

Attention Sinks in Diffusion Language Models

Maximo Eduardo Rulli^{†*} Simone Petruzzi^{†*} Edoardo Michielon[‡]
Fabrizio Silvestri[†] Simone Scardapane[†] Alessio Devoto[†]

[†]Sapienza University of Rome [‡]Fastweb

 <https://github.com/Maximo-Rulli/dlms-sinks>

Abstract

Masked Diffusion Language Models (DLMs) have recently emerged as a promising alternative to traditional Autoregressive Models (ARMs). DLMs employ transformer encoders with bidirectional attention, enabling parallel token generation while maintaining competitive performance. Although their efficiency and effectiveness have been extensively studied, the internal mechanisms that govern DLMs remain largely unexplored. In this work, we conduct an empirical analysis of DLM attention patterns, focusing on the attention sinking phenomenon, an effect previously observed in various transformer-based architectures. Our findings reveal that DLMs also exhibit attention sinks, but with distinct characteristics. First, unlike in ARMs, the sink positions in DLMs tend to shift throughout the generation process, displaying a dynamic behaviour. Second, while ARMs are highly sensitive to the removal of attention sinks, DLMs remain robust: masking sinks leads to only a minor degradation in performance. These results provide new insights into the inner workings of diffusion-based language models and highlight fundamental differences in how they allocate and utilize attention compared to autoregressive models.

1 Introduction

Large Language Models (LLMs) have driven a paradigm shift across numerous scientific and industrial domains, demonstrating remarkable capabilities in language understanding, generation, and reasoning (Achiam et al., 2023; Anthropic, 2025; Yang et al., 2025a; MetaAI, 2024). This rapid progress is rooted in the transformer architecture and the attention mechanism (Vaswani et al., 2017). While attention is a critical aspect of the

* Equal contribution.

Correspondence to: rulli.2154435@studenti.uniroma1.it, alessio.devoto@uniroma1.it




Figure 1: Cumulative attention scores for each token in LLaDA-8B (Nie et al., 2025) across denoising steps. Unlike autoregressive models, DLMs exhibit attention sinks that shift across the sequence as tokens are progressively unmasked.

transformer’s effectiveness, it also gives rise to complex and often non-intuitive emergent phenomena.

One of the most striking traits of these behaviours is the "attention sink" (Xiao et al., 2023; Miller, 2023). This consists in the fact that, in most autoregressive models (ARMs), a small subset of tokens consistently receives a disproportionate amount of attention from other tokens in the sequence. The pattern is not limited to language, and similar patterns have been observed in Vision Transformers (Darcet et al., 2024) and encoder-only transformers (Ruscio et al., 2025), suggesting it may be a fundamental property of attention-based deep networks.

Recently, masked Diffusion Language Models (DLMs) have emerged as an alternative to the dominant autoregressive paradigm (Nie et al., 2025; Ye et al., 2025; Team, 2025; Labs et al., 2025; Yang et al., 2025b; Wang et al., 2025; Song et al., 2025; Sahoo et al., 2024; Zhu et al., 2025; Liu et al., 2025). Unlike Autoregressive Models (ARMs), which generate text strictly from left to

right, DLMs iteratively refine a fully masked sequence through successive denoising steps (Nie et al., 2025; Ye et al., 2025; Yang et al., 2025b). Generation is based on the unmasking of an initial fully masked sequence of tokens, that the model progressively "denoises" over multiple steps to produce a coherent fully unmasked output. Crucially, DLMs employ a bidirectional attention mechanism. While this bidirectional information flow is key to their parallel, non-causal generation process, the precise impact of this architecture on the inner workings of DLMs remains largely unexplored.

In this work, we present an empirical study of attention patterns in DLMs, focusing specifically on the attention sink phenomenon. We analyse three state-of-the-art open-source masked DLMs: Dream-7B (Ye et al., 2025), a model initialized from a pre-trained ARM; LLaDA-8B (Nie et al., 2025), a large-scale model trained from scratch; and MMaDA-8B (Yang et al., 2025b), a multi-modal DLM trained from LLaDA-8B. Our analysis reveals that DLMs do exhibit attention sinks, but these sinks possess unique dynamic properties rarely seen in their autoregressive counterparts.

Unlike the static attention sinks well-documented in ARMs, most of the sinks in DLMs are unstable and their position actively shifts across the iterative denoising process. Additionally, while ARMs are extremely sensitive to removing the sink tokens, we find that DLMs are significantly more robust to this intervention. We attribute this property to their decoding strategy that unmask only the tokens with highest probabilities in the sequence, and the lack of a causal mask that limits the attention interaction among tokens. To summarize, our primary contributions are the following:

- We conduct an empirical study on attention patterns in DLMs, and provide empirical evidence that attention sinks consistently emerge in these models.
- We characterize the dynamic properties of these sinks, showing they can disappear and shift positions during inference, and we introduce a metric to track their intensity and location across denoising steps.
- We investigate how model performance is affected by removing sinks, and show DLMs are robust to sink masking.

2 Related Work

2.1 Diffusion Language Models

Language modelling has traditionally been dominated by autoregressive models that generate text sequentially, one token at a time. While this paradigm has proven highly successful, DLMs have emerged as an alternative, offering token generation through an iterative denoising processes with potential efficiency advantages (Li et al., 2025b; Wu et al., 2025b; Kim et al., 2025; Liu et al., 2025; Li et al., 2025a; Wu et al., 2025a). Some applications of diffusion to language modelling operate in continuous space, first embedding discrete tokens into continuous vectors, applying diffusion-based denoising, and then mapping back to discrete tokens (Li et al., 2022; Strudel et al., 2022; Gong et al., 2022; Dieleman et al., 2022).

