

MITIGATING ATTENTION SINKS AND MASSIVE ACTIVATIONS IN AUDIO-VISUAL SPEECH RECOGNITION WITH LLMs

Anand ♣ Umberto Cappellazzo ♠ Stavros Petridis ♠ Maja Pantic ♠

♣ University of British Columbia, Canada ♠ Imperial College London, UK

ABSTRACT

Large language models (LLMs) have recently advanced auditory speech recognition (ASR), visual speech recognition (VSR), and audio-visual speech recognition (AVSR). However, understanding of their internal dynamics under fine-tuning remains limited. In natural language processing, recent work has revealed attention sinks, tokens that attract disproportionately high attention, and associated massive activations in which some features of sink tokens exhibit huge activation in LLMs. In this work, we are the first to study these phenomena in multimodal speech recognition. Through a detailed analysis of audio-visual LLMs, we identify attention sinks and massive activations not only at the BOS token but also at intermediate low-semantic tokens across ASR, VSR, and AVSR. We show that massive activations originate in the MLP layers and correspond to fixed feature indices across all sink tokens. We further show that intermediate sink tokens exhibit high cosine similarity to the BOS token, thereby amplifying attention and activation. Building on these insights, we introduce a simple decorrelation loss that reduces cosine similarity between BOS and other tokens, effectively mitigating intermediate sinks and massive activations. Furthermore, our method improves word error rate (WER) under high audio-visual feature downsampling while remaining stable at lower downsampling rates.

Index Terms— Audio-Visual Speech Recognition, Attention Sinks, Massive Activations, Large Language Models

1. INTRODUCTION

Pre-trained Large Language Models (LLMs) have shown remarkable ability to adapt to new domains through parameter-efficient fine-tuning [1–7]. Recent works demonstrate their effectiveness in Auditory Speech Recognition (ASR), Visual Speech Recognition (VSR), and Audio-Visual Speech Recognition (AVSR) [8–15]. These approaches extract modality-specific embeddings from pre-trained encoders, downsample them for efficiency, and map them into the LLM embedding space through projectors. The resulting audio and video tokens, concatenated with an instruction prompt, are fed to the LLM, which generates transcriptions autoregressively and is fine-tuned via LoRA [16]. Despite these advances, the internal mechanisms underlying audio-visual LLMs remain poorly understood.

Studies of LLMs in NLP and vision have revealed that, within self-attention, the BOS (Beginning of Sentence) special token and certain semantically uninformative intermediate tokens often attract disproportionately large attention [17–22]. These sink tokens give rise to the phenomenon of attention sinks [17]. While the BOS sink can be useful, acting as key biases that stabilize predictions [18] and mitigate forgetting in long contexts [17], the role of intermediate sinks is less understood. Analyses in NLP suggest they may harm performance [19]. Moreover, sink tokens exhibit massive activations, where a small subset of hidden-state features reach mag-

nitudes up to four orders larger than the median [23–25]. Yet, the interaction between attention sinks and massive activations remains unclear.

Understanding the internal dynamics of audio-visual LLMs is essential for both interpretability and performance. In this context, BOS sinks may aid performance [17, 18], but intermediate sinks risk disrupting the alignment between audio and visual streams by diverting attention from phonetic or lip-movement cues. Similarly, massive activations can over-amplify irrelevant features. Despite their potential impact, these phenomena remain unexplored in speech recognition, leaving a gap in how we understand audio-visual LLMs’ integration of heterogeneous signals.

We present the first extensive analysis of the internal mechanisms of multimodal speech recognition with LLMs. Using Llama-AVSR [8], we reveal the presence of attention sinks at both BOS and intermediate tokens across ASR, VSR, and AVSR tasks. Unlike the BOS sink, which exists in the pre-trained LLM, we observed that intermediate sinks emerge during fine-tuning. We further show that massive activations in these sink tokens originate as early as layer 2 from the MLP component of the transformer block, and that the massively activated feature indices are shared between BOS and intermediate sink tokens. This stems from our key observation that *intermediate sink hidden states exhibit high cosine similarity with the BOS hidden state*.

