.33	63.64	41.47	65.52	59.98	48.15	58.35	82.29
SWA-only	77.15	80.05	54.61	50.31	72.22	45.28	63.27	88.38
5 epochs								
Linear+SWA	78.13	80.72	56.23	72.84	72.61	50.97	68.58	95.97
+ Overlap	71.55	63.30	41.13	68.13	60.38	49.75	59.04	82.78
SWA-only	77.48	80.43	55.12	56.11	72.61	46.16	64.65	90.65

Table 10: Finetuning using SSD for a dropout schedule of 0.9, 0.75, 0.5 and a fixed sliding window size of 32.

Active Attn Modules	PIQA	ARC-E	ARC-C	HellaSwag	WG	MMLU	AVG	Rec. Perf
Mistral-7B	79.27	80.01	52.22	74.60	69.93	53.51	68.26	100.00
1 epoch								
Linear+SWA	78.24	77.95	50.34	67.30	68.27	43.93	64.34	94.26
+ Overlap	77.37	77.19	47.10	66.63	59.91	42.36	61.76	90.48
SWA-only	77.69	78.28	50.00	58.59	68.27	41.65	62.41	91.44
2 epochs								
Linear+SWA	78.78	79.34	53.33	70.47	69.06	44.95	65.99	96.68
+ Overlap	77.31	77.44	48.46	66.74	61.48	43.04	62.41	91.44
SWA-only	78.29	79.46	52.56	61.13	69.06	42.40	63.82	93.50
5 epochs								
Linear+SWA	78.51	79.12	52.13	68.30	69.06	42.69	64.97	95.18
+ Overlap	76.82	76.52	47.95	62.36	60.54	39.30	60.58	88.76
SWA-only	78.18	79.08	51.62	62.49	69.06	41.52	63.66	93.26
Llama3-8B	78.13	81.69	56.66	75.94	71.67	63.85	71.32	100.00
1 epoch								
Linear+SWA	78.07	80.64	55.38	72.23	72.38	50.38	68.18	95.59
+ Overlap	71.16	65.03	42.15	63.66	60.22	48.08	58.38	81.86
SWA-only	76.88	79.76	54.27	48.21	72.38	44.59	62.68	87.88
2 epochs								
Linear+SWA	78.07	80.77	54.86	72.41	71.59	50.46	68.03	95.38
+ Overlap	71.38	64.56	41.55	65.71	60.77	48.20	58.70	82.29
SWA-only	77.04	80.05	53.75	50.37	71.59	45.41	63.04	88.38
5 epochs								
Linear+SWA	78.35	80.43	55.97	72.89	72.14	50.90	68.45	95.97
+ Overlap	71.38	63.47	42.06	68.12	60.30	49.65	59.16	82.95
SWA-only	77.64	80.18	54.86	56.07	72.14	46.20	64.52	90.45

Table 11: Finetuning using SSD for a fixed dropout of 0.5 and sliding window size of 16.

Active Attn Modules	PIQA	ARC-E	ARC-C	HellaSwag	WG	MMLU	AVG	Rec. Perf
Mistral-7B	79.27	80.01	52.22	74.60	69.93	53.51	68.26	100.00
1 epoch								
Linear+SWA	77.75	77.74	50.17	66.81	68.43	43.28	64.03	93.81
+ Overlap	77.26	76.56	46.67	66.01	59.91	42.24	61.44	90.02
SWA-only	77.42	77.99	49.74	58.31	68.43	41.71	62.27	91.22
2 epochs								
Linear+SWA	76.82	74.07	46.08	66.44	68.43	42.73	62.43	91.46
+ Overlap	75.84	72.31	40.96	65.79	58.88	41.48	59.21	86.75
SWA-only	76.82	74.71	46.50	58.98	68.43	41.70	61.19	89.65
5 epochs								
Linear+SWA	76.33	73.65	46.08	63.96	67.17	41.34	61.42	89.99
+ Overlap	75.63	72.31	42.41	62.89	59.27	38.66	58.53	85.75
SWA-only	76.22	73.91	46.25	58.38	67.17	41.01	60.49	88.62
Llama3-8B	78.13	81.69	56.66	75.94	71.67	63.85	71.32	100.00
1 epoch								
Linear+SWA	77.91	80.39	54.27	71.49	72.06	51.00	67.85	95.13
+ Overlap	71.00	63.64	40.87	66.37	61.40	48.88	58.69	82.29
SWA-only	77.26	80.13	53.75	55.03	72.06	45.93	64.03	89.77
2 epochs								
Linear+SWA	77.48	79.50	53.67	71.63	72.45	50.35	67.51	94.66
+ Overlap	70.08	60.23	39.08	66.63	60.93	48.06	57.50	80.62
SWA-only	77.31	79.34	53.16	56.75	72.45	45.64	64.11	89.88
5 epochs								
Linear+SWA	78.02	79.59	54.10	71.65	71.67	49.67	67.45	94.57
+ Overlap	68.82	58.50	38.14	66.24	58.88	46.99	56.26	78.88
SWA-only	77.37	79.59	52.99	57.16	71.67	45.37	64.03	89.77

Table 12: Finetuning using SSD for a fixed dropout of 0.5 and a sliding window size schedule of 4, 8, 16, 32, 64.