While theoretically elegant, this approach introduces additional complexity in handling the discrete nature of language. A more direct approach emerged with discrete diffusion models, which operate directly on token vocabularies (Austin et al., 2021; Gong et al., 2023; Hoogeboom et al., 2021; Campbell et al., 2022). Starting from fully masked sequences of [MASK] tokens, these models iteratively predict and refine tokens through a process reminiscent of BERT-style masked language modelling (He et al., 2023; Gong et al., 2025). Several works (Austin et al., 2021; He et al., 2023; Gong et al., 2025) have adopted this paradigm but faced significant scaling challenges, remaining limited in size while autoregressive models scaled to billions of parameters.

Recently, discrete DLMs have gained traction thanks to open-source models like Dream-7B (Ye et al., 2025), MMaDA-8B (Yang et al., 2025b) and LLaDA-8B (Nie et al., 2025; Zhu et al., 2025; Liu et al., 2025), which have successfully scaled to 7 billion parameters and beyond, narrowing the performance gap with ARMs.

In this work, we investigate attention patterns in large discrete DLMs, that operate directly on the vocabulary space.

2.2 Attention Sink in Transformers

Attention Sink refers to the common phenomenon observed in transformers where a small subset of tokens consistently receives a disproportionate amount of attention from other tokens in the sequence. This behaviour was initially discovered in Xiao et al. (2023), and leveraged for efficiency.

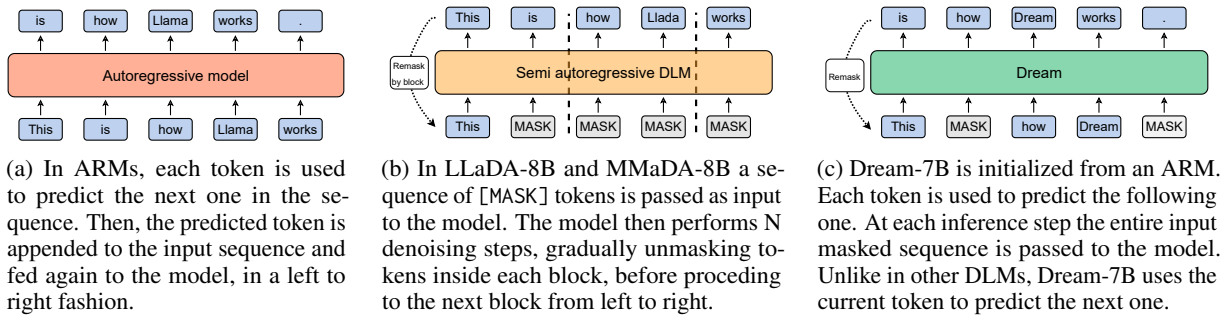


Figure 2: Snapshot of an inference step for different language models. ARMs and Dream-7B predict the next token, while MMaDA-8B and LLaDA-8B predict the current one. MMaDA-8B and LLaDA-8B perform semi-autoregressive block decoding, where only tokens in the current block are unmasked, while Dream-7B may unmask a token at any position.

After this, other works have then explored the sink phenomenon, characterizing properties of sink tokens like high L_2 norm in the hidden state activations (Sun et al., 2024; Cancedda, 2024) or low L_2 norm in the key projection (Devoto et al., 2024; Gu et al., 2024). Similar properties have been also observed in the vision domain (Darcet et al., 2024). Several works have attempted to explain the emergence of attention sinks in transformers. Gu et al. (2025) offers an empirical study of how attention sinks manifest in transformer models, specifically focusing on ARMs. Barbero et al. (2024, 2025) and Pappone (2025) investigate the phenomenon analytically and show how attention sinks act as a bias for ARMs and can mitigate information over-squashing. Finally, Ruscio et al. (2025) analyses attention sinks from a geometric perspective, and shows that they emerge to establish stable coordinate systems in the model’s high-dimensional latent space. While these works analyse sinks in both decoder and encoder transformers, we are the first to observe and investigate this phenomenon in the context of DLMs.

3 Background on Masked Discrete Diffusion

Traditional ARMs model the probability of a text sequence $\mathbf{x} = (x_1, x_2, \dots, x_L)$ of length L by decomposing the joint probability into a product of conditional probabilities, generated in a strict, left-to-right order (Jelinek, 1980; Bengio et al., 2000). This decomposition is given by:

$$p(\mathbf{x}) = p(x_1) \prod_{i=2}^L p(x_i | x_1, \dots, x_{i-1}) \quad (1)$$

where x_i is the token at position i , and $p(x_i | x_1, \dots, x_{i-1})$ is the probability of the current

token conditioned only on all preceding tokens. Masked discrete DLMs offer a non-autoregressive, parallel alternative. Instead of generating tokens one by one, they model a Markov diffusion process over discrete token sequences. This consists of two complementary phases: a fixed **forward corruption process** and a learned **reverse denoising process**. The forward process systematically corrupts a clean data sequence \mathbf{x}_0 (the original text) over a series of time steps $t \in [0, T]$ by progressively replacing tokens with a special mask token [MASK]. Starting with the clean sequence \mathbf{x}_0 , a noisy sequence \mathbf{x}_t at time step t is generated by a Markov transition $q(\mathbf{x}_t | \mathbf{x}_{t-1})$. The marginal distribution of a token x_t^i at time t conditioned on its clean version x_0^i is defined by a masking schedule $\alpha_t \in [0, 1]$. The complete forward process is the joint distribution over all intermediate noisy states, a product of the Markov transitions:

$$q(\mathbf{x}_{1:T} | \mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1}) \quad (2)$$

In the denoising process, a model p_θ , parametrized by θ , reverses this noising process, generating new data from a fully masked sequence \mathbf{x}_T back to a clean sequence \mathbf{x}_0 . More specifically, reverse transition $p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)$ is parameterized by the model, which is trained to estimate the true reverse conditional probability $q(\mathbf{x}_{t-1} | \mathbf{x}_t)$.