To probe the role of intermediate sinks, we evaluate performance after mitigating them. We propose a lightweight decorrelation loss that reduces cosine similarity between BOS and other tokens, thereby addressing both attention sinks and massive activations. Prior sink-mitigation strategies, such as prepending placeholder tokens or modifying softmax (e.g., Softmax-off-by-one, SoftPick [17, 26]), require full pretraining and mainly target BOS sinks in long-context settings. Attention Calibration (ACT) [19] adjusts attention during inference but adds overhead and does not address massive activations. In contrast, our method integrates seamlessly with LoRA-based fine-tuning, incurs no inference-time cost, and improves WER across ASR, VSR, and AVSR, even under high compression.

Our key contributions are: (1) We provide the first analysis of attention sinks and massive activations in audio-visual LLMs across ASR, VSR, and AVSR. (2) We identify the origin of massive activations and explain the co-existence of massive activations and attention sinks via cosine similarity. (3) We introduce a novel decorrelation loss that mitigates intermediate sinks and massive activations while improving WER at high compression rates.

2. PRELIMINARIES

Llama-AVSR. We begin our analysis by revisiting the architecture of Llama-AVSR [8], which forms the foundation of our study. In this setting, raw audio and video inputs are first encoded into modality-

specific embeddings using pre-trained encoders. In our setting, we use AV-HuBERT [27] as video encoder and Whisper [28] as the audio encoder. Since these embeddings are high-dimensional and temporally dense, directly feeding them into the LLM would be computationally expensive. To address this, Llama-AVSR applies a compression step that temporally downsamples the embeddings via average pooling, before projecting them into the LLM embedding space using lightweight linear projectors. We denote compression rates as (a, v) for AVSR, where a and v are the downsampling factors for audio and video tokens respectively (e.g., AVSR (16,5)), and as a single value (a) or (v) for unimodal ASR and VSR. The compressed audio $\mathbf{X}_{\text{aud}}$ and video $\mathbf{X}_{\text{vid}}$ tokens are then concatenated with an instruction prompt $\mathbf{X}_{\text{inst}}$ and passed to the LLM, which autoregressively generates the target output $\mathbf{Y}$ as:

$$p(\mathbf{Y}|\mathbf{X}_{\text{aud}}, \mathbf{X}_{\text{vid}}, \mathbf{X}_{\text{inst}}) = \prod_{i=1}^N p(y_i|\mathbf{X}_{\text{aud}}, \mathbf{X}_{\text{vid}}, \mathbf{X}_{\text{inst}}, y_{<i}), \quad (1)$$

where N is number of tokens and $y_{<i}$ is the generated output sequence up to token $i - 1$.

Autoregressive LLMs. Autoregressive LLMs are typically constructed by stacking L transformer decoder blocks [29]. Each block consists of a Multi-Head Self Attention (MHSA) module followed by a Multilayer Perceptron (MLP). Given hidden states $\mathbf{H}^{l-1} \in \mathbb{R}^{N \times d}$ at layer $l - 1$, MHSA computes pair-wise relationships between the tokens with help of queries $\mathbf{Q}_h^l$, keys $\mathbf{K}_h^l$, and values $\mathbf{V}_h^l$ computed for each head h from linear projection of each token's d -dimensional hidden state. The attention map is then computed with:

$$\mathbf{A}_h^l = \text{Softmax} \left(\frac{\mathbf{Q}_h^l \mathbf{K}_h^{l\top}}{\sqrt{d_h}} + \mathbf{M} \right), \quad (2)$$

where $d_h = d/H$ where H is number of heads and $\mathbf{M} \in \mathbb{R}^{N \times N}$ is the causal mask. The head outputs $\mathbf{O}_h^l = \mathbf{A}_h^l \mathbf{V}_h^l$, are concatenated and projected to get the output $\mathbf{O}^l$ of MHSA. In pre-norm LLM blocks, the hidden state for next layer is then computed by performing MHSA on Layer Normalized (LN) hidden state $\mathbf{H}^{l-1}$ as:

$$\mathbf{H}^l = \mathbf{H}^{l-1} + \mathbf{O}^l + \text{MLP}(\text{LN}(\mathbf{H}^{l-1} + \mathbf{O}^l)). \quad (3)$$

