

Garbage Attention in Large Language Models: <BOS> Sink Heads and Sink-aware Pruning

Jaewon Sok¹ Jewon Yeom² Seonghyeon Park³ Jeongjae Park² Taesup Kim^{2,*}

¹Department of Rural Systems Engineering, Seoul National University

²Graduate School of Data Science, Seoul National University

³Department of Aerospace Engineering, Seoul National University

Abstract

Large Language Models (LLMs) are known to contain significant redundancy, yet a systematic explanation for why certain components, particularly in higher layers, are more redundant has remained elusive. In this work, we identify the <BOS> sink phenomenon as a key mechanism driving this layer-wise sensitivity. We show that attention heads with high <BOS> sink scores are strongly associated with functional redundancy: such heads, especially in deeper layers, contribute little to predictive performance and effectively serve as *dumping grounds* for superfluous attention weights. Leveraging this insight, we introduce a simple pruning strategy that removes high-<BOS> sink heads. Experiments on Gemma-3, Llama-3.1, and Qwen3 demonstrate that this approach identifies redundant transformer components more reliably than weight- and activation-based criteria in terms of downstream task retention, remaining close to dense baselines at low-to-moderate pruning ratios. We further find that high-scoring sink heads sustain their focus on <BOS> as context length grows. Overall, our results suggest that structural properties of attention offer a more direct basis for model compression than magnitude-based methods.

1 Introduction

The advent of Transformers (Vaswani et al., 2017) has enabled Large Language Models (LLMs) to achieve unprecedented performance on various tasks (Brown et al., 2020; Zhao et al., 2026). However, the massive parameter spaces of these models contain significant redundancy (Michel et al., 2019; Fan et al., 2025; He et al., 2024; Kaushik et al., 2025), serving as a primary bottleneck for efficient deployment. Consequently, while much effort has been devoted to model compression and pruning (Hoeffer et al., 2021; Zhu et al., 2024; Song et al.,




Figure 1: Attention weights of a representative <BOS> sink head (L26, H0) in Gemma-3-4B. The heatmap reveals a stark attention sink pattern where query tokens disproportionately attend to the <BOS> token during an MMLU task.

2024), the underlying nature and origins of this redundancy remain relatively unexplored.

Previous research on the attention mechanism, the cornerstone of the Transformer architecture, has demonstrated that individual attention heads perform distinct functional roles (Clark et al., 2019b; Voita et al., 2019). Notably, a phenomenon known as the *attention sink* has been identified, where attention weights are disproportionately concentrated on specific tokens, particularly the Beginning-of-Sentence (<BOS>) token (Xiao et al., 2024). Further analysis reveals that this behavior is driven by specialized attention heads that exhibit a remarkably high propensity for such effects (de Llano et al., 2026). The importance of attention sinks is evidenced by the substantial performance degradation observed when omitting the <BOS> token at inference time in models trained with it (Barbero et al., 2025).

In parallel, structural pruning studies have increasingly focused on the layer-wise sensitivity of LLMs to identify redundant components. A

* Corresponding author.

recurring observation in this line of work is that LLMs exhibit a non-uniform distribution of importance across their depth, with higher layers often being more amenable to pruning or skipping without catastrophic performance loss (Gromov et al., 2025; Zhang et al., 2024; Men et al., 2025). Recent studies further demonstrate that attention heads in higher layers can be effectively pruned, particularly when combined with rescaling to preserve representation stability (Liu and Liu, 2026).

Despite their importance during training, the functional role of attention sink heads at inference remains largely underexplored. Moreover, while prior pruning literature consistently reports redundancy in higher layers, clear criteria for identifying such layers are still lacking, and proposed explanations are often post hoc.

In this work, we make the following contributions:

- We identify <BOS> sink heads (Figure 1) as structurally specialized attention heads that act as *dumping grounds* for superfluous attention weights, providing a mechanistic explanation for redundancy in LLMs.
- We show that the emergence and concentration of these heads in higher layers explain the layer-wise sensitivity consistently observed in prior pruning literature.
- We demonstrate that <BOS> sink scores constitute a simple yet robust structural metric for identifying redundant heads and layers, enabling effective structured pruning while preserving predictive performance.

2 Related Work

2.1 Attention Sink

Attention sinks describe the tendency of LLMs to allocate disproportionate mass to early tokens, especially <BOS>, as a stable outlet for the surplus attention probability (Xiao et al., 2024; Gu et al., 2025). Previous work further suggests that this pattern is not merely incidental noise, but a learned allocation strategy that emerges during training to satisfy normalization constraints while preserving useful token-to-token interactions. Motivated by this account, Miller (2023) proposes adding a constant to the softmax denominator, which allows a head’s attention weights to approach zero so that no token need serve as the outlet.

Recent analysis shows that the phenomenon is concentrated in specialized sink heads and aligned with highly redundant regions of the network (de Llano et al., 2026), supporting the view that many of such heads operate with limited task-specific utility. At the same time, removing <BOS> at inference can substantially degrade models trained to rely on it, indicating that structural anchoring and functional redundancy can coexist (Barbero et al., 2025). This duality motivates a more granular perspective: the anchor token may remain globally important even when a subset of sink-specialized heads is locally expendable. Building on this tension, we ask whether sink heads, rather than the <BOS> token itself, can be pruned safely and systematically.

2.2 Pruning Methods

Prior work reduces LLM cost through two broad paradigms: unstructured weight sparsification and structured removal of architectural units (Han et al., 2015; Ma et al., 2023). Unstructured methods target individual parameters and can achieve high sparsity, as exemplified by magnitude pruning, Wanda, SparseGPT, and Wanda++ (Sun et al., 2024; Frantar and Alistarh, 2023; Yang et al., 2025c). In contrast, structured methods remove entire components (e.g., heads) or full layers, often yielding clearer hardware compatibility and more direct latency gains (An et al., 2024; Ashkboos et al., 2024; Men et al., 2025). Within structured pruning, this distinction is commonly framed as component-level versus layer-level pruning (Song et al., 2024; Yang et al., 2024). Across these settings, multiple studies report that deeper blocks are often more prunable than earlier ones; however, selection criteria are typically derived from magnitude, saliency, similarity, or representation heuristics rather than from an explicit mechanism of functional redundancy (Men et al., 2025; Yang et al., 2025b). Our work complements this line by introducing the <BOS> sink score as a simple structural signal that links head-level redundancy to layer-wise pruning sensitivity.

3 Preliminary

3.1 Attention Sink Metrics

To analyze how attention is allocated across different tokens, we adopt the metrics from Gu et al. (2025). Let $\alpha_{t,k}^{(\ell,h)}$ denote the attention weight from the query token at position t to the key token at




Figure 2: Single-head ablation on MMLU versus $\langle \text{BOS} \rangle$ sink score for Gemma-3-4B, Llama-3.1-8B, and Qwen3-4B. Dots denote heads and dashed lines denote model baselines; high-score heads consistently have minimal impact when ablated.




Figure 3: Relationship between $\langle \text{BOS} \rangle$ sink score and single-head ablation effect on MMLU. Colored and black lines denote layer-mean sink score (right axis) and ablated MMLU accuracy (left axis), respectively. In deeper layers, higher sink scores coincide with lower functional necessity, indicating redundant heads.

position k in the h -th head of the ℓ -th layer.

Attention Sink Score First, we use the attention sink score to quantify the degree to which a specific token position k attracts attention weights within a head. For a sequence of length T , the score for token k in head h of layer ℓ is defined as the average attention weight it receives from all tokens in the sequence:

$$\text{sink-score}_k^{(\ell,h)} = \frac{1}{T} \sum_{t=0}^{T-1} a_{t,k}^{(\ell,h)} \quad (1)$$

A high score indicates that the token at position k serves as a primary destination for attention weights within that specific head.