In practice, the model p_θ is often trained to predict the clean data \mathbf{x}_0 from the noisy input \mathbf{x}_t at a given time t , and this prediction is then used to approximate the reverse transition. The model output is a distribution over the original tokens, from which the next, less-noisy state \mathbf{x}_{t-1} is sampled.

In this work we consider three masked discrete DLMs: LLaDA-8B (Nie et al., 2025), MMaDA-

8B (Yang et al., 2025b) and Dream-7B (Ye et al., 2025). LLaDA-8B and MMaDA-8B are trained from scratch, with a masked language modelling loss where a token x_i is masked during the forward process, and the model learns to predict the token itself ($x_i \rightarrow [\text{MASK}] \rightarrow x_i$). At inference time, LLaDA-8B and MMaDA-8B use semi-autoregressive block diffusion, where the input sequence is divided into blocks, and the model gradually unmask all tokens inside the corresponding block in a left-to-right manner (Arriola et al., 2025), (see Figure 2). Dream-7B, on the other hand, is initialized from an autoregressive model to leverage the pretrained weights and its training objective employs a "shift operation" (Ye et al., 2025; Gong et al., 2025). More specifically, when a token x_i is masked, Dream-7B is trained to predict x_{i+1} , similarly to an autoregressive model ($x_i \rightarrow [\text{MASK}] \rightarrow x_{i+1}$). In Figure 2 we provide a comparison and visual explanation of how the different types of inference are implemented.

4 Analysis of Attention Sinks in Masked Diffusion Language Models

Previous work has shown that attention sinks emerge in most transformer-based architectures, regardless of the data domain and training strategy (Gu et al., 2025; Ruscio et al., 2025; Xiao et al., 2023; Darcet et al., 2024). Attention sinks are characterized by the disproportionate attention score they receive from all the tokens in the sequence, and can be easily identified as vertical bright lines in attention maps (like the one we show in Figure 1). To validate the presence of attention sinks in DLMs, we first analyse the distribution of attention scores in LLaDA-8B. We see that only a few tokens, the sinks, capture a very high attention score consistently. Similar patterns emerge for Dream-7B and MMaDA-8B (see Appendix A). We now define a metric to characterize and locate attention sinks in DLMs.

4.1 Definition of Attention Sink

Consider an encoder-only transformer model. For a single attention head h and layer l , we have that the attention score is defined as:

$$A_{ij} = \text{softmax}_j \left(\frac{q_i^\top k_j}{\sqrt{d}} \right)$$

where q_i and k_j are the query and key projections for token i and j respectively, and A_{ij} represents the

amount of attention that token i pays to token j . In a DLM attention is bidirectional, and we obtain a distribution of attention scores across the entire sequence *at each denoising step*. Given the attention scores, we define the cumulative attention score for a token j as the average attention it receives from all tokens in a specific denoising step t :

$$\bar{A}_j^{(t,l,h)} = \frac{1}{S} \sum_{i=1}^S A_{ij}^{(t,l,h)}$$

where S is the sequence length, and $A_{i,j}^{(t,l,h)}$ represents the attention score from token i to token j at denoising step t , in head h of layer l . We then identify attention sinks as tokens that receive a cumulative attention score substantially larger than the average.

Attention Sink. We formally define a token j at a specific denoising step t , in head h of layer l to be a **sink token**, if its cumulative attention score exceeds the average cumulative attention score of all other tokens by at least a threshold ϵ :

$$j \text{ is a sink token if } \bar{A}_j^{(t,l,h)} > \frac{1}{S-1} \sum_{k \neq j} \bar{A}_k^{(t,l,h)} + \epsilon \quad (3)$$

This definition ensures that sink tokens represent significant outliers in the attention distribution. In all our claims and experiments we use $\epsilon = 3$ —unless otherwise specified—, which we selected to filter out at least the 96% of tokens in sequence, and empirically showed a sufficient robustness to detect sinks while also serving as a filter for tokens that did not exhibit a sink characteristic. We further discuss the value of ϵ in Appendix B.

4.2 Sink Patterns

Our analysis reveals that DLMs exhibit distinct types of attention sinks with unique dynamic properties not observed in ARMs. We find that sinks do not necessarily appear in the beginning of the sentence, but also show up in the middle or towards the end, which is possible as attention in DLMs is bidirectional. Along with the typical static sink that is frequently observed in ARMs, we identify a new kind of attention sinks that we call **moving sinks**.

Moving sinks appear at different positions during denoising and exhibit widely different patterns according to layer depth and backbone model. Moving sinks are not consistent across diffusion steps, i.e. they do not remain at the same position

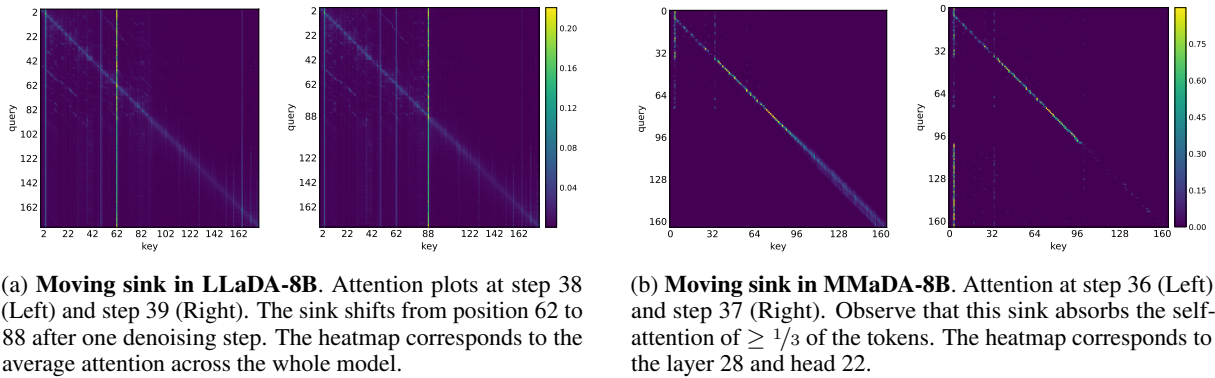


Figure 3: Moving sink in LLaDA-8B, and MMaDA-8B.