The MLP processes each token independently, often in a gated linear unit (GLU) form as follows:

$$\text{MLP}(h) = ((h\mathbf{W}_{\text{up}}) \odot \sigma(h\mathbf{W}_{\text{gate}}))\mathbf{W}_{\text{down}}, \quad (4)$$

where σ is a non-linear activation function and $\mathbf{W}_{\text{gate}}, \mathbf{W}_{\text{up}}, \mathbf{W}_{\text{down}} \in \mathbb{R}^{d \times d'}$ and $\odot$ is element-wise product.

3. ANALYSIS OF AUDIO-VISUAL LLMS

3.1. Attention Sinks

Each element of the attention map $\mathbf{A}_h^l \in \mathbb{R}^{N \times N}$ given by $\mathbf{A}_h^l[i, j]$ represents the attention token i gives to token j . The causal mask $\mathbf{M}$ ensures $\mathbf{A}_h^l[i, j] = 0$ for all $i < j$. Thus, we compute attention score of token i at layer l as average attention it receives from other tokens across all the heads as

$$\alpha_i^l := \frac{1}{H(N - i + 1)} \sum_{h=1}^H \sum_{k=1}^N \mathbf{A}_h^l[k, i]. \quad (5)$$




Fig. 1. (a,b,c) Attention sinks present in BOS and intermediate tokens for different tasks and compression rates (e.g., AVSR task at audio-video compression rates of (16, 5)). (d) Activation magnitudes (z-axis) of the hidden state in Llama-AVSR (16, 5) at layer 5 reveal some features with massive activation in sink tokens.

We perform our analysis on Llama-based LLMs [30] (since the trend is similar across LLMs, we report the results with Llama 3.2-3B) on ASR, VSR, and AVSR tasks with different compression rates and compute the attention scores α_i^l for each token across all layers. Figure 1 (a–c) presents the average attention maps aggregated across all heads and layers. We observe that the initial token consistently receives substantially higher attention compared to other tokens across layers, confirming the presence of a BOS sink. This observation is consistent with prior work [17], which demonstrated that the BOS token serves as an attention sink. Furthermore, after layer 2, certain intermediate tokens begin to attract elevated attention scores as shown in Figure 2(a), suggesting the emergence of intermediate attention sinks. These patterns highlight both BOS and intermediate sinks in ASR, VSR, and AVSR under different compression rates. Computing the attention score at different epochs during fine-tuning showed that these intermediate sinks appear as a result of fine-tuning unlike the BOS sink which we noticed is already present in the pre-trained LLM.

Finally, we also notice that the intermediate sink tokens occur on tokens with low semantic value. This includes special tokens like `<audio>`, `</audio>`, `<video>`, `</video>` and prompt tokens. We hypothesize that this phenomenon occurs because these tokens are consistently present during training, leading the LLM to use them as anchors to absorb excessive attention while optimizing the loss function. Later, in Section 5, we observe that mitigating these intermediate sinks leads to improved WER at high compression rates.




Fig. 2. (a) Intermediate attention sinks originate after layer 2 in Llama-AVSR (16,5). (b) Massive activations in Llama-AVSR (16,5) originate from the MLP of Layer 2.

3.2. Massive Activations

In addition to attention sinks, we analyze *massive activations*, a phenomenon where certain hidden-state features of a token exhibit extremely large magnitudes compared to the median [23], despite the presence of normalization layers. Formally, for token i at layer l , we define the massively activated feature index set as:

$$\Theta_i^l := \left\{ j \in \{1, \dots, d\} \mid |\mathbf{H}^l[i, j]| \geq \tau \cdot \text{median}(|\mathbf{H}^l|) \right\}, \quad (6)$$

where $\mathbf{H}^l[i, j]$ denotes the j -th feature of the hidden state of token i at layer l , and $\text{median}(|\mathbf{H}^l|)$ is computed over the magnitudes of all features across all tokens in layer l . For our analysis, we set $\tau = 10^3$.