$\langle \text{BOS} \rangle$ Sink Score Given that the $\langle \text{BOS} \rangle$ token (at index $k = 0$) often acts as a focal point for attention, we define the $\langle \text{BOS} \rangle$ sink score to measure its influence at both the head and layer levels. The head-wise $\langle \text{BOS} \rangle$ sink score $S_{\langle \text{BOS} \rangle}^{(\ell,h)}$ and the layer-wise $\langle \text{BOS} \rangle$ sink score $S_{\langle \text{BOS} \rangle}^{(\ell)}$ are defined as

follows:

$$\begin{aligned} S_{\langle \text{BOS} \rangle}^{(\ell,h)} &= \text{sink-score}_0^{(\ell,h)} \\ S_{\langle \text{BOS} \rangle}^{(\ell)} &= \frac{1}{H} \sum_{h=1}^H S_{\langle \text{BOS} \rangle}^{(\ell,h)} \end{aligned} \quad (2)$$

where H is the total number of attention heads in layer ℓ . The layer-wise $\langle \text{BOS} \rangle$ sink score $S_{\langle \text{BOS} \rangle}^{(\ell)}$ serves as an aggregate metric to evaluate the dependence of layer ℓ on the $\langle \text{BOS} \rangle$ token.

3.2 Block Influence Score

To identify redundant layers within LLMs, Men et al. (2025) leverages the Block Influence (BI) score within the ShortGPT method. This metric quantifies the degree of representation transformation performed by each layer by measuring the distance between its input and output hidden states.

Definition Let $X_{i,t} \in \mathbb{R}^d$ denote the t -th row of the hidden state matrix (corresponding to the t -th token) at the i -th layer. The BI score for the i -th layer is defined as:

$$\text{BI}_i = 1 - \mathbb{E}_{X,t} \frac{X_{i,t}^T X_{i+1,t}}{\|X_{i,t}\|_2 \|X_{i+1,t}\|_2} \quad (3)$$

where $\|\cdot\|_2$ denotes the L_2 norm, and the expectation $\mathbb{E}_{X,t}$ is taken over a dataset X and token positions t .

4 <BOS> Sink as a Marker of Functional Redundancy

<BOS> Sink Score Extraction To analyze the model’s internal behavior, we calculate the <BOS> sink score defined in Equation 2. Specifically, we perform 0-shot inference on the MMLU dataset and extract attention weights directly from the attention maps during the initial decoding step.

Head-wise Ablation Figure 2 plots single-head ablation impact on MMLU versus <BOS> sink score for Llama-3.1-8B, Qwen3-4B, and Gemma-3-4B. Across all models, sensitivity separates into two regimes: high-<BOS> heads are consistently near-baseline when removed, whereas low-<BOS> heads induce larger variance and occasional accuracy drops. This low-score sensitivity is most pronounced in Gemma-3-4B, which we attribute to lower head density (272 heads) relative to Llama-3.1-8B (1,024) and Qwen3-4B (1,152), leaving less functional redundancy per head. Accordingly, high-<BOS> redundancy appears architecture-invariant, while vulnerability to low-<BOS> ablation increases as head density decreases.

Layer-wise Concentration in Deeper Blocks To characterize layer-wise redundancy, we measure the mean ablation impact of heads by layer. Figure 3 shows that <BOS> sink scores are concentrated in deeper layers. In these final blocks, higher mean <BOS> scores coincide with near-zero average accuracy loss under single-head ablation, indicating strong functional redundancy.

Early layers exhibit model-specific behavior. In Qwen3-4B and Gemma-3-4B, ablation impact remains small despite low <BOS> scores, whereas Llama-3.1-8B shows elevated early-layer <BOS> scores with occasional accuracy decline. These results suggest that high <BOS> scores are a reliable redundancy signal in deeper layers, while the role of <BOS> in initial layers is architecture-dependent.

5 Mechanistic Evidence: Sink Heads as Attention Dumping Grounds

Operational Characteristics A <BOS> sink head is defined by two primary traits: (i) a disproportionately high <BOS> sink score $S_{\text{<BOS>}}^{(\ell,h)}$, and (ii) near-zero functional sensitivity under ablation. As illus-

trated in Figure 2, heads in the high-<BOS> regime across Gemma-3, Llama-3.1, and Qwen3 consistently show negligible impact on MMLU accuracy when removed. This behavior suggests that these heads do not transform or convey critical semantic information but instead serve as a stable outlet within the attention mechanism.

The Structural Uniqueness of <BOS> and Redundancy

A key enabler of this phenomenon is the structural uniqueness of <BOS>. Unlike other tokens, it is fixed at position 0 and intentionally non-semantic, encoding neither lexical meaning nor contextual information. It thus provides a stable, content-free target for attention mass under softmax normalization, which forces weights to sum to one even when no semantically relevant target exists (Xiao et al., 2024). We hypothesize that as representations stabilize in deeper layers, some heads increasingly exploit this property and converge on <BOS> as a non-semantic *dumping ground*, whose invariant, context-independent status offers a consistent outlet for surplus attention probability. These heads then contribute to the residual stream in a near-invariant way, shifting from semantic processing to structural maintenance—a transition reflected in Appendix A.

Layer-wise Localization of Redundancy The concentration of these dumping grounds in deeper blocks is driven by the stabilization of representations in higher layers. As the model reaches a consensus on the input’s meaning, additional attention heads become functionally redundant, repurposed as numerical outlets to satisfy the softmax-induced requirement. This localized redundancy explains the minimal performance impact of pruning deeper layers. Consequently, high <BOS> sink scores offer a robust proxy for identifying this layer-specific redundancy, providing a clear empirical basis for efficient structural pruning in the model’s final stages.

6 <BOS> Sink-based Structured Pruning

Based on the evidence of <BOS> sink heads as redundancy repositories, we introduce a structured pruning strategy that employs the <BOS> sink score to identify and eliminate redundant components.

Pruning Criterion We rank architectural components by their <BOS> sink scores and prune those with the highest values. Attention heads are selected by $S_{\text{<BOS>}}^{(\ell,h)}$, and transformer blocks are selected by the layer-wise score $S_{\text{<BOS>}}^{(\ell)}$ (Equation 2).

To maintain stable representations, we exclude the first and last transformer blocks from pruning. The first block is preserved because it builds core lexical and positional features for later layers. The last block is preserved following ShortGPT (Men et al., 2025), which reports that removing its FFN can cause catastrophic performance degradation.

Implementation (Head Ablation) For head-wise pruning, we compute <BOS> sink scores on MMLU (0-shot), rank heads by sink intensity, and select the highest-scoring heads as pruning candidates. We ablate each selected head by zeroing its slice in the output projection matrix (W_O). For a target head h in layer ℓ , we set $W_O^{(\ell,h)} \in \mathbb{R}^{d_{\text{head}} \times d_{\text{model}}}$ to zero:

$$W_O^{(\ell,h)} \leftarrow \mathbf{0} \quad (4)$$

This operation nullifies the contribution of head h before residual aggregation, effectively removing its functional effect while preserving model structure. We do not apply post-pruning rescaling or layer-norm recalibration, so performance changes directly reflect the quality of <BOS> sink score-based head selection.

Pruning Granularity Attention and FFN write to the residual stream as separate sublayers, yet we remove them together as a block. In pre-normalized architectures the update each module writes into the residual stream shrinks relative to the stream itself as depth increases (Men et al., 2025); at those depths the hidden state is carried mostly by the residual path, so deleting a block’s two sublayers together perturbs it little more than deleting either one alone. The block is also the more economical unit: on Gemma-3-4B a 0.125 ratio corresponds to 4 blocks but to 23 of the 34 attention sublayers, and Appendix G shows that removing the latter under the identical <BOS> ranking is far more damaging.

7 Experiments

7.1 Experimental Setup

Models and Benchmarks We perform evaluations using three LLMs that employ Grouped-Query Attention (GQA): Gemma-3-4B/12B (Team et al., 2025), Llama-3.1-8B (Grattafiori et al., 2024), and Qwen3-4B (Yang et al., 2025a). All experiments are conducted with a maximum sequence length of 4,096 tokens.

To assess performance, we use the lm-evaluation-harness¹ (Gao et al., 2024) across multiple benchmark groups. For core language understanding, we evaluate WikiText-2 (Merity et al., 2017), ARC-Easy (Clark et al., 2018), BoolQ (Clark et al., 2019a), OpenbookQA (Mihaylov et al., 2018), and PIQA (Bisk et al., 2020). For reasoning-focused evaluation, we use Winogrande (Sakaguchi et al., 2021), MMLU (Hendrycks et al., 2021), HellaSwag (Zellers et al., 2019), and ARC-Challenge (Clark et al., 2018). For generation, we evaluate HumanEval (Chen et al., 2021), TriviaQA (Joshi et al., 2017), and GSM8K (Cobbe et al., 2021). For long-context and retrieval, we use LongBench (Bai et al., 2024) and report HotpotQA (Yang et al., 2018), QASPER (Dasigi et al., 2021), and GovReport (Huang et al., 2021), following the LongBench configuration. Detailed dataset settings are provided in Appendix C.