across all diffusion steps and may move or even vanish throughout the denoising process. We show an example in Figure 3a. We now analyse how attention sinks appear in the considered pre-trained models.

LLaDA-8B exhibits diverse moving sink patterns with consistency across different sequences. Moving sinks often remain at a specific position for some consecutive denoising steps, before vanishing. Nonetheless, we also find some edge cases in which the moving sinks behave extremely unstably, as we see in Figure 5b, where a sink appears for only one timestep before vanishing on the next one.

As we progress to deeper layers, the number of sinks decreases, converging to one or two sinks per layer (see Figure 4). The deepest layers showcase a particular type of moving sinks, where masked and unmasked tokens maintain separate attention sinks, and switch gradually. We show an example of this phenomenon in Figure 5a. Notably, LLaDA-8B demonstrates a strong semantic basis for sink selection as sinks consistently form on punctuation marks (periods, commas), whitespace, and end-of-sequence tokens. This pattern suggests that LLaDA-8B, trained from scratch as a diffusion model, developed semantically-aware attention mechanisms that identify structurally important tokens as reference points for attention.

Dream-7B showcases a sink behaviour that follows primarily a positional rather than a semantic pattern. Unlike LLaDA-8B, Dream-7B’s sinks often originate at the rightmost masked token and shift leftward as tokens are progressively unmasked, regardless of the token content (see example in Appendix C). This right-to-left migration is most prominent in early layers and creates a dynamic attention flow that follows the unmask-




Figure 4: **Cumulative attention score for LLaDA-8B’s sink across heads and layers.** The variation of the cumulative attention score of the model’s main sink token is displayed across the different heads and layers, averaged through time. In later layers there are usually fewer sinks and the attention score is therefore higher, as it is shared among fewer sink tokens.

ing frontier. This positional nature of Dream-7B’s sinks likely stems from its initialization from a pre-trained autoregressive model. The inherited representations may be less refined for bidirectional attention, causing the model to rely on positional cues rather than semantic content for sink formation. Dream-7B’s positional bias represents a difference from LLaDA-8B’s semantic approach and suggests that initialization strategy and positional embeddings significantly influences attention organization in diffusion models (Ruscio et al., 2025).

MMaDA-8B presents the most stable sink behaviour among the three models, with sinks that are generally static and less frequent. When sinks do manifest they often remain fixed at their initial positions throughout the entire generation process




Figure 5: **Different types of moving sinks in LLaDA-8B.** (a) A particular kind of moving sink in which attention is split according to token type. Some heads exhibit this behaviour in which the masked tokens heavily attend to a specific sink, while the unmasked ones are more concentrated on another one. This heatmap is from layer 30, head 18, and step 32, at the precise end of a block, explaining why we have a perfect line separating all the unmasked and masked tokens. (b) A sink appears at layer 28, head 23, step 96, but suddenly disappears at step 97.




Figure 6: **Example of how sinks move over time.** The largest sink from each model’s specific heads is selected at each iteration. See how the attention shifts according to the explained phenomena. Note that these are sinks for a specific head of the model and not the actual averaged one.

(see example in Appendix C). The model exhibits minimal moving sinks, with most layers showing no clear sink patterns at all. This stability contrasts with the dynamic patterns in LLaDA-8B and Dream-7B, potentially reflecting MMaDA-8B’s different multimodal training data. The static nature of MMaDA-8B’s sinks more closely resembles traditional autoregressive models, though the bidirectional attention mechanism still allows for unique patterns not possible in causal models. For instance, in Figure 3b we show that a considerable amount of tokens shift their attention towards an already unmasked token from one step to the other.

In Figure 6 we show an example of the phenomena described previously for the different models. We select a specific head from each model and compare the position of the largest sink detected

by our metric. We observe that while MMaDA-8B exhibits a mostly static sinking behaviour, sinks tend to shift position in Dream-7B and LLaDA-8B. More specifically, we observe that in LLaDA-8B the sink tends to shift right as more blocks are denoised, while it moves from right to left in Dream-7B.

4.3 Robustness of DLMs to Masking Sinks

Previous studies have demonstrated that attention sinks play a crucial role in transformer-based models, with their removal typically causing catastrophic performance degradation (Xiao et al., 2023; Gu et al., 2025; Barbero et al., 2024). However, given that attention sinks in DLMs exhibit markedly different and more dynamic patterns compared to ARMs, we investigate whether DLMs demonstrate similar sensitivity to sink masking during generation.

