

One Token Is Enough: Improving Diffusion Language Models with a Sink Token

Zihou Zhang^{1,2}, Zheyong Xie², Li Zhong², Haifeng Liu², Shaosheng Cao^{2*}

¹University of Chinese Academy of Sciences ²Xiaohongshu Inc.

Correspondence: ssskywalkerzzh@gmail.com

Abstract

Diffusion Language Models (DLMs) have emerged as a compelling alternative to autoregressive approaches, enabling parallel text generation with competitive performance. Despite these advantages, there is a critical instability in DLMs: the moving sink phenomenon. Our analysis indicates that sink tokens exhibit low-norm representations in the Transformer’s value space, and that the moving sink phenomenon serves as a protective mechanism in DLMs to prevent excessive information mixing. However, their unpredictable positions across diffusion steps undermine inference robustness. To resolve this, we propose a simple but effective extra sink token implemented via a modified attention mask. Specifically, we introduce a special token constrained to attend solely to itself, while remaining globally visible to all other tokens. Experimental results demonstrate that introducing a single extra token stabilizes attention sinks, substantially improving model performance. Crucially, further analysis confirms that the effectiveness of this token is independent of its position and characterized by negligible semantic content, validating its role as a robust and dedicated structural sink. Our code is available at <https://github.com/skywalker0523/OneTokenIsEnough>.

1 Introduction

Diffusion Language Models (DLMs) have emerged as a compelling alternative to autoregressive (AR) approaches, introducing a paradigm shift towards parallel text generation. Unlike the left-to-right sequential decoding of AR models, DLMs utilize bidirectional attention to enable global context modeling. This mechanism has demonstrated remarkable potential, with recent works such as LLaDA (Nie et al., 2025b) and Dream (Ye et al., 2025) achieving competitive performance in complex generation tasks.

Nevertheless, despite these capabilities, DLMs remain constrained by practical limitations, including inefficient KV caching (Song et al., 2025; Wu et al., 2025) and complex optimization strategies (Zhao et al., 2025). Beyond these general issues, a more specific instability arises from the standard Transformer backbone: the **attention sink** phenomenon. While AR models inadvertently stabilize this phenomenon by anchoring excessive attention to a fixed initial token via causal masking (Xiao et al., 2023; Barbero et al., 2025), DLMs lack such a structural constraint. In the absence of a consistent start token and causal mask, the sink position in DLMs shifts erratically across diffusion timesteps and layers (Rulli et al., 2025). This unpredictable behavior contrasts sharply with the stability of AR models, posing a unique challenge to inference robustness.

In this work, we analyze the characteristics of sink tokens in DLMs through the lens of the transformer value space. As shown in Figure 1, we observe that sink tokens in DLMs consistently exhibit the lowest L_2 norms. This behavior functions as an implicit regularization mechanism, where the model directs attention to these low-norm tokens to mitigate excessive information mixing (Gu et al., 2024; Barbero et al., 2025), relevant details are shown in Appendix A.3. Consequently, the inference process inherently offloads excess attention onto tokens with negligible semantic information.

However, this dependency on low-information tokens creates instability due to the stochastic input masking in DLMs. Unlike autoregressive models that utilize a fixed initial token as a stable anchor, DLMs lack a static, information-sparse token. Instead, they typically utilize masked tokens as attention sinks. Since the set of masked tokens changes dynamically across diffusion steps, there is no consistent position for attention offloading. This moving sink phenomenon introduces structural instability during inference and potentially

*Corresponding author.