Analyzing Θ_i^l across multiple LLMs trained for ASR, VSR, and AVSR on different compression rates, we empirically observed that Θ_i^l is non-empty if and only if i is a sink token and $l \in \{2, 3, \dots, L-1\}$ i.e., massive activations do not happen in first and last layer of the LLM. *This suggests that the phenomena of attention sinks and massive activations co-exist in intermediate layers of the LLM.* Moreover, we found that Θ_i^l is identical for all the sink tokens. Figure 1(d) shows presence of massive activations in all 3 sink tokens with $\Theta_0^5 = \Theta_{20}^5 = \Theta_{21}^5$ in layer 5 of LLM as sinks were present at token indices $\{0, 20, 21\}$.

To further investigate the origin of massive activations, we analyzed the contributions of different components in layer 2 and found that it arises from the MLP module as shown in Figure 2(b). Specifically, in layer 2 of LLaMA3.2-3B, we observed that within the GLU, the term $h\mathbf{W}_{\text{gate}}$ exhibits large positive values for a fixed set of features across all sink tokens, and negative values for non-sink tokens. Due to the non-linear activation function σ (typically SiLU), only the positive values attain high magnitudes. These are further amplified by the element-wise product with $h\mathbf{W}_{\text{up}}$, resulting in massive activations in the MLP latent space $\mathbb{R}^{d'}$. This amplified signal is then propagated to the LLM’s hidden state, producing the observed massive activations in sink tokens through down projection of MLP.

3.3. Cosine Similarity with BOS

To understand why intermediate sink tokens share the same set of massively activated features as the BOS token and the co-existence of massive activations and attention sinks, we analyzed the directional similarity of their hidden states. Specifically, we computed the cosine similarity between the hidden state of each intermediate sink token $\mathbf{H}^l[i, :]$ and that of the BOS token $\mathbf{H}^l[0, :]$ across layers:

$$\cos\text{-sim}(\mathbf{H}^l[i, :], \mathbf{H}^l[0, :]) = \frac{\mathbf{H}^l[i] \cdot \mathbf{H}^l[0]}{\|\mathbf{H}^l[i]\|_2 \|\mathbf{H}^l[0]\|_2}. \quad (7)$$




Fig. 3. (a) Cosine similarity of intermediate sink tokens and non-sink tokens with BOS token across layers of Llama-AVSR (16, 5). (b) Pairwise cosine similarity heatmap of hidden-state embeddings of tokens of Llama-AVSR (16, 5) in layer 5.

We observe that hidden states of intermediate sink tokens are highly aligned with the BOS token from layer 2 onwards as shown in Figure 3(a). This directional alignment explains why intermediate sink tokens exhibit identical massively activated feature indices Θ_i^l and receive excessive attention; as their hidden states point in nearly the same direction as the BOS token, they activate the same set of features and inherit the same attention patterns. *These findings indicate that the root cause of both attention sinks and massive activations in intermediate tokens is the alignment of their hidden states with the BOS token with high cosine similarity.* Figure 3(b) illustrates the cosine similarity between tokens in Llama-AVSR (16, 5). Sink tokens, located at indices $\{0, 20, 21\}$, exhibit very high pairwise cosine similarity, indicating that their hidden states are closely aligned. In contrast, all other tokens show orthogonal behavior with these sink tokens, suggesting that the directional alignment is specific to sink tokens and does not extend to regular tokens.