We evaluate the three DLM variants — LLaDA-8B, Dream-7B, and MMaDA-8B — on both coding and mathematical reasoning tasks using the GSM8K (Cobbe et al., 2021) and HumanEval (Chen et al., 2021) datasets. GSM8K contains grade-school level math word problems, while HumanEval comprises programming problems designed to evaluate code generation and reasoning capabilities. We evaluate each model in two configurations: (1) the original, unmodified model, and (2) with the top-K sink tokens masked. To determine which tokens to mask, we analyse the attention map from step $t - 1$ across each head and layer. We identify the top-K sink tokens based on our metric (Equation 3) and mask them by setting their attention to 0 during the computation of step t . We vary the threshold parameter ϵ , where smaller values result in masking a larger proportion

Dataset	Sinks	DREAM-7B [58]	LLaDA-8B [38]	MMaDA-8B [57]	LLAMA-3.1-8B [35]
GSM8K	Unmasked	0.82 \pm 0.01	0.76 \pm 0.01	0.54 \pm 0.01	0.85 \pm 0.01
	Masked ϵ_0	0.79 \pm 0.01	0.75 \pm 0.01	0.53 \pm 0.01	0.02 \pm 0.00
	Masked ϵ_1	0.78 \pm 0.01	0.73 \pm 0.01	0.54 \pm 0.01	0.02 \pm 0.00
	Masked ϵ_2	0.75 \pm 0.01	0.55 \pm 0.01	0.37 \pm 0.01	0.01 \pm 0.03
HumanEval	Unmasked	0.60 \pm 0.03	0.37 \pm 0.03	0.16 \pm 0.02	0.66 \pm 0.04
	Masked ϵ_0	0.64 \pm 0.03	0.37 \pm 0.03	0.16 \pm 0.03	0.00 \pm 0.00
	Masked ϵ_1	0.61 \pm 0.03	0.39 \pm 0.03	0.18 \pm 0.03	0.00 \pm 0.00
	Masked ϵ_2	0.57 \pm 0.03	0.35 \pm 0.03	0.09 \pm 0.02	0.00 \pm 0.00

Table 1: **Performance of DLMs and ARMs under attention sink masking.** Thresholds ϵ_0 , ϵ_1 , and ϵ_2 correspond to masking 1, 5, and 10 top-ranked attention sinks, respectively. While LLama-3.1-8B exhibits severe degradation when masking a single sink token, LLaDA-8B, Dream-7B, and MMaDA-8B maintain competitive performance across masking settings, suggesting that parallel inference in DLMs provides robustness to attention sink removal.

of sinks. Specifically, we select ϵ_0 , ϵ_1 and ϵ_2 to mask the top 1, 5 and 10 sinks respectively.

Surprisingly, the tested DLMs exhibit only modest performance degradation when sinks are masked (Table 1). For all the tested DLMs, masking one sink leads to a degradation in performance smaller than 1%. Substantial degradation occurs only when ϵ is decreased further to mask 10 sinks, and mostly in MMaDA-8B. To make a fair comparison with ARMs, we conduct a similar experiment with LLama-3.1-8B, where instead of analyzing the entire attention map to determine the token to mask at the next iteration, we focus on the attention distribution of the last token. We adopt this approach because ARMs, due to their causal nature, cannot ‘shift’ their attention as flexibly as DLMs. In contrast to DLMs, applying this masking procedure to LLama-3.1-8B results in severe performance drops even when masking a single sink token, confirming prior findings that ARMs are highly sensitive to attention sink removal (Xiao et al., 2023; Gu et al., 2025).

We hypothesize that this increased robustness stems from the parallel inference mechanism inherent to DLMs, which may provide alternative attention pathways when primary sinks are unavailable. We explore this hypothesis further in Section 5.2.

Implementation details. We evaluate our models in PyTorch (Paszke et al., 2019) using the checkpoints released on Hugging Face transformers (Wolf et al., 2020) and the official lm evaluation harness scripts (Gao et al., 2024). We use the same hyper-parameters specified in the respective original papers. For LLaDA-8B, we use a block size of 32 and a generation length of 256 tokens for GSM8K and 512 for HumanEval. For

Dream-7B, which does not use semi-autoregressive block generation, we adjust only the generation length and diffusion step parameters according to the original settings. We successfully reproduce the reported results for LLaDA-8B, Dream-7B, and LLama-3.1-8B using these configurations. However, we were unable to reproduce the original results for MMaDA-8B despite following the published implementation details, and we therefore report our own evaluation results for this model.

5 Discussion

5.1 Dynamic Sinks and Positional Encoding

Recent work on encoder-only models notes that attention sinks can shift usually around special markers like [CLS] or [EOS] and connects this behaviour to the use of absolute positional embeddings (Ruscio et al., 2025). However, we find that DLMs, despite using Rotary Positional Embeddings (RoPE, Su et al. 2023), show extremely varied and dynamic sink patterns, including sinks that move and others that split attention between masked and unmasked tokens. These appear all over the text sequence, often on important structural tokens (like punctuation). We further explain these observations in Table 2. The emergence of sink tokens on semantic and formatting markers suggests that the sinking behaviour is driven not only by the positional encoding or token index in the sequence (Ruscio et al., 2025; Barbero et al., 2025), but also by training dynamics and frequency of the token in the training corpus (Sun et al., 2024; Land and Bartolo, 2024).

5.2 Robustness to Masking Sinks

A notable result from Section 4.3 is that DLMs keep working, although with a drop in performance, even when we mask their attention sinks, which would cause an ARM to fail completely. We believe this robustness comes from the bidirectional attention and the iterative denoising process working together to create stability that ARMs lack. In ARMs, attention is causal, and the sink token is usually a single, static anchor, that all future tokens rely on. The next token to predict is therefore usually highly dependent on the sink, and cutting its attention score causes the model to fail.

However, the bidirectional attention in DLMs lets every token see the full context at every denoising step. Additionally, at each step all tokens are considered for unmasking, and only the ones with highest probability (i.e., where the model is most confident) are actually unmasked. This iterative denoising process might ensure higher stability: when a sink is masked, the model may become less confident about those tokens that are highly affected by the sink, and therefore not consider them for unmasking.