Baselines and Implementation We compare our method with Wanda-SP (An et al., 2024), Mag-SP (An et al., 2024), and ShortGPT (Men et al., 2025). For layer-level pruning, we also include two positional baselines: Bottom-Up (pruning from early to late layers) and Top-Down (from late to early layers), to test whether redundancy reflects functional importance or layer position. In all settings, we keep the first and last layers to preserve stable input feature extraction and output representation.

Because the FLAP implementations of Wanda-SP and Mag-SP² do not natively support Grouped-Query Attention (GQA), we adapted their core logic for GQA and reimplemented all baselines under the same framework.

Because Wanda-SP and Mag-SP are component-level criteria, they do not define an ordering over full transformer blocks in their original formulation. To obtain layer-level counterparts, we aggregate their per-component importance scores into a single scalar per block by averaging over all output channels of the attention and MLP projections in that block, and then remove the blocks with the lowest aggregated importance. We denote these adapted baselines Wanda-Layer and Mag-Layer. Full implementation details of Wanda-SP and Mag-SP, together with a discussion of how the structured variants relate to their unstructured counterparts, are given in Appendix F.

¹<https://github.com/EleutherAI/lm-evaluation-harness>

²<https://github.com/CASIA-LMC-Lab/FLAP>

Model	Strategy	Pruning Ratio	Wikitext-2 ↓	ARC-c	ARC-e	BoolQ	HellaS.	MMLU	OBQA	PIQA	WG	Avg.
Gemma-3-4B	Dense	0	10.21	0.533	0.816	0.791	0.580	0.597	0.352	0.788	0.725	0.648
	Mag-sp	0.125	17.98	0.421	0.754	0.624	0.451	0.442	0.324	0.767	0.559	0.543
	Wanda-sp		13.53	0.470	0.786	0.724	0.508	0.540	0.318	0.771	0.635	0.594
	Ours (<BOS>)		11.27	0.522	0.813	0.770	0.571	0.594	0.348	0.788	0.718	0.641
Gemma-3-12B	Dense	0	10.184	0.650	0.872	0.854	0.645	0.741	0.354	0.807	0.803	0.716
	Mag-sp	0.250	25.582	0.389	0.815	0.699	0.420	0.316	0.314	0.788	0.628	0.546
	Wanda-sp		24.619	0.411	0.824	0.661	0.447	0.595	0.310	0.768	0.564	0.573
	Ours (<BOS>)		13.607	0.622	0.846	0.820	0.619	0.711	0.342	0.803	0.792	0.694
Llama-3.1-8B	Dense	0	7.71	0.550	0.822	0.829	0.618	0.653	0.336	0.791	0.781	0.673
	Mag-sp	0.250	315.21	0.225	0.612	0.414	0.321	0.229	0.220	0.695	0.507	0.403
	Wanda-sp		10.75	0.410	0.796	0.760	0.487	0.344	0.330	0.779	0.668	0.572
	Ours (<BOS>)		20.55	0.521	0.810	0.787	0.589	0.524	0.316	0.778	0.722	0.631
Qwen3-4B	Dense	0	17.31	0.584	0.804	0.851	0.532	0.701	0.294	0.753	0.672	0.649
	Mag-sp	0.125	30.91	0.386	0.662	0.569	0.401	0.459	0.248	0.708	0.585	0.502
	Wanda-sp		18.48	0.542	0.760	0.804	0.513	0.653	0.306	0.734	0.637	0.618
	Ours (<BOS>)		19.87	0.568	0.794	0.817	0.534	0.690	0.302	0.755	0.685	0.643

Table 1: Performance comparison of component-level pruning strategies on Gemma-3-4B/12B, Llama-3.1-8B, and Qwen3-4B. We report WikiText-2 perplexity and MCQA task accuracy, where our <BOS>-based strategy consistently outperforms Mag-SP and Wanda-SP while remaining close to dense baselines.

7.2 Perplexity and MCQA Results

Component-level Pruning Table 1 summarizes component-level pruning results. Our <BOS>-based method achieves the highest MCQA average among the pruning baselines while remaining close to the dense baselines.

At 12.5% pruning, <BOS>-based pruning achieves average accuracies of 0.641 and 0.643 on Gemma-3-4B and Qwen3-4B, respectively, compared with dense baselines of 0.648 and 0.649. At 25.0%, it also performs best on Gemma-3-12B (0.694) and Llama-3.1-8B (0.631), consistently outperforming Mag-SP and Wanda-SP. Although Wanda-SP occasionally yields lower WikiText-2 perplexity, our method better preserves MCQA accuracy, suggesting that <BOS> sink scores are a stronger criterion for maintaining reasoning performance under structured pruning. Additional results are provided in Appendix B.

Perplexity and Task Retention Perplexity degrades more sharply than task accuracy under component-level pruning. Under our criterion it stays within one to three points of the dense baseline on Gemma-3-4B, Gemma-3-12B, and Qwen3-4B, whereas Llama-3.1-8B rises from 7.71 to 20.55 at 25%. The two measures also order the criteria differently: on Llama-3.1-8B, Wanda-SP attains a lower perplexity than our method while retaining less MCQA accuracy, and at 12.5% it also trails on GSM8K and HumanEval, leading only on TriviaQA. Perplexity on a single held-out corpus is therefore a limited proxy for task-relevant capa-

bility under structured pruning, which is why we evaluate generation (Table 3) and long-context (Table 4) performance alongside it. The magnitude of the perplexity increase and its dependence on pruning ratio and model scale are examined in Appendix H.

Layer-level Pruning Table 2 reports layer-level pruning, where each step removes full transformer blocks (attention + FFN). Besides the positional baselines and ShortGPT, we compare against Mag-Layer and Wanda-Layer, the block-level adaptations. The <BOS>-based strategy attains the best MCQA average on Gemma-3-4B (0.589), Gemma-3-12B (0.547), and Qwen3-4B (0.570); on Llama-3.1-8B it surpasses every baseline except Top-Down (0.390 vs. 0.426).

The two adapted baselines separate language-modeling loss from task-relevant capability. Wanda-Layer attains the lowest perplexity on Gemma-3-4B (15.18 vs. our 20.02) and Mag-Layer on Qwen3-4B (39.67 vs. our 49.89), yet both trail our MCQA average by 7 and 5 points; at 25.0% they collapse to perplexities of 10^4 – 10^7 and averages of 0.390/0.351 (Gemma-3-12B) and 0.313/0.315 (Llama-3.1-8B), against 0.547 and 0.390 for ours. Averaging a per-channel saliency into a block score thus transfers poorly across depth. The layer-wise <BOS> score instead measures the share of a block’s attention mass that is allocated to the anchor token, a quantity we relate to ablation insensitivity in Section 4, and it is the only criterion in Table 2 never worse than second on MCQA average.