To test whether BOS alignment drives attention sinks and massive activations, we perform controlled rotations of tokens. Specifically, for an intermediate sink token i , we rotate its hidden state towards the nearest non-sink token $f(i)$ as:

$$\mathbf{H}^l[i] \leftarrow \|\mathbf{H}^l[i]\|_2 \cdot \frac{\mathbf{H}^l[f(i)]}{\|\mathbf{H}^l[f(i)]\|_2}. \quad (8)$$

Applied to sink tokens at indices $\{20, 21\}$ in Llama-AVSR (16, 5), this operation removes both attention sinks and massive activations, as shown in Figure 4(a,b). Conversely, when we rotate a non-sink token towards the BOS token direction,

$$\mathbf{H}^l[i] \leftarrow \|\mathbf{H}^l[i]\|_2 \cdot \frac{\mathbf{H}^l[0]}{\|\mathbf{H}^l[0]\|_2}, \quad (9)$$

we observe the emergence of attention sink behavior and massive activations at that position as shown in Figure 4(c,d) where a sink emerged at index 10.

4. PROPOSED METHOD

Our analysis in Section 3.3 reveals that the root cause of both attention sinks and massive activations in intermediate tokens is their directional alignment with the BOS token in hidden-state space. While these phenomena emerge naturally during training, whether their existence is beneficial for model’s performance is a natural question. Motivated by this, we propose a simple yet effective **decorrelation loss** that explicitly discourages alignment between the BOS token and other tokens, thereby mitigating both attention sinks and massive activations in intermediate tokens.




Fig. 4. (a,b) Sinks and massive activations vanish when disaligned from BOS. (c,d) They emerge when aligned with BOS.

4.1. Decorrelation Loss

Let $\mathbf{H}^l \in \mathbb{R}^{N \times d}$ denote the hidden states at layer l . Our decorrelation loss penalizes similarity between BOS and non-BOS tokens:

$$\mathcal{L}_{\text{decor}} = \frac{1}{(N-1)(L-2)} \sum_{l=2}^{L-1} \sum_{i=1}^{N-1} \cos\text{-sim}(\mathbf{H}^l[i], \mathbf{H}^l[0])^2. \quad (10)$$

We exclude the first and last layers, where massive activations do not occur. Squaring the cosine similarity yields smoother gradients and stronger penalties for highly aligned tokens.

4.2. Final Training Loss

We combine the decorrelation loss with the standard cross-entropy loss for autoregressive speech recognition:

$$\mathcal{L} = \mathcal{L}_{\text{CE}} + \lambda \cdot \mathcal{L}_{\text{decor}}, \quad (11)$$

where λ is a hyperparameter that controls the strength of decorrelation regularization. We report the results of our experiments by selecting the value of $\lambda \in \{10, 10^2, 10^4\}$ that yields the best performance. Importantly, our method requires no modification to the model architecture and adds negligible computational overhead, as cosine similarity is computed directly from hidden states.

5. EXPERIMENTAL RESULTS

Based on Equation 11, we investigate whether the removal of the intermediate attention sinks and massive activations help the model in terms of WER performance. All our experiments follow the training details and the code provided in [8]. We report the results using Llama 3.2-3B [30] as LLM. We observed similar trends for other LLMs (i.e., Llama 3.2-1B and Llama 2-7B).

Table 1. Llama-AVSR WER (%) results with and without decorrelation loss.

Task	Compression	Base	Decorr.	Δ
ASR	(4)	2.62	2.61	+0.01
	(16)	4.83	3.91	+0.92
	(32)	12.92	11.50	+1.42
VSR	(1)	25.84	25.63	+0.21
	(5)	45.19	34.08	+11.11
	(1,1)	2.26	2.22	+0.04
AVSR	(4,2)	2.44	2.42	+0.02
	(16,5)	4.15	3.72	+0.43

Table 2. Comparison between our proposed method and ACT [19].

Task	Compression	Base	ACT	Decorr. (Ours)
ASR	(32)	12.92	12.81	11.50
AVSR	(16,5)	4.15	4.08	3.72

5.1. Decorrelation Loss for Intermediate Attention Sinks

By penalizing BOS alignment, the decorrelation loss encourages intermediate tokens to occupy distinct representational directions in hidden-state space than the initial token. Using our proposed decorrelation loss, we successfully mitigated both attention sinks and massive activations from intermediate tokens. We conducted extensive experiments on the LRS2 dataset for both AVSR and ASR tasks. For VSR, we opted for the LRS3 dataset due to the increased challenge of the task. As shown in Table 1, we observe consistent improvements in WER at high compression rates, while performance remains comparable to the baseline at lower compression rates. These results demonstrate that intermediate attention sinks are detrimental to model robustness under compression, and that decorrelation loss provides an effective and lightweight solution.