Model	Token	Freq.
LLaDA	Ĉ	0.368
	< mdm_mask >	0.366
	.	0.050
	< start_header_id >	0.008
	(white-space)	0.008
Dream	< mask >	0.321
	Ĉ	0.090
	.	0.046
	,	0.044
	< endoftext >	0.040
MMaDA	Ĉ	0.180
	< mdm_mask >	0.166
	< end_header_id >	0.160
	< startoftext >	0.108
	istant	0.066

Table 2: **Relative Frequency of Top-5 Sink Tokens.**

Relative frequency of the top-5 most frequent tokens detected as sinks by our metric. Across all models the Ĉ (a token representing white-spaces) and [MASK] token are the ones at which the models most often sink their attention. We note that in MMaDA the `istant` token comes from `assistant`, which is part of the prompt template. Frequencies computed across 30 generations with generation length of 128, 128 denoising steps and block size of 32.

5.3 Practical Implications

Extending the analyses presented in this study, several concurrent works have further characterized and exploited attention sinks in DLMs. For instance, Zhang et al. (2026) add a new sink token to the model and show that it improves generation quality. Myrzakhan et al. (2026) use DLMs’ attention sinks to perform online pruning of unstable sink tokens, enabling a better quality-efficiency trade-off. While Dai et al. (2026) show how DLMs’ sink mechanism helps them in RAG-related tasks, even outperforming ARMs.

Moreover, in ARMs, attention sinks have been proven to act as a tool to control over-mixing and avoid representation collapse, especially in long contexts (Barbero et al., 2025; Di Giovanni et al., 2023). However, attention sinks in ARMs are usually present only at the beginning of the sequence and represent a single point of reference for the entire generation. In contrast, DLMs offer a flexible inference and their sinks often shift position during generation. Recent work of Long et al. (2026) provides a method for efficiently handling long-context generation by dynamically identifying attention sinks, and retaining them during denoising.

6 Conclusion

We presented the first empirical analysis of attention sinks in Diffusion Language Models, showing that they consistently emerge but behave differently from those in autoregressive models. In DLMs, sinks are dynamic, often shifting across denoising steps and aligning with semantic or structural tokens rather than fixed positions. Moreover, DLMs remain remarkably robust to sink masking, suggesting that their bidirectional and iterative generation distributes attention more evenly and avoids reliance on single anchor tokens. These findings reveal that diffusion models organize attention through flexible mechanisms, offering new insights into their internal dynamics and interpretability.

7 Future Work

While our empirical analysis offers a general overview of sink behaviour in DLMs, it also raises several open questions. It remains unclear what type of information the model stores in the sinks that correspond to future positions. A promising direction to investigate this would be a mechanistic analysis, for instance using the Logit Lens (Nostalgebraist, 2023).

Moreover, although we observed several sink behaviours (e.g., Figure 5a), we did not attempt to provide a detailed explanation of these phenomena. While such an investigation would be valuable, it would require an interpretability-focused study, which lies beyond the scope of this primarily empirical work.

8 Limitations

While this study provides an extensive analysis across three DLMs, it remains a purely empirical work. We focus on documenting the emergence and dynamic behavior of attention sinks through observation rather than providing a formal theoretical framework or a detailed mechanistic explanation of these phenomena. Furthermore, while we conducted an extensive study across three DLMs, our analysis is limited to instruct models, as we did not perform experiments on their corresponding base versions. Finally, we focused on attention sinks in pre-trained models and did not explore how modifications to the training procedure might influence their behaviour, an aspect that has recently been investigated for ARMs by Miller (2023); OpenAI et al. (2025).

9 Acknowledgements

We thank Fastweb S.p.a. for providing the computational resources used in this paper. This work was supported by Sapienza grant RG123188B3EF6A80 (CENTS). We also thank Jary Pomponi, Pasquale Minervini and Emile van Krieken for helpful discussions and valuable feedback.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Anthropic. 2025. System card: Claude opus 4 & claude sonnet 4. *arxiv*.

Marianne Arriola, Aaron Gokaslan, Justin T Chiu, Zhihan Yang, Zhixuan Qi, Jiaqi Han, Subham Sekhar Sahoo, and Volodymyr Kuleshov. 2025. Block diffusion: Interpolating between autoregressive and diffusion language models. In *The Thirteenth International Conference on Learning Representations*.

Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. 2021. Structured denoising diffusion models in discrete state-spaces.

Advances in neural information processing systems, 34:17981–17993.

- Federico Barbero, Andrea Banino, Steven Kapturovski, Dharshan Kumaran, João Guilherme Madeira Araújo, Alex Vitvitskyi, Razvan Pascanu, and Petar Veličković. 2024. Transformers need glasses! information over-squashing in language tasks. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- 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.
- Yoshua Bengio, Réjean Ducharme, and Pascal Vincent. 2000. A neural probabilistic language model. *Advances in neural information processing systems*, 13.
- Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and Arnaud Doucet. 2022. A continuous time framework for discrete denoising models. *Advances in Neural Information Processing Systems*, 35:28266–28279.
- Nicola Cancedda. 2024. Spectral filters, dark signals, and attention sinks. *Preprint*, arXiv:2402.09221.
- 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. *arXiv preprint*.
- 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. *arXiv preprint arXiv:2110.14168*.
- Xin Dai, Pengcheng Huang, Zhenghao Liu, Shuo Wang, Yukun Yan, Chaojun Xiao, Yu Gu, Ge Yu, and Maosong Sun. 2026. Revealing the attention floating mechanism in masked diffusion models. *arXiv preprint arXiv:2601.07894*.
- Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. 2024. Vision transformers need registers. In *The Twelfth International Conference on Learning Representations*.
- Alessio Devoto, Yu Zhao, Simone Scardapane, and Pasquale Minervini. 2024. A simple and effective l_2 norm-based strategy for kv cache compression. *The 2024 Conference on Empirical Methods in Natural Language Processing*.
- Francesco Di Giovanni, Lorenzo Giusti, Federico Barbero, Giulia Luise, Pietro Lio, and Michael M. Bronstein. 2023. On over-squashing in message passing