Model	Strategy	Pruning Ratio	Wikitext-2 ↓	ARC-c	ARC-e	BoolQ	HellaS.	MMLU	OBQA	PIQA	WG	Avg.
Gemma-3-4B	Dense	0	10.21	0.533	0.816	0.791	0.580	0.597	0.352	0.788	0.725	0.648
	Bottom Up	0.125	50643.64	0.206	0.314	0.466	0.260	0.241	0.146	0.549	0.488	0.334
	Top Down		45.68	0.379	0.629	0.254	0.427	0.178	0.274	0.717	0.657	0.439
	ShortGPT		20.82	0.307	0.673	0.412	0.431	0.283	0.254	0.743	0.549	0.456
	Mag-Layer		33.28	0.213	0.627	0.619	0.349	0.269	0.242	0.705	0.531	0.444
	Wanda-Layer		15.18	<u>0.400</u>	0.739	<u>0.633</u>	<u>0.478</u>	<u>0.300</u>	<u>0.280</u>	0.762	0.564	<u>0.519</u>
	Ours (<BOS>)		<u>20.02</u>	0.456	<u>0.715</u>	0.701	0.507	0.584	0.300	<u>0.749</u>	0.701	0.589
Gemma-3-12B	Dense	0	10.18	0.650	0.872	0.854	0.645	0.741	0.354	0.807	0.803	0.716
	Bottom Up	0.250	3968.23	0.189	0.312	0.497	0.262	0.250	0.156	0.542	0.509	0.340
	Top Down		191.38	<u>0.407</u>	<u>0.628</u>	0.256	<u>0.365</u>	<u>0.452</u>	<u>0.272</u>	<u>0.670</u>	<u>0.681</u>	<u>0.466</u>
	ShortGPT		305692485.89	0.228	0.342	0.433	0.268	0.234	0.172	0.563	0.489	0.341
	Mag-Layer		173294.34	0.261	0.452	<u>0.506</u>	0.310	0.249	0.188	0.641	0.511	0.390
	Wanda-Layer		18253938.76	0.245	0.350	0.454	0.295	0.236	0.190	0.553	0.481	0.351
	Ours (<BOS>)		63.98	0.430	0.665	0.655	0.450	0.501	0.274	0.708	0.690	0.547
Llama-3.1-8B	Dense	0	7.71	0.550	0.822	0.829	0.618	0.653	0.336	0.791	0.781	0.673
	Bottom Up	0.250	129826.91	0.206	0.270	0.456	0.260	0.228	0.132	0.536	0.493	0.323
	Top Down		1202.52	0.270	0.362	0.622	0.249	0.523	0.220	0.578	0.582	0.426
	ShortGPT		7211.05	<u>0.269</u>	<u>0.436</u>	0.375	<u>0.271</u>	<u>0.342</u>	0.194	<u>0.614</u>	<u>0.542</u>	0.380
	Mag-Layer		136054.85	0.201	0.264	0.383	0.260	0.243	0.116	0.540	0.501	0.313
	Wanda-Layer		85784.66	0.210	0.276	0.402	0.259	0.230	0.132	0.532	0.477	0.315
	Ours (<BOS>)		<u>2017.63</u>	0.212	0.518	<u>0.478</u>	0.284	0.266	<u>0.210</u>	0.635	0.521	<u>0.390</u>
Qwen3-4B	Dense	0	17.31	0.584	0.804	0.851	0.532	0.701	0.294	0.753	0.672	0.649
	Bottom Up	0.125	120.72	0.260	0.513	0.572	0.301	0.293	0.178	0.586	0.499	0.400
	Top Down		72.55	0.408	0.614	<u>0.717</u>	0.392	0.682	0.236	0.648	<u>0.633</u>	0.541
	ShortGPT		41.62	0.460	0.662	0.564	0.450	<u>0.673</u>	<u>0.240</u>	0.680	0.623	<u>0.544</u>
	Mag-Layer		39.67	0.428	0.670	0.708	0.403	0.475	0.218	<u>0.677</u>	0.572	0.519
	Wanda-Layer		6686.82	0.195	0.275	0.388	0.260	0.253	0.124	0.540	0.458	0.312
	Ours (<BOS>)		49.89	<u>0.451</u>	<u>0.666</u>	0.826	<u>0.441</u>	0.611	0.252	0.674	0.642	0.570

Table 2: Performance comparison of layer-level pruning strategies on Gemma-3-4B/12B, Llama-3.1-8B, and Qwen3-4B. Mag-Layer and Wanda-Layer are block-level adaptations of Mag-SP and Wanda-SP. Bold and underline denote the best and second-best result per column within each model. Our <BOS>-based strategy attains the best or second-best MCQA average in every setting, whereas the weight- and activation-derived baselines are competitive on perplexity only.

Model	Strategy	Pruning Ratio	GSM8K	HumanEval	TriviaQA
Gemma-3-4B	Dense	0	0.392	0.348	0.605
	Mag-sp	0.125	0.010	0.049	0.489
	Wanda-sp		0.024	0.171	0.571
	Ours (<BOS>)		0.263	0.341	0.588
Gemma-3-12B	Dense	0	0.715	0.427	0.747
	Mag-sp	0.125	0.040	0.189	0.674
	Wanda-sp		0.287	0.250	0.682
	Ours (<BOS>)		0.667	0.402	0.714
Llama-3.1-8B	Dense	0	0.491	0.366	0.704
	Mag-sp	0.125	0	0.073	0.517
	Wanda-sp		0.150	0.274	0.687
	Ours (<BOS>)		0.445	0.329	0.634
Qwen3-4B	Dense	0	0.861	0.750	0.442
	Mag-sp	0.125	0.017	0.134	0.232
	Wanda-sp		0.188	0.494	0.408
	Ours (<BOS>)		0.317	0.543	0.376

Table 3: Performance comparison of generation tasks under component-level pruning on Gemma-3-4B/12B, Llama-3.1-8B, and Qwen3-4B. We report GSM8K and TriviaQA accuracy, and HumanEval Pass@1.

7.3 Generation Results

Table 3 summarizes generation performance under component-level pruning. We report GSM8K and TriviaQA accuracy, and HumanEval Pass@1. Overall, <BOS>-based pruning is substantially more robust than Mag-SP and remains competitive with Wanda across model families.

Across all four models, <BOS> yields the highest GSM8K and HumanEval scores among the pruned variants, with a particularly wide margin over Mag-SP, which reduces GSM8K accuracy to near zero in every case. TriviaQA behaves differently: <BOS> leads on both Gemma models but trails Wanda-SP on Llama-3.1-8B (0.634 vs. 0.687) and Qwen3-4B (0.376 vs. 0.408), the two families for which Wanda-SP also attains the lower perplexity in Table 1. Retention relative to the dense model is likewise uneven, ranging from 93% of dense GSM8K accuracy on Gemma-3-12B to 37% on Qwen3-4B. These results indicate that <BOS> sink scores preserve generation quality more reliably than the baselines: the margin over the strongest baseline exceeds 12 points on GSM8K for every model, while the two TriviaQA deficits stay within 6 points.

7.4 Long-Context and Retrieval Results

Table 4 summarizes long-context and retrieval performance under component-level pruning. We report HotpotQA and QASPER F1 scores, and GovReport ROUGE. Overall, <BOS>-based pruning consistently outperforms Mag-SP and remains competitive with Wanda across LongBench tasks.

Model	Strategy	Pruning Ratio	HotpotQA	QASPER	GovReport
Gemma-3-4B	Dense	0	0.071	0.101	0.270
	Mag-sp	0.125	0.015	0.023	0.071
	Wanda-sp		0.047	0.043	0.153
	Ours (<BOS>)		0.048	0.063	0.140
Llama-3.1-8B	Dense	0	0.108	0.134	0.302
	Mag-sp	0.125	0.043	0.048	0.098
	Wanda-sp		0.087	0.090	0.191
	Ours (<BOS>)		0.093	0.134	0.237
Qwen3-4B	Dense	0	0.128	0.104	0.300
	Mag-sp	0.125	0.034	0.043	0.132
	Wanda-sp		0.069	0.100	0.247
	Ours (<BOS>)		0.084	0.087	0.235

Table 4: Performance comparison of long-context and retrieval tasks under component-level pruning on Gemma-3-4B, Llama-3.1-8B, and Qwen3-4B. We report HotpotQA and QASPER F1 scores, and GovReport ROUGE.

In particular, <BOS> shows strong gains on retrieval-oriented tasks for Gemma-3-4B and Llama-3.1-8B, including matching the dense QASPER score on Llama-3.1-8B at a pruning ratio of 0.125. On Qwen3-4B, our method remains robust and consistently outperforms Mag-SP, although Wanda performs better on QASPER and GovReport. This pattern indicates model-family sensitivity in the reliability of the <BOS> pruning signal.

8 Validation of the <BOS> Sink-based Pruning Criterion

8.1 Distinctiveness of <BOS> -based Head Selection

To directly examine whether our pruning criterion selects a head set that is truly distinct from existing methods, we compute the intersection-over-union (IoU) between the attention heads pruned by our method and those selected by Mag-SP and Wanda-SP under component-level pruning. As shown in Table 5, our method has consistently low overlap with Mag-SP (0.125, 0.113, and 0.104 for Gemma-3-4B, Llama-3.1-8B, and Qwen3-4B, respectively). Although overlap with Wanda-SP is higher (0.219, 0.453, and 0.431), a substantial portion of heads remains non-overlapping. These results suggest that the <BOS> sink signal captures structural redundancy that is not fully explained by weight magnitude or activation-based criteria alone. Additional experimental details are presented in Appendix E.