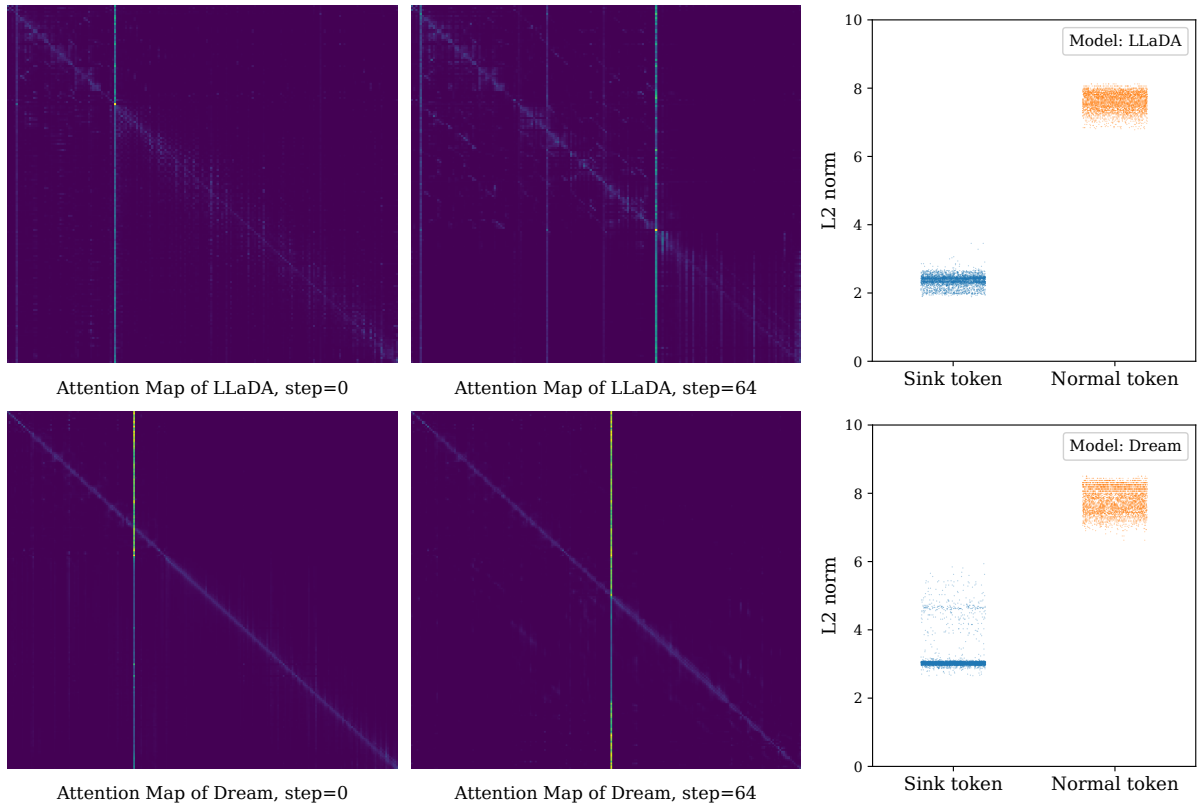


Figure 1: Moving attention sinks and their low-norm property in diffusion language models. Attention maps show representative patterns of LLaDA-Instruct (top) and Dream-Instruct (bottom) at two denoising steps ($step=0$ and $step=64$). Bright vertical stripes mark sink tokens that attract concentrated attention; their positions shift across steps, demonstrating moving sinks. Norm scatter plots compare the L_2 norms of value vectors for 10K sink tokens and 10K normal tokens sampled during inference on GSM8K (Cobbe et al., 2021) and HumanEval (Chen et al., 2021). Sink tokens exhibit significantly lower mean L_2 norms compared to normal tokens (2.36 vs. 7.60 for LLaDA; 3.15 vs. 7.79 for Dream). This consistent disparity suggests that DLMs preferentially select low-norm tokens as sinks during generation.

constrains model performance. This raises a pivotal question: Can we introduce a dedicated, static sink token to effectively regulate information flow and stabilize the moving sink behavior?

To validate this hypothesis, we introduce an extra position-stable sink token explicitly designed for Diffusion Language Models. This is implemented by prepending the sink token to the beginning of the sequence, utilizing a modified attention mask where the sink token is constrained to attend solely to itself, while remaining globally visible to all other tokens. We demonstrate that this simple architectural modification yields substantial performance improvements across DLMs with various initialization strategies. By providing a stable, dedicated anchor, our method effectively mitigates the instability caused by moving sinks, leading to consistently better generation quality regardless of the underlying model configuration.