- neural networks: The impact of width, depth, and topology. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 7865–7885. PMLR.
- Sander Dieleman, Laurent Sartran, Arman Roshanai, Nikolay Savinov, Yaroslav Ganin, Pierre H Richemond, Arnaud Doucet, Robin Strudel, Chris Dyer, Conor Durkan, and 1 others. 2022. Continuous diffusion for categorical data. *arXiv preprint arXiv:2211.15089*.
- 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](#).
- Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An, Peilin Zhao, Wei Bi, Jiawei Han, Hao Peng, and Lingpeng Kong. 2025. [Scaling diffusion language models via adaptation from autoregressive models](#). *Preprint*, arXiv:2410.17891.
- Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An, Peilin Zhao, Wei Bi, Jiawei Han, and 1 others. 2023. Scaling diffusion language models via adaptation from autoregressive models. In *The Thirteenth International Conference on Learning Representations*.
- Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and Lingpeng Kong. 2022. Diffuseq: Sequence to sequence text generation with diffusion models. In *The Eleventh International Conference on Learning Representations*.
- Xiangming Gu, Tianyu Pang, Chao Du, Qian Liu, Fengzhuo Zhang, Cunxiao Du, Ye Wang, and Min Lin. 2024. When attention sink emerges in language models: An empirical view. *arXiv preprint arXiv:2410.10781*.
- 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*.
- Zhengfu He, Tianxiang Sun, Qiong Tang, Kuanning Wang, Xuanjing Huang, and Xipeng Qiu. 2023. Diffusionbert: Improving generative masked language models with diffusion models. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. 2021. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in neural information processing systems*, 34:12454–12465.
- Frederick Jelinek. 1980. Interpolated estimation of markov source parameters from sparse data. In *Proc. Workshop on Pattern Recognition in Practice, 1980*.
- Jaeyeon Kim, Lee Cheuk-Kit, Carles Domingo-Enrich, Yilun Du, Sham Kakade, Timothy Ngotiaoco, Sit-tan Chen, and Michael Albergo. 2025. Any-order flexible length masked diffusion. *arXiv preprint arXiv:2509.01025*.
- Inception Labs, Samar Khanna, Siddhant Kharbanda, Shufan Li, Harshit Varma, Eric Wang, Sawyer Birnbaum, Ziyang Luo, Yanis Miraoui, Akash Palrecha, Stefano Ermon, Aditya Grover, and Volodymyr Kuleshov. 2025. [Mercury: Ultra-fast language models based on diffusion](#). *arXiv*.
- Sander Land and Max Bartolo. 2024. [Fishing for magikarp: Automatically detecting under-trained tokens in large language models](#). *Preprint*, arXiv:2405.05417.
- Jinsong Li, Xiaoyi Dong, Yuhang Zang, Yuhang Cao, Jiaqi Wang, and Dahua Lin. 2025a. Beyond fixed: Training-free variable-length denoising for diffusion large language models. *arXiv preprint arXiv:2508.00819*.
- Tianyi Li, Mingda Chen, Bowei Guo, and Zhiqiang Shen. 2025b. [A survey on diffusion language models](#). *Preprint*, arXiv:2508.10875.
- Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. 2022. Diffusion-lm improves controllable text generation. *Advances in neural information processing systems*, 35:4328–4343.
- Xiaoran Liu, Zhigeng Liu, Zengfeng Huang, Qipeng Guo, Ziwei He, and Xipeng Qiu. 2025. [Longllada: Unlocking long context capabilities in diffusion llms](#). *Preprint*, arXiv:2506.14429.
- Lingkun Long, Yushi Huang, Shihao Bai, Ruihao Gong, Jun Zhang, Ao Zhou, and Jianlei Yang. 2026. Focus-dllm: Accelerating long-context diffusion llm inference via confidence-guided context focusing. *arXiv preprint arXiv:2602.02159*.
- MetaAI. 2024. Introducing llama 4: Advancing multimodal intelligence. *arXiv*.
- MetaAI. 2025. The llama 3 herd of models. *arXiv*.
- Evan Miller. 2023. Attention is off by one. <https://www.evanmiller.org/attention-is-off-by-one.html>.
- Aidar Myrzakhan, Tianyi Li, Bowei Guo, Shengkun Tang, and Zhiqiang Shen. 2026. Sink-aware pruning for diffusion language models. *arXiv preprint arXiv:2602.17664*.
- Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. 2025. Large language diffusion models. *arXiv*.