8.2 Consistency Across Context Lengths

To further validate the reliability of <BOS> sink scores as a pruning metric, we investigate their stability across varying input sequence lengths T .

Model	Pruning Ratio	Mag-sp	Wanda-sp
Gemma-3-4B	0.125	0.125	0.219
Llama-3.1-8B	0.25	0.113	0.453
Qwen3-4B	0.125	0.104	0.431

Table 5: IoU of pruned attention heads between our method and Mag-SP/Wanda-SP under component-level pruning.

Scaling and Structural Normalization By definition, the <BOS> sink score represents the average of attention weights assigned to the <BOS> token throughout a sequence. Due to the Softmax operation’s normalization—where the sum of attention weights for each query must equal 1—the average attention weight per token naturally scales with $1/T$. Consequently, as T increases, absolute <BOS> sink scores exhibit an inevitable downward trend across all heads. This reduction should be viewed as a structural consequence of the attention mechanism’s normalization rather than a loss of functional importance.

Differential Stability of Sink Heads Notably, we observe a clear correlation between a head’s average <BOS> sink score μ and its robustness to increasing sequence lengths. As illustrated in Figure 4, the rate of decline in attention scores is inversely related to the magnitude of μ : while the average score for all heads drops sharply, heads with $\mu \geq 0.6$ show more tempered decreases, and those with $\mu \geq 0.8$ exhibit the most gradual downward trajectories. This differential stability indicates that higher-scoring <BOS> sink heads appear to partially mitigate the $1/T$ scaling effect, maintaining a disproportionately high focus on the <BOS> token even as the total attention budget is distributed over an increasingly large number of tokens.

Positional Bias and Convergence Furthermore, our analysis reveals a strong positional bias in the emergence and stability of these heads. To quantify this, we use the coefficient of variation ($CV = \sigma/\mu$), where σ and μ represent the standard deviation and mean of the <BOS> sink scores across different sequence lengths. As shown in Figure 5, a clear inverse correlation exists between μ and CV . Crucially, heads with high μ and near-zero CV are predominantly concentrated in the deeper layers of the architecture (indicated by the yellow and green hues). This suggests that as the model’s layers deepen, certain heads converge into stable <BOS> sinks that maintain consistent attention pat-




Figure 4: Mean $\langle \text{BOS} \rangle$ sink score versus sequence length T . Points show input-level means, and dashed, dash-dotted, and solid lines are linear fits for all heads, $\mu \geq 0.6$, and $\mu \geq 0.8$, respectively. Although softmax implies $1/T$ scaling, higher- μ sink heads show shallower slopes, indicating more stable $\langle \text{BOS} \rangle$ focus at longer contexts.




Figure 5: Relationship between mean $\langle \text{BOS} \rangle$ sink score (μ) and coefficient of variation (CV). Each dot denotes an attention head, colored by layer. Higher- μ heads exhibit lower CV and are concentrated in deeper layers, indicating stable sink behavior and layer-wise redundancy.

terns regardless of the input sequence length.

9 Conclusion

In this study, we established the $\langle \text{BOS} \rangle$ sink phenomenon as a primary driver of functional redundancy in LLMs, revealing a strong correlation between high $\langle \text{BOS} \rangle$ sink scores and negligible predictive contribution. Our findings provide a systematic explanation for the high redundancy observed in deeper layers, where specialized heads act as stable repositories for unnecessary attention weights, effectively functioning as internal *dumping grounds*. By bridging the gap between internal attention patterns and layer-wise sensitivity, we demonstrated that structural attention properties can serve as a more efficient and direct alternative to traditional weight-based metrics for model compression. Experiments across Gemma-3, Llama-3.1, and Qwen3 confirm that our straightforward $\langle \text{BOS} \rangle$ -based pruning matches or exceeds the efficacy of established methods like Wanda-SP and Mag-SP.

Future work will focus on validating the $\langle \text{BOS} \rangle$ sink phenomenon across different attention architectures, including MoE and MHA, as well as under modified attention operators that give heads an ex-

plicit outlet for surplus attention mass, to assess whether sink behavior is architecture-dependent. Furthermore, we aim to investigate whether similar attention-anchoring mechanisms exist in other Transformer-based architectures beyond text, such as Vision-Language Models (VLMs). Additionally, investigating the developmental dynamics of these sinks during the pre-training phase may offer deeper insights into how models establish representational stability through structural redundancy.

Limitations

Despite the robust findings presented in this study, several limitations remain.

First, while we established a strong correlation between high <BOS> sink scores and functional redundancy, our work is primarily based on empirical observations. We demonstrated that these specialized heads predominantly emerge and concentrate in the deeper layers of the architecture, yet we do not provide a formal theoretical derivation for the underlying cause of this positional concentration.

Second, although we validated the efficacy of our <BOS>-based pruning strategy across diverse GQA-based architectures—specifically Gemma-3-4B/12B, Llama-3.1-8B, and Qwen3-4B—further experiments are needed on ultra-large-scale models (e.g., 70B and above) and on a broader range of attention architectures beyond GQA, including MoE and MHA, to verify whether these attention patterns and redundancy characteristics generalize.

Third, structured pruning removes entire architectural units rather than individual weights, which is what makes the efficiency gains of Appendix D realizable but also makes the intervention coarse. Degradation therefore grows quickly with the pruning ratio, most visibly in perplexity: on Llama-3.1-8B at 25%, every layer-level strategy including ours exceeds a perplexity of 1,000 while MCQA accuracy remains relatively stable. Our results support the <BOS> sink score in the low-to-moderate regime rather than under aggressive compression.

Fourth, the <BOS> sink score presupposes the standard softmax, whose normalization constraint forces every query’s weights to sum to one and therefore requires surplus mass to land on an existing token. In architectures that supply an explicit outlet instead—for instance, a constant in the softmax denominator (Miller, 2023)—the surplus need not accumulate at position 0, and our score would return uniformly low values and lose its ability to rank heads. Whether the redundancy it tracks disappears in such models, or merely moves to that outlet, remains untested.

Ethical Considerations

Fairness and Potential Bias The pruning process, by removing model parameters, carries the risk of inadvertently amplifying biases toward specific demographic groups inherent in the training data. While our study demonstrates robust performance retention on general reasoning benchmarks

such as MMLU and ARC, the specific impact on social bias and toxicity remains to be fully characterized. Consequently, we strongly advise conducting a comprehensive evaluation to detect any discriminatory behaviors against marginalized communities before deploying the pruned models in production environments.

Use of AI Assistants In accordance with the ACL Policy on AI Assistance, we acknowledge the use of Gemini 3 Pro³ to assist with code debugging and improving the readability and grammar of this manuscript. The final content and responsibility for the work rest solely with the authors.

Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2024-00345809, Research on AI Robustness Against Distribution Shift in Real-World Scenarios; and No. RS-2023-00222663, Center for Optimizing Hyperscale AI Models and Platforms), and the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2026-25507282, Development of AI learning technologies based on brain cognitive science and biological learning principles).