5.2. Comparison with Prior Sink Mitigation Methods

Most prior sink mitigation methods target only the BOS sink token in streaming applications with very long context windows and require full model pre-training [17, 26], making them incompatible with our LoRA-based setting. Another approach, Attention Calibration (ACT) [19], focuses on intermediate sink mitigation in NLP tasks by redistributing attention in some selected attention heads. We applied ACT on Llama-AVSR (16,5) and ASR (32) settings on LRS2. Unlike in NLP tasks, we observe only marginal improvements for audio-visual speech recognition tasks. Additionally, ACT fails to mitigate massive activations in sink tokens. Table 2 summarizes the comparison, highlighting that our method effectively addresses both attention sinks and massive activations while providing larger performance gains under high compression rates.

6. CONCLUSION

In this work, we presented the first study of attention sinks and massive activations in multimodal speech recognition LLMs. Our analysis revealed that these phenomena arise from the directional alignment of intermediate tokens with the BOS token. To address this, we proposed a lightweight decorrelation loss that mitigates both effects without architectural changes. Experiments across AVSR, ASR, and VSR showed consistent WER improvements under high compression, demonstrating the effectiveness of our approach.

7. REFERENCES

- [1] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi, “Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models,” in *International conference on machine learning*. PMLR, 2023, pp. 19730–19742.
- [2] Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xiangyu Yue, et al., “Llama-adapter v2: Parameter-efficient visual instruction model,” *arXiv preprint arXiv:2304.15010*, 2023.
- [3] Yassir Fathullah, Chunyang Wu, Egor Lakomkin, Junteng Jia, Yuan Shangguan, Ke Li, Jinxi Guo, Wenhan Xiong, Jay Mahadeokar, Ozlem Kalinli, et al., “Prompting large language models with speech recognition abilities,” in *ICASSP*, 2024, pp. 13351–13355.
- [4] Ziyang Ma, Guanrou Yang, Yifan Yang, Zhifu Gao, Jiaming Wang, Zhihao Du, Fan Yu, Qian Chen, Siqi Zheng, Shiliang Zhang, et al., “An embarrassingly simple approach for llm with strong asr capacity,” *arXiv preprint arXiv:2402.08846*, 2024.
- [5] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan, “Video-chatgpt: Towards detailed video understanding via large vision and language models,” *arXiv preprint arXiv:2306.05424*, 2023.
- [6] Xi Chen, Songyang Zhang, Qibing Bai, Kai Chen, and Satoshi Nakamura, “Llast: Improved end-to-end speech translation system leveraged by large language models,” *arXiv preprint arXiv:2407.15415*, 2024.
- [7] Yuhang Zang, Wei Li, Jun Han, Kaiyang Zhou, and Chen Change Loy, “Contextual object detection with multimodal large language models,” *International Journal of Computer Vision*, vol. 133, no. 2, pp. 825–843, 2025.
- [8] Umberto Cappellazzo, Minsu Kim, Honglie Chen, Pingchuan Ma, Stavros Petridis, Daniele Falavigna, Alessio Brutti, and Maja Pantic, “Large language models are strong audio-visual speech recognition learners,” in *ICASSP*, 2025.
- [9] Marshall Thomas, Edward Fish, and Richard Bowden, “Vallr: Visual asr language model for lip reading,” *arXiv preprint arXiv:2503.21408*, 2025.
- [10] J. Yeo et al., “Where visual speech meets language: Vsp-llm framework for efficient and context-aware visual speech processing,” in *EMNLP Findings*, 2024.