Motivated by these significant performance

gains, we further explore the mechanism behind this extra sink token through extensive experiments. We first observe that when this low-information token is introduced during pretraining, the model spontaneously learns to offload a significant portion of its attention mass to it. Moreover, our analysis reveals that this effectiveness is position-invariant. Whether the sink token is placed at the beginning or the end of the sequence, the model achieves comparable improvements, confirming that the benefit stems from the token’s functional role as a fixed attention sink rather than its specific position in the sequence.

Our main contributions are as follows:

- We provide a systematic analysis of *moving sinks* in DLMs from the *value-space* perspective, revealing that sink tokens consistently act as an implicit mechanism to mitigate excessive information mixing.

- Based on our hypothesis, we propose a simple yet effective *extra sink token* which not only taming moving sink issue in DLMs but also improves DLMs’ performance. Extensive experiments are conducted to prove the effectiveness of our approach.

2 Related Work

2.1 Diffusion Language Model

Diffusion models have achieved substantial progress and attracted extensive exploration in the image domain (Rombach et al., 2022; Ci et al., 2025). In addition, Diffusion Language Models have been studied from both continuous and discrete perspectives. Gulrajani and Hashimoto (2023) analyzed scaling laws for continuous diffusion language models and found that, to achieve favorable compute efficiency, continuous diffusion models typically require substantially longer training than autoregressive counterparts. Building on this line of work, recent efforts have scaled DLMs to billions of parameters. Nie et al. (2025a) demonstrated that 1.1B-scale DLMs are effective for downstream language tasks such as question answering. Rather than training from scratch, Gong et al. (2025) proposed converting pretrained autoregressive language models into DLMs (DiffuGPT and DiffuLLaMA). In parallel, Nie et al. (2025b) introduced LLaDA, an 8B-parameter DLM trained from scratch, achieving performance competitive with LLaMA3-8B (Team, 2024).

Beyond general language modeling, DLMs have been extended to other settings. Mercury Coder (Khanna et al., 2025) showed practical applicability in code generation. Zhang et al. (2025) explored DLMs for text embedding, and recent studies further investigated post-training techniques for DLMs (Zhao et al., 2025; Lin et al., 2025; Zhu et al., 2025). DLMs have also been extended to multimodal modeling (Yang et al., 2025; Xin et al., 2025; You et al., 2025).

2.2 Attention Sink

Attention sinks in autoregressive Large Language Models: Attention sinks refer to tokens that attract a disproportionate amount of attention despite carrying limited semantic content (Xiao et al., 2023; Barbero et al., 2025). In decoder-only autoregressive LLMs, this often manifests as many heads allocating substantial attention to the first token (Barbero et al., 2025; Gu et al., 2024). Be-

yond being an analysis artifact, sinks have been connected to several practical considerations, including streaming/sliding-window attention (Xiao et al., 2023), KV-cache efficiency and quantization robustness (Liu et al., 2024; Ge et al., 2024). Recent analyses suggest that concentrating attention on a low-information token can help mitigate excessive information mixing in deep and long-context Transformers, and can serve as an approximate no-op for some heads when strong updates are unnecessary (Barbero et al., 2025; Gu et al., 2024). Empirically, removing or weakening the sink pattern (e.g., dropping <bos> token at inference) can substantially degrade performance, especially in long-context settings (Barbero et al., 2025). Alternative mechanisms such as gated attention explicitly regulate information flow and can suppress sink behaviors, but introduce additional parameters and architectural complexity (Qiu et al., 2025).

Attention sinks in Diffusion Language Model:

Recent study (Rulli et al., 2025) shows that attention sinks also arise in Diffusion Language Models, but with distinct characteristics. Unlike the largely static sinks observed in autoregressive models, sinks in DLMs are often dynamic, shifting across denoising steps due to the bidirectional and iterative nature of diffusion generation. Moreover, sink positions in DLMs are not confined to the beginning of the sequence and frequently align with meaningless tokens such as mask token and white spaces token. Recent work has begun to exploit sink tokens as pivotal tokens to construct sparse attention patterns for accelerating inference process of DLMs (Wang et al., 2025; Song et al., 2025).