References

- Yongqi An, Xu Zhao, Tao Yu, Ming Tang, and Jinqiao Wang. 2024. [Fluctuation-based adaptive structured pruning for large language models](#). In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI’24/IAAI’24/EAAI’24. AAAI Press.
- Saleh Ashkboos, Maximilian L. Croci, Marcelo Genari Do Nascimento, Torsten Hoeffler, and James Hensman. 2024. [Slicept: Compress large language models by deleting rows and columns](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024. [LongBench: A bilingual, multi-task benchmark for long context understanding](#). In

³<https://deepmind.google/models/gemini/>

- Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3119–3137, Bangkok, Thailand. Association for Computational Linguistics.
- Federico Barbero, Álvaro Arroyo, Xiangming Gu, Christos Perivolaropoulos, Michael Bronstein, Petar Veličković, and Razvan Pascanu. 2025. [Why do llms attend to the first token?](#) *Preprint*, arXiv:2504.02732.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. 2020. [PIQA: reasoning about physical commonsense in natural language](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7432–7439. AAAI Press.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. [Evaluating large language models trained on code](#). *Preprint*, arXiv:2107.03374.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019a. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). 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)*, pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019b. [What does BERT look at? an analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#). *Preprint*, arXiv:1803.05457.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. 2021. [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*, pages 4599–4610, Online. Association for Computational Linguistics.
- Enrique Queipo de Llano, Álvaro Arroyo, Federico Barbero, Xiaowen Dong, Michael Bronstein, Yann LeCun, and Ravid Shwartz-Ziv. 2026. [Attention sinks and compression valleys in llms are two sides of the same coin](#). *Preprint*, arXiv:2510.06477.
- Siqi Fan, Xin Jiang, Xiang Li, Xuying Meng, Peng Han, Shuo Shang, Aixin Sun, and Yequan Wang. 2025. [Not all layers of llms are necessary during inference](#). In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI ’25*.
- Elias Frantar and Dan Alistarh. 2023. Sparsegpt: massive language models can be accurately pruned in one-shot. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, and 5 others. 2024. [The language model evaluation harness](#).
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Andrey Gromov, Kushal Tirumala, Hassan Shapourian, Paolo Glorioso, and Daniel A. Roberts. 2025. [The unreasonable ineffectiveness of the deeper layers](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Xiangming Gu, Tianyu Pang, Chao Du, Qian Liu, Fengzhuo Zhang, Cunxiao Du, Ye Wang, and Min Lin. 2025. [When attention sink emerges in language models: An empirical view](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.

- Song Han, Jeff Pool, John Tran, and William J. Dally. 2015. Learning both weights and connections for efficient neural networks. In *Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’15, page 1135–1143, Cambridge, MA, USA. MIT Press.
- Shwai He, Guoheng Sun, Zheyu Shen, and Ang Li. 2024. [What matters in transformers? not all attention is needed](#). *Preprint*, arXiv:2406.15786.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Torsten Hoefer, Dan Alistarh, Tal Ben-Nun, Nikoli Dryden, and Alexandra Peste. 2021. Sparsity in deep learning: pruning and growth for efficient inference and training in neural networks. *J. Mach. Learn. Res.*, 22(1).
- Luyang Huang, Shuyang Cao, Nikolaus Parulian, Heng Ji, and Lu Wang. 2021. [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*, pages 1419–1436, Online. Association for Computational Linguistics.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [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)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Prakhar Kaushik, Shravan Chaudhari, Ankit Vaidya, Rama Chellappa, and Alan Yuille. 2025. [The universal weight subspace hypothesis](#). *Preprint*, arXiv:2512.05117.
- Songtao Liu and Peng Liu. 2026. [High-layer attention pruning with rescaling](#). *Preprint*, arXiv:2507.01900.
- Xinyin Ma, Gongfan Fang, and Xinchao Wang. 2023. Llm-pruner: on the structural pruning of large language models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc.
- Xin Men, Mingyu Xu, Qingyu Zhang, Qianhao Yuan, Bingning Wang, Hongyu Lin, Yaojie Lu, Xianpei Han, and Weipeng Chen. 2025. [ShortGPT: Layers in large language models are more redundant than you expect](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 20192–20204, Vienna, Austria. Association for Computational Linguistics.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. [Are sixteen heads really better than one?](#) In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. [Can a suit of armor conduct electricity? a new dataset for open book question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391, Brussels, Belgium. Association for Computational Linguistics.
- Evan Miller. 2023. [Attention is off by one](#).
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. [Winogrande: an adversarial winograd schema challenge at scale](#). *Commun. ACM*, 64(9):99–106.
- Jiwon Song, Kyungseok Oh, Taesu Kim, Hyungjun Kim, Yulhwa Kim, and Jae-Joon Kim. 2024. Sleb: streamlining llms through redundancy verification and elimination of transformer blocks. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org.
- Mingjie Sun, Zhuang Liu, Anna Bair, and J. Zico Kolter. 2024. [A simple and effective pruning approach for large language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024. [Efficient streaming language models with attention sinks](#). In *International*

- Conference on Learning Representations*, volume 2024, pages 21875–21895.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025a. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Liangwei Yang, Yuhui Xu, Juntao Tan, Doyen Sahoo, Silvio Savarese, Caiming Xiong, Huan Wang, and Shelby Heinecke. 2025b. [Entropy-based block pruning for efficient large language models](#). *Preprint*, arXiv:2504.03794.
- Yifan Yang, Kai Zhen, Bhavana Ganesh, Aram Galstyan, Goeric Huybrechts, Markus Müller, Jonas M. Kübler, Rupak Vignesh Swaminathan, Athanasios Mouchtaris, Sravan Babu Bodapati, Nathan Susanj, Zheng Zhang, Jack FitzGerald, and Abhishek Kumar. 2025c. [Wanda++: Pruning large language models via regional gradients](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 4321–4333, Vienna, Austria. Association for Computational Linguistics.
- Yifei Yang, Zouying Cao, and Hai Zhao. 2024. [LaCo: Large language model pruning via layer collapse](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 6401–6417, Miami, Florida, USA. Association for Computational Linguistics.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a machine really finish your sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.
- Yang Zhang, Yawei Li, Xinpeng Wang, Qianli Shen, Barbara Plank, Bernd Bischl, Mina Rezaei, and Kenji Kawaguchi. 2024. [Finercut: Finer-grained interpretable layer pruning for large language models](#). *Preprint*, arXiv:2405.18218.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Zican Dong, Yupeng Hou, Beichen Zhang, Yingqian Min, Junjie Zhang, Peiyu Liu, Xiaolei Wang, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, and 4 others. 2026. [A survey of large language models](#). *Frontiers of Computer Science*, 20(12).
- Xunyu Zhu, Jian Li, Yong Liu, Can Ma, and Weiping Wang. 2024. [A survey on model compression for large language models](#). *Transactions of the Association for Computational Linguistics*, 12:1556–1577.




Figure 6: Visualization of four representative attention patterns in Gemma-3-4B during MMLU tasks. From left to right, the maps illustrate: (1) **<BOS> Sink** showing intense focus on the initial token (Layer 31, Head 1); (2) **Diagonal (Self-Attention)** with weights concentrated strictly on the query token (Layer 33, Head 7); (3) **Uniform** broad attention distribution primarily observed in lower layers (Layer 2, Head 6); and (4) **Random** sparse attention patterns (Layer 12, Head 7).

A Patterns of Attention

To qualitatively characterize the attention behaviors underlying model redundancy, we visualize the attention maps of Gemma-3-4B during inference. As illustrated in Figure 6, we identify four representative patterns that define the functional landscape of the attention mechanism: <BOS> Sink, Diagonal, Uniform, and Random. Our qualitative analysis of these signatures yields three primary findings regarding the structural properties of these attention heads.

Dominance of Non-Transformative Patterns A significant majority of attention heads exhibit either <BOS> sink or Diagonal patterns. While <BOS> sinks are known as repositories for unnecessary attention weights, we hypothesize that Diagonal patterns may serve a similar function. In a residual architecture, a head attending solely to its own position effectively acts as an identity map, contributing no new contextual information to the hidden state and thus remaining functionally redundant.

Layer-wise Specialization Uniform patterns, which facilitate broad information mixing, are predominantly concentrated in the lower layers (e.g., Layer 2). This suggests that global contextual integration occurs early in the forward pass, while deeper layers transition toward specialized *dumping grounds* behaviors like <BOS> or Diagonal sinks.

Functional Stability We observed that these attention signatures tend to remain input-invariant. Regardless of changes in the input sequence, in-

dividual heads appear to maintain their specific patterns. This stability suggests that attention behaviors could be interpreted as structural traits inherent to the pre-trained architecture, rather than purely dynamic responses to specific semantic content. While these observations provide meaningful insights, a more rigorous validation remains for future work to confirm the extent of this inherent stability.


B Robustness to Model Pruning

To evaluate our method, we perform component-level pruning on Gemma-3-4B, Qwen3-4B, and Llama-3.1-8B, and test on MMLU (5-shot) and WikiText-2. Across all three model families, the same trend appears: performance degrades more gracefully under our criterion than under baseline pruning rules. This margin is especially clear in moderate pruning regimes, where accuracy remains close to each dense baseline. Figure 7 shows that our approach is notably more robust on MMLU than strong baselines (Wanda-SP, Mag-SP).


Table 6 reports downstream results for component-level pruning. At pruning ratios of 0.125 and 0.25 across WikiText-2 and eight MCQA benchmarks, our <BOS>-based method achieves the best average retention in most settings.