- [11] Umberto Cappellazzo, Minsu Kim, and Stavros Petridis, “Adaptive audio-visual speech recognition via matryoshka-based multimodal llms,” in *ASRU*, 2025.
- [12] Jeong Hun Yeo, Hyeongseop Rha, Se Jin Park, and Yong Man Ro, “Mms-llama: Efficient llm-based audio-visual speech recognition with minimal multimodal speech tokens,” *arXiv preprint arXiv:2503.11315*, 2025.
- [13] Umberto Cappellazzo, Minsu Kim, Stavros Petridis, Daniele Falavigna, and Alessio Brutti, “Scaling and enhancing llm-based avsr: A sparse mixture of projectors approach,” in *Inter-speech*, 2025.
- [14] J. Yeo et al., “Zero-avsr: Zero-shot audio-visual speech recognition with llms by learning language-agnostic speech representations,” in *ICCV*, 2025.
- [15] Umberto Cappellazzo, Minsu Kim, Pingchuan Ma, Honglie Chen, Xubo Liu, Stavros Petridis, and Maja Pantic, “Mome: Mixture of matryoshka experts for audio-visual speech recognition,” in *Advances in neural information processing systems*, 2025.
- [16] E. Hu et al., “Lora: Low-rank adaptation of large language models,” in *ICLR*, 2022.
- [17] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis, “Efficient streaming language models with attention sinks,” *arXiv*, 2023.
- [18] Xiangming Gu, Tianyu Pang, Chao Du, Qian Liu, Fengzhuo Zhang, Cunxiao Du, Ye Wang, and Min Lin, “When attention sink emerges in language models: An empirical view,” *arXiv preprint arXiv:2410.10781*, 2024.
- [19] Zhongzhi Yu, Zheng Wang, Yonggan Fu, Huihong Shi, Khalid Shaikh, and Yingyan Celine Lin, “Unveiling and harnessing hidden attention sinks: Enhancing large language models without training through attention calibration,” *arXiv preprint arXiv:2406.15765*, 2024.
- [20] Nicola Cancedda, “Spectral filters, dark signals, and attention sinks,” *arXiv preprint arXiv:2402.09221*, 2024.
- [21] Federico Barbero, Alvaro Arroyo, Xiangming Gu, Christos Perivolaropoulos, Michael Bronstein, Petar Veličković, and Razvan Pascanu, “Why do llms attend to the first token?,” *arXiv preprint arXiv:2504.02732*, 2025.
- [22] Zihan Qiu, Zekun Wang, Bo Zheng, Zeyu Huang, Kaiyue Wen, Songlin Yang, Rui Men, Le Yu, Fei Huang, Suozhi Huang, et al., “Gated attention for large language models: Non-linearity, sparsity, and attention-sink-free,” *arXiv preprint arXiv:2505.06708*, 2025.
- [23] Mingjie Sun, Xinlei Chen, J. Zico Kolter, and Zhuang Liu, “Massive activations in large language models,” *arXiv preprint arXiv:2402.17762*, 2024.
- [24] Louis Owen, Nilabhra Roy Chowdhury, Abhay Kumar, and Fabian Gra, “A refined analysis of massive activations in llms,” *arXiv preprint arXiv:2503.22329*, 2025.
- [25] Prannay Kaul, Chengcheng Ma, Ismail Elezi, and Jiankang Deng, “From attention to activation: Unravelling the enigmas of large language models,” *arXiv preprint arXiv:2410.17174*, 2024.
- [26] Zayd MK Zuhri, Erland Hilman Fuadi, and Alham Fikri Aji, “Softpick: No attention sink, no massive activations with rectified softmax,” *arXiv preprint arXiv:2504.20966*, 2025.
- [27] Bowen Shi, Wei-Ning Hsu, Kushal Lakhota, and Abdelrahman Mohamed, “Learning audio-visual speech representation by masked multimodal cluster prediction,” *arXiv preprint arXiv:2201.02184*, 2022.
- [28] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever, “Robust speech recognition via large-scale weak supervision,” in *International conference on machine learning*. PMLR, 2023, pp. 28492–28518.
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [30] A. Dubey et al., “The llama 3 herd of models,” *arXiv preprint arXiv:2407.21783*, 2024.