In this work, we analyze the characteristics of sink tokens in DLMs from the perspective of the Transformer value space. We find that, in DLMs, the moving sink token exhibits the same behavior observed in autoregressive models: it consistently attains the smallest global norm. This phenomenon can be interpreted as a protective mechanism learned during training, which prevents token-wise information from being excessively mixed. We further hypothesize that the position instability of moving sinks introduces additional inference-time instability, potentially constraining further gains in model capacity. Motivated by this hypothesis, we conduct experiments that replace the moving sink with a position-stable sink token, and we observe consistent improvements in performance.

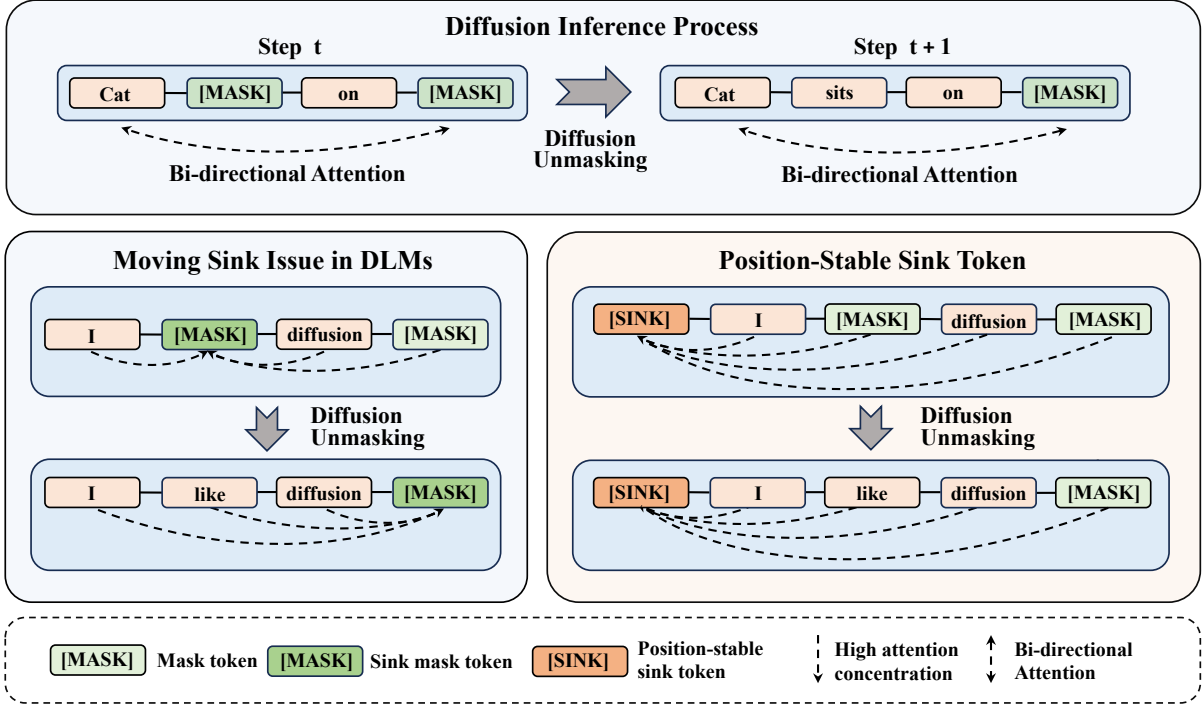


Figure 2: **Overview of the moving-sink phenomenon in diffusion language models and our stable sink token.** **Top:** Diffusion inference iteratively denoises a partially masked sequence; at each step, the model predicts masked positions with bi-directional attention. **Middle (left):** In vanilla DLMs, attention can concentrate on low-norm tokens that act as an implicit sink; since the masked set changes across steps, the sink position shifts over time (*moving sink*), which may increase inference instability. **Middle (right):** We prepend an extra [SINK] token, turning the moving sink into a stable sink. **Bottom:** Legend of the symbols used in the figure.

3 Method

As shown in Figure 2, during both training and inference of DLMs, we introduce an additional sink token to convert the unstable moving attention sink into a position-stable attention sink, thereby improving DLM performance.