Although <BOS> perplexity rises more sharply beyond a 0.5 pruning ratio, it remains competitive in the practically useful low-to-moderate range (≤ 0.5). Given that heavier pruning often yields non-functional models, this result supports <BOS> as an effective strategy for real-world compression.

Gemma-3-4B



Llama-3.1-8B



Qwen3-4B




Figure 7: Performance comparison of component-level pruning methods on Gemma-3-4B, Qwen3-4B, and Llama-3.1-8B. We report MMLU 5-shot accuracy and WikiText-2 perplexity as a function of the pruning ratio. Our method consistently demonstrates stronger accuracy robustness across all models over the full range of pruning ratios compared to the baselines.

Model	Strategy	Pruning Ratio	Wikitext-2 ↓	ARC-c	ARC-e	BoolQ	HellaS.	MMLU	OBQA	PIQA	WG	Avg.
Gemma-3-4B	Dense	0	10.211	0.533	0.816	0.791	0.580	0.597	0.352	0.788	0.725	0.648
	Mag	0.125	17.983	0.421	0.754	0.624	0.451	0.442	0.324	0.767	0.559	0.543
	Wanda		13.532	0.470	0.786	0.724	0.508	0.540	0.318	0.771	0.635	0.594
	Ours (<BOS>)		11.271	0.522	0.813	0.770	0.571	0.594	0.348	0.788	0.718	0.641
Llama-3.1-8B	Dense	0	7.713	0.550	0.822	0.829	0.618	0.653	0.336	0.791	0.781	0.673
	Mag	0.125	15.408	0.384	0.773	0.580	0.445	0.248	0.326	0.762	0.591	0.514
	Wanda		8.501	0.533	0.812	0.812	0.559	0.585	0.346	0.788	0.743	0.647
	Ours (<BOS>)		11.249	0.559	0.825	0.823	0.611	0.634	0.340	0.789	0.770	0.669
Qwen3-4B	Dense	0	17.311	0.584	0.804	0.851	0.532	0.701	0.294	0.753	0.672	0.649
	Mag	0.125	30.915	0.386	0.662	0.569	0.401	0.459	0.248	0.708	0.585	0.502
	Wanda		18.485	0.542	0.760	0.804	0.513	0.653	0.306	0.734	0.637	0.618
	Ours (<BOS>)		19.867	0.568	0.794	0.817	0.534	0.690	0.302	0.755	0.685	0.643
Qwen3-4B	Dense	0	63.242	0.334	0.547	0.453	0.353	0.258	0.210	0.652	0.511	0.415
	Mag	0.25	20.515	0.486	0.760	0.668	0.450	0.550	0.290	0.729	0.590	0.565
	Wanda		34.869	0.534	0.741	0.747	0.494	0.653	0.246	0.729	0.672	0.602
	Ours (<BOS>)											

Table 6: Downstream performance under component-level pruning at ratios of 0.125 and 0.25. We compare Mag, Wanda, and our <BOS>-based criterion across Gemma-3-4B, Llama-3.1-8B, and Qwen3-4B, reporting Wikitext-2 perplexity (↓) and the average score over eight zero/few-shot benchmarks (↑).

Category	Dataset	Year	Description	Shots	License
Perplexity	WikiText-2	2017	Language modeling/perplexity evaluation on Wikipedia text passages.	0-shot	CC BY-SA 3.0
MCQA	ARC-Easy	2018	Grade-school science multiple-choice questions (easy split).	0-shot	CC BY-SA 4.0
	BoolQ	2019	Yes/no reading comprehension based on short passages.	0-shot	CC BY-SA 3.0
	OpenBookQA	2018	Science multiple-choice reasoning with open-book style facts.	0-shot	Apache 2.0
	PIQA	2020	Physical commonsense reasoning for everyday situations.	0-shot	Apache 2.0
	Winogrande	2019	Pronoun/coreference disambiguation under commonsense constraints.	5-shot	Apache 2.0
	MMLU	2021	Multi-domain knowledge and reasoning across academic/professional subjects.	5-shot	MIT
	HellaSwag	2019	Commonsense sentence completion and grounded continuation selection.	10-shot	MIT
	ARC-Challenge	2018	More difficult split of grade-school science QA.	25-shot	CC BY-SA 4.0
Generation	HumanEval	2021	Code generation from docstrings, evaluated by unit tests.	0-shot	MIT
	TriviaQA	2017	Open-domain question answering using evidence documents.	5-shot	Apache 2.0
	GSM8K	2021	Multi-step grade-school math word-problem solving.	8-shot	MIT
Long Context and Retrieval	HotpotQA	2018	Multi-hop question answering requiring evidence integration across documents.	0-shot	CC BY-SA 4.0
	QASPER	2021	Question answering over scientific papers (long-context comprehension).	0-shot	CC BY 4.0
	GovReport	2021	Long-document government report summarization.	0-shot	CC BY 4.0

Table 7: Summary of datasets used in our pruning experiments. For long-context evaluation, we follow the LongBench configuration and report results on HotpotQA, QASPER, and GovReport.




Figure 8: Wall-clock inference latency (seconds) of original and pruned models under layer-level pruning across sequence lengths from 1,024 to 65,536 tokens. Results are reported for Llama-3.1-8B, Qwen3-4B, and Gemma-3-4B at pruning ratios of 0.125 and 0.25. The proposed method consistently reduces latency across all settings, yielding about 7–13% speedup at 0.125 pruning and 20–25% speedup at 0.25 pruning.

C Dataset Summary for Models and Benchmarks

Table 7 summarizes the benchmark suite used in our experiments, including dataset name, release year, task description, and shot setting. The protocol covers diverse capability axes—language modeling, commonsense and scientific reasoning, coding, mathematical reasoning, multi-hop retrieval, and long-document understanding—to evaluate pruning robustness across domains. Most benchmarks are evaluated in 0-shot to measure retained capability after pruning, while selected reasoning-intensive tasks (e.g., MMLU, ARC-Challenge, GSM8K, and HellaSwag) use few-shot settings following standard LLM practice. This breadth supports the consistency of our pruning observations across evaluation formats. For long-context evaluation, we follow LongBench and report HotpotQA, QASPER, and GovReport, covering both short- and long-context scenarios.

D Efficiency Analysis under Layer-level Pruning

Inference Latency. Figure 8 compares wall-clock inference latency between original and

pruned models under layer-level pruning across sequence lengths from 1,024 to 65,536 tokens. We evaluate Llama-3.1-8B, Qwen3-4B, and Gemma-3-4B at pruning ratios of 0.125 and 0.25. Across all models and sequence lengths, our approach consistently reduces inference time, delivering approximately 7–13% speedup at a pruning ratio of 0.125 and 20–25% speedup at 0.25. These results indicate that the proposed layer-level pruning strategy provides robust and scalable latency gains, especially in long-context inference settings where computational overhead is most substantial.

Parameter Efficiency. Table 8 summarizes parameter efficiency under layer-level pruning by reporting the total number of parameters (in billions) before and after pruning for Llama-3.1-8B, Qwen3-4B, and Gemma-3-4B. Our method consistently reduces parameter counts across all models, with reductions of approximately 9–11% at a pruning ratio of 0.125 and 18–22% at a pruning ratio of 0.25. These results confirm that layer-level pruning yields substantial model compression while preserving a consistent reduction trend across different model families.

Model	Pruning Ratio	Original (B)	Pruned (B)	Reduction (%)
Llama-3.1-8B	0.125	8.03	7.16	10.86
	0.25	8.03	6.29	21.73
Qwen3-4B	0.125	4.02	3.62	10.04
	0.25	4.02	3.11	22.58
Gemma-3-4B	0.125	4.30	3.92	8.78
	0.25	4.30	3.55	17.56

Table 8: Parameter efficiency of our method under layer-level pruning. We report the total number of parameters (in billions) before and after pruning for Llama-3.1-8B, Qwen3-4B, and Gemma-3-4B. Our method achieves parameter reductions of approximately 9–11% at a pruning ratio of 0.125 and 18–22% at a pruning ratio of 0.25.