- Nostalgebraist. 2023. Interpreting gpt: the logit lens.
- OpenAI, :, Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K. Arora, Yu Bai, Bowen Baker, Haiming Bao, Boaz Barak, Ally Bennett, Tyler Bertao, Nivedita Brett, Eugene Brevdo, Greg Brockman, Sebastien Bubeck, and 108 others. 2025. [gpt-oss-120b and gpt-oss-20b model card](#). *Preprint*, arXiv:2508.10925.
- Francesco Pappone. 2025. Attention sinks from the graph perspective. <https://publish.obsidian.md/the-tensor-throne/Transformers+as+GNNs/Attention+sinks+from+the+graph+perspective>.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, and 2 others. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). *Preprint*, arXiv:1912.01703.
- Valeria Ruscio, Umberto Nanni, and Fabrizio Silvestri. 2025. What are you sinking? a geometric approach on attention sink. *arXiv preprint arXiv:2508.02546*.
- Subham Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin Chiu, Alexander Rush, and Volodymyr Kuleshov. 2024. Simple and effective masked diffusion language models. *Advances in Neural Information Processing Systems*, 37:130136–130184.
- Yuxuan Song, Zheng Zhang, Cheng Luo, Pengyang Gao, Fan Xia, Hao Luo, Zheng Li, Yuehang Yang, Hongli Yu, Xingwei Qu, and 1 others. 2025. Seed diffusion: A large-scale diffusion language model with high-speed inference. *arXiv preprint arXiv:2508.02193*.
- Robin Strudel, Corentin Tallec, Florent Althé, Yilun Du, Yaroslav Ganin, Arthur Mensch, Will Grathwohl, Nikolay Savinov, Sander Dieleman, Laurent Sifre, and 1 others. 2022. Self-conditioned embedding diffusion for text generation. *arXiv preprint arXiv:2211.04236*.
- Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. 2023. [Roformer: Enhanced transformer with rotary position embedding](#). *Preprint*, arXiv:2104.09864.
- Mingjie Sun, Xinlei Chen, J. Zico Kolter, and Zhuang Liu. 2024. [Massive activations in large language models](#). *Preprint*, arXiv:2402.17762.
- Gemini Team. 2025. Gemini diffusion. <https://deepmind.google/models/gemini-diffusion/>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Proceedings of the 31st International Conference on Neural Information Processing Systems*.
- Yinjie Wang, Ling Yang, Bowen Li, Ye Tian, Ke Shen, and Mengdi Wang. 2025. Revolutionizing reinforcement learning framework for diffusion large language models. *arXiv preprint arXiv:2509.06949*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Chengyue Wu, Hao Zhang, Shuchen Xue, Shizhe Diao, Yonggan Fu, Zhijian Liu, Pavlo Molchanov, Ping Luo, Song Han, and Enze Xie. 2025a. Fast-dllm v2: Efficient block-diffusion llm. *arXiv preprint arXiv:2509.26328*.
- Chengyue Wu, Hao Zhang, Shuchen Xue, Zhijian Liu, Shizhe Diao, Ligeng Zhu, Ping Luo, Song Han, and Enze Xie. 2025b. Fast-dllm: Training-free acceleration of diffusion llm by enabling kv cache and parallel decoding. *arXiv preprint arXiv:2505.22618*.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2023. Efficient streaming language models with attention sinks. *International Conference on Learning Representations*.
- 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. *arXiv*.
- Ling Yang, Ye Tian, Bowen Li, Xinchen Zhang, Ke Shen, Yunhai Tong, and Mengdi Wang. 2025b. Mmada: Multimodal large diffusion language models. *arXiv preprint arXiv:2505.15809*.
- Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng Kong. 2025. Dream 7b: Diffusion large language models. *arXiv*.
- Zihou Zhang, Zheyong Xie, Li Zhong, Haifeng Liu, and Shaosheng Cao. 2026. One token is enough: Improving diffusion language models with a sink token. *arXiv preprint arXiv:2601.19657*.
- Fengqi Zhu, Zebin You, Yipeng Xing, Zenan Huang, Lin Liu, Yihong Zhuang, Guoshan Lu, Kangyu Wang, Xudong Wang, Lanning Wei, Hongrui Guo, Jiaqi Hu, Wentao Ye, Tiejuan Chen, Chenchen Li, Chengfu Tang, Haibo Feng, Jun Hu, Jun Zhou, and 7 others. 2025. [Llada-moe: A sparse moe diffusion language model](#). *Preprint*, arXiv:2509.24389.




Figure 7: Distribution of attention scores in LLaDA-8B (left), Dream-7B (centre), and MMaDA-8B (right). The y-axis is the frequency with which a token receives the cumulative attention score indicated in the x-axis. Note that the majority of tokens don't receive a cumulative attention score bigger than 5, while there are evident outliers, which are the attention sinks.

A Cumulative Attention Scores Distribution

In Figure 7 we show additional plots of the cumulative attention score distribution across all denoising steps, displaying how a only a few tokens, the sinks, receive a disproportionate high attention score.

B Selection of Sink Threshold

In Equation 3 we defined ϵ to be the threshold for classifying a token as a sink. In Figure 8 we show how the value of ϵ affects sink selection. We see that all the analysed DLMs filter out at least 96% of tokens when using $\epsilon = 3$.




Figure 8: Percentage of tokens not considered as sinks when increasing the value of ϵ , for a sequence of 64 tokens. A balanced threshold is found at $\epsilon = 3$, which we mainly used in this investigation to define that a token is a sink.


C Characteristic Sinks

We provide examples of characteristic moving sinks in Dream-7B and MMaDA-8B, see Figure 9.


D Prompt Template

Across this study we used the standard Hugging Face prompt formatting for instruct models, which defaults to:

```
<|startoftext|><|start_header_id|>user<|end_header_id|>
User's prompt
<|eot_id|><|start_header_id|>assistant<|end_header_id|>
Model's response
```



(a) **Fixed sink in MMaDA-8B.** MMaDA-8B often exhibits a static sink at the beginning of the sequence. In different denoising steps (0 and 127), the sink stays consistently at the beginning of the sequence. Heatmap extracted from layer 12 and head 9.



(b) **Moving sinks in Dream-7B** typically shift from right to left. The sink moving is on step 32 (Left) and at the rightmost position. While at step 33 (Right) the sink has moved towards the centre. The heatmap was extracted from layer 5 and head 6.

Figure 9: Fixed sink in MMaDA-8B and moving sink in Dream-7B.