3.1 Diffusion Language Modeling

Diffusion Language Models (DLMs) define a distribution over complete token sequences via (i) a forward masking corruption process and (ii) a learned reverse denoising process. Let $x_0 = (x_0^1, \dots, x_0^L)$ denote a clean sequence of length L , where x_0^i is the token at position i . We use a special mask token [MASK]. The forward process yields progressively corrupted variables $x_{1:T} = \{x_1, \dots, x_T\}$, and the reverse process reconstructs less corrupted sequences from more corrupted ones. Unlike autoregressive models, DLMs predict masked tokens in parallel at each denoising step.

3.1.1 Forward Process

Starting from x_0 , the forward process progressively corrupts the sequence by replacing tokens with

[MASK], producing $x_{1:T}$. It factorizes as

$$q(x_{1:T} | x_0) = \prod_{t=1}^T q(x_t | x_{t-1}), \quad (1)$$

where $t \in \{1, \dots, T\}$ indexes diffusion steps and T is the total number of steps. Each transition $q(x_t | x_{t-1})$ applies independent masking per position under a predefined noise schedule: at step t , a token remains *unmasked* with some probability, otherwise it is replaced by [MASK]. We set x_T to be fully corrupted (all positions are [MASK]), so generation can start from a fixed maximally corrupted sequence and then denoise.

3.1.2 Reverse Process

The reverse process iteratively denoises from x_T back to x_0 . We parameterize reverse transitions with a neural model p_θ that predicts tokens conditioned on the current corrupted sequence, i.e., $p_\theta(x_{t-1} | x_t)$. Operationally, p_θ performs *mask prediction*: it predicts token identities at masked positions conditioned on x_t , and it predicts all masked tokens in parallel. A common choice treats posi-

tions independently given x_t , yielding

$$p_\theta(x_0 | x_t) = \prod_{i=1}^L p_\theta(x_0^i | x_t). \quad (2)$$

During generation, these conditionals are used to fill masked tokens to obtain a less corrupted sequence, repeating over steps until a fully specified x_0 is reached.

3.1.3 Training Objective

DLMs are trained with a cross-entropy objective computed only on masked positions. Let $\{\tau_t\}_{t=1}^T$ denote a predefined noise schedule, where $\tau_t \in (0, 1]$ is the masking probability at diffusion step t . For a corrupted sequence x_t produced from x_0 , we define the masked-token log-likelihood as

$$\ell(x_0, x_t; \theta) = \frac{1}{\tau_t} \sum_{i=1}^L \mathbf{1}[x_t^i = [\text{MASK}]] \log p_\theta(x_0^i | x_t). \quad (3)$$

The indicator restricts the loss to masked positions. The overall training objective minimizes the negative expected masked-token log-likelihood:

$$\mathcal{L}(\theta) = -\mathbb{E}_{t, x_0, x_t} [\ell(x_0, x_t; \theta)]. \quad (4)$$

This objective learns a denoising function over partially observed sequences, enabling parallel token updates and avoiding strict left-to-right dependencies.

3.2 Extra Sink Token for DLM

As illustrated in Figure 1, moving sink behavior is widely observed in existing DLMs. In the absence of a fixed low-norm anchor, the attention sink tends to shift unpredictably across tokens. These temporary sink tokens typically contain low semantic information and exhibit a lower L_2 norm. The root cause lies in the softmax operation which necessitates that attention weights sum to one, mirroring the attention sink behavior observed in autoregressive models (Barbero et al., 2025). When a token lacks a strong semantic match within the context, the model is forced to assign redundant attention mass to globally visible tokens, often turning content tokens into implicit sinks.

However, this shifting behavior complicates the modeling of the diffusion process and hinders the improvement of DLMs. To address this issue and stabilize the attention mechanism, we introduce a

dedicated extra sink token to the input sequence. This token is designed to serve as a stable, low-information target to absorb excess attention. To strictly enforce its role as a pure sink and prevent it from aggregating semantic information, we apply a structured attention mask with the following constraints: (1) the sink token is restricted to attend only to itself; and (2) all other tokens in the sequence are allowed to attend to the sink token.