Model	Pruning Ratio	BOS-Wanda	BOS-Mag	Wanda-Mag
Gemma-3-4B	0.125	0.2188	0.1250	0.3438
	0.25	0.3750	0.1354	0.3125
Gemma-3-12B	0.125	0.3043	0.1159	0.2246
	0.25	0.3992	0.1708	0.2741
Llama-3.1-8B	0.125	0.3803	0.0810	0.1384
	0.25	0.4679	0.1128	0.1913
Qwen3-4B	0.125	0.4445	0.1045	0.1297
	0.25	0.4851	0.1433	0.1755

Table 9: Pairwise IoU overlap of pruned attention-head sets across <BOS>-based pruning, Wanda-SP, and Mag-SP under component-level pruning.

E Extended IoU Analysis

To further characterize what our criterion removes, Table 9 reports pairwise IoU overlap between pruned head sets from <BOS>-based pruning, Wanda-SP, and Mag-SP. A consistent pattern emerges across all model families and sparsity settings: overlap between our method and Mag-SP remains low, while overlap with Wanda-SP is moderate. This indicates that our method is not a simple proxy for weight magnitude and only partially aligns with activation-based saliency. In other words, <BOS> sink scores identify a structurally different subset of redundant heads that conventional criteria do not fully capture. Even when sparsity increases from 0.125 to 0.25 and overlap values rise overall, this relative ordering is preserved, reinforcing that our method captures complementary redundancy signals rather than reproducing baseline selections.

F Implementation Details of Wanda-SP and Mag-SP

Wanda-SP and Mag-SP are the structured-pruning surrogates of Wanda and magnitude pruning introduced as baselines in FLAP (An et al., 2024, Appendix B.1). Because structured pruning must remove an entire output channel (equivalently, an entire head) rather than individual weights, the per-

weight Wanda saliency is aggregated along the row dimension of the weight matrix:

$$S_{i,j}^{\ell} = \left(\sum_{i=1}^{\text{row}} |W_{ij}^{\ell}| \right) \cdot \|X_j^{\ell}\|_2 \quad (5)$$

where $|\cdot|$ denotes the elementwise absolute value and $\|X_j^{\ell}\|_2$ is the ℓ_2 norm of the j -th input feature aggregated over $N \times L$ calibration tokens. Mag-SP uses the same aggregation with the activation term removed. Channels are ranked by S and the lowest-scoring ones are removed. Because the FLAP reference implementation assumes multi-head attention, we adapted the score to Grouped-Query Attention by computing it per query head and removing a key/value group only when all query heads sharing it are selected; all baselines were re-implemented under this common framework.

Divergence from the original Wanda results.

Our Wanda-SP results show substantially larger degradation than the near-lossless numbers reported for Wanda (Sun et al., 2024) at comparable sparsity. This gap is expected, for four reasons.

Granularity. Unstructured Wanda prunes individual weights under a per-output-row comparison group: every row retains its most salient weights, so the row space of each projection is preserved and the head keeps operating, only more sparsely. Wanda-SP deletes whole heads and channels, removing entire subspaces of the residual-stream update. The same saliency signal is thus applied to a strictly coarser and more destructive decision.

Sparsity levels are not comparable across the two settings. The commonly cited Wanda operating point (50% unstructured sparsity) zeroes half of the parameters while keeping every head and channel alive. Our 12.5–25% structured ratios remove far fewer parameters but eliminate complete architectural units, which is the more aggressive intervention in functional terms.

No compensation. FLAP couples structured removal with a bias-compensation term that restores the mean output of the pruned channels; Wanda-SP as a baseline does not include this, and we deliberately apply no post-pruning rescaling or layer-norm recalibration to any method—including ours—so that the reported numbers reflect selection quality alone.

Model family. The original Wanda evaluation is centred on LLaMA-1/2, whereas we evaluate Gemma-3, Llama-3.1, and Qwen3, all of which

Target	Pruning Ratio	# Layers Pruned	Wikitext-2 ↓	Avg. ↑
Attention only	0.125	23	2528008.47	0.331
Attention + MLP	0.125	4	20.02	0.589

Table 10: Effect of the pruning unit on Gemma-3-4B at a 0.125 parameter budget. Both rows use the same layer-wise <BOS> sink ranking and differ only in what is removed. **Avg.** is the mean over the eight MCQA benchmarks of Table 2.

use GQA and therefore carry a markedly higher functional load per head. Comparable drops for Wanda-SP and Mag-SP on generation tasks are also reported by HARP (Liu and Liu, 2026), consistent with our measurements.

Wanda-SP should therefore not be read as a reproduction of unstructured Wanda; it is the closest structured analogue of the same criterion, and it is included so that magnitude- and activation-based selection can be compared against the <BOS> sink score under an identical removal operator.

G Granularity of the Pruning Unit

Section 6 removes full transformer blocks even though the <BOS> sink score is computed from attention alone. To evaluate this choice, we fix the criterion and the parameter budget and vary only the unit that is removed: Attention-only removes attention sublayers ranked by the layer-wise <BOS> score while leaving every FFN intact, and Attention+MLP removes whole blocks under the identical ranking.

Because attention accounts for a minority of the parameters in Gemma-3-4B, a 12.5% parameter budget requires removing 23 of the 34 attention sublayers under Attention-only, against 4 blocks under Attention+MLP. Table 10 reports the outcome: Attention-only yields a WikiText-2 perplexity of 2.5×10^6 and an MCQA average of 0.331, close to the chance level of these eight benchmarks, whereas Attention+MLP yields 20.02 and 0.589. Meeting the budget through attention alone therefore requires removing attention from most of the network, including the layers the criterion ranks as least prunable, whereas block-level removal draws its 4 blocks from the high-scoring deeper layers (Figure 3).

Model	Pruning Ratio	Dense	Ours (<BOS>)	ShortGPT	Bottom Up	Top Down
Gemma-3-4B	0	10.21	—	—	—	—
	0.0625	—	12.62	<u>12.90</u>	16.52	19.35
	0.125	—	20.02	<u>20.82</u>	50,643.64	45.68
Gemma-3-12B	0	10.18	—	—	—	—
	0.0625	—	14.42	14.42	16.00	20.51
	0.125	—	17.27	<u>34.53</u>	147.38	46.85
	0.25	—	63.98	3.06×10^8	3,968.23	191.38
	0.375	—	<u>2,327.06</u>	6.84×10^{12}	3.72×10^7	896.58
	0.5	—	5,292.32	1.40×10^{10}	1.95×10^{11}	312,570.92
Gemma-3-27B	0	7.45	—	—	—	—
	0.0625	—	<u>9.26</u>	9.50	8.43	17.52
	0.125	—	12.01	<u>13.73</u>	107.19	197.59
	0.25	—	22.41	<u>24.96</u>	273.96	953.51
	0.375	—	86.62	46,875.04	48,826.19	22,201.11
	0.5	—	1,038.98	<u>14,295.21</u>	1.13×10^{14}	8.08×10^5

Table 11: WikiText-2 perplexity (↓) under layer-level pruning across model scales and pruning ratios. Bold and underline denote the best and second-best result per row.

H Perplexity across Model Scales and Pruning Ratios

The perplexity cost of structured pruning depends jointly on the pruning ratio and on model size. To characterise this, we extend the layer-level comparison of Table 2 to pruning ratios from 0.0625 to 0.5 and add Gemma-3-27B, evaluated under the same WikiText-2 protocol as the other models.

Table 11 reports the results. At 0.0625 all three models stay within roughly two to four points of their dense baselines (12.62 against 10.21 on Gemma-3-4B, 14.42 against 10.18 on Gemma-3-12B, and 9.26 against 7.45 on Gemma-3-27B). The ratio at which perplexity leaves a usable range increases with model size: at 0.125 the pruned-to-dense ratio is 1.96 on Gemma-3-4B, 1.70 on Gemma-3-12B, and 1.61 on Gemma-3-27B, and Gemma-3-27B still reaches 22.41 at a 0.25 ratio. Beyond 0.25 every strategy degrades sharply, but the gap between them widens rather than closing: at 0.375 on Gemma-3-12B, ShortGPT reaches 6.84×10^{12} while our criterion stays at 2.33×10^3 . Across the twelve pruned settings, our criterion attains the lowest perplexity in ten, counting the tie with ShortGPT at 0.0625 on Gemma-3-12B.




Figure 9: Full Attention Maps for Gemma-3-4B.




Figure 10: Full Attention Maps for Llama-3.1-8B.




Figure 11: Full Attention Maps for Qwen3-4B.