Formally, given an input sequence of latent representation $X \in \mathbb{R}^{L \times d_{\text{model}}}$, we prepend an extra sink token embedding $s \in \mathbb{R}^{d_{\text{model}}}$ to the sequence. The augmented input $\tilde{X} \in \mathbb{R}^{(L+1) \times d_{\text{model}}}$ is given by:

$$\tilde{X} = [s; X]. \quad (5)$$

Let $k = 0$ denote the index of the sink token s , we enforce the sink constraints by defining the attention mask bias M_{ij} as:

$$M_{ij} = \begin{cases} -\infty, & \text{if } i = k \text{ and } j \neq k \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

Under this formulation, the sink token is effectively isolated from aggregating sequence information, creating an asymmetric dependency where content tokens retain the ability to allocate attention mass to it. This constraint ensures that the sink token remains semantically neutral, functioning solely as a target to absorb excess attention weights. In our DLM framework, this configuration is applied consistently across all diffusion timesteps. Given that it introduces only a single additional token and utilizes standard masking, the computational overhead is negligible, providing an efficient and effective solution to regulate attention behavior and enhance model robustness.

4 Experiment

4.1 Experiment Settings

4.1.1 Benchmarks

To enable a thorough evaluation, we evaluate DLMs on widely adopted benchmarks spanning commonsense reasoning and reading comprehension: HellaSwag (Zellers et al., 2019), ARC-e (Clark et al., 2018), ARC-c (Clark et al., 2018), PIQA (Bisk et al., 2020), SIQA (Sap et al., 2019), RACE (Lai et al., 2017), and LAMBADA (Paperno et al., 2016). Following prior studies (Nie et al., 2025a), we also evaluate mathematical reasoning ability of DLMs on GSM8K (Cobbe et al., 2021) after the supervised fine-tune (SFT)

Model	Tokens	ARC-e	ARC-c	HellaSwag	PIQA	RACE	SIQA	LAMBADA	GSM8K
0.5B Models									
DLM	30B	44.70	27.13	35.21	55.28	31.29	40.28	44.09	46.92
DLM + extra token	30B	55.47	32.17	48.91	63.44	33.79	37.82	45.24	49.20
DLM + GA	30B	47.85	29.27	41.95	61.92	34.07	36.08	45.10	47.01
1.5B Models									
DLM	100B	65.49	41.47	59.50	66.49	37.80	36.18	66.58	57.77
DLM + extra token	100B	68.18	43.43	61.31	68.17	37.61	39.97	66.41	58.45
DLM + GA	100B	53.03	32.27	52.90	60.33	37.42	38.26	51.10	52.35

Table 1: Evaluation of our DLMs (autoregressive LLMs initialized). We report results for two model scales (0.5B and 1.5B), separated into two blocks for clarity. The “Tokens” column denotes the total number of training tokens used for each setting. “DLM + extra token” denotes the DLM augmented with an additional introduced sink token, while “DLM + GA” denotes the DLM equipped with the gated attention (GA) mechanism.

Model	Tokens	ARC-e	ARC-c	HellaSwag	PIQA	RACE	SIQA	LAMBADA	GSM8K
DLM	100B	32.20	25.09	31.56	54.08	31.77	35.98	29.26	34.95
DLM + extra token	100B	40.91	22.35	40.43	62.24	32.73	37.92	47.16	35.17
DLM + GA	100B	36.28	23.46	37.61	57.07	33.68	34.44	42.27	36.39

Table 2: Evaluation of DLMs (trained from scratch) on different benchmarks. The “Tokens” column denotes the total number of training tokens used for each setting. “DLM + extra token” denotes the DLM augmented with an additional introduced sink token, while “DLM + GA” denotes the DLM equipped with the gated attention (GA) mechanism.

process.

4.1.2 Implementation Details

We conduct experiments under two training settings: (1) training a DLM initialized from an autoregressive LLM, and (2) training a DLM from scratch. For the former, we initialize the DLM with Qwen2.5 Base model weights (Qwen et al., 2025) and consider two model scales, 0.5B¹ and 1.5B² parameters. These models are trained on the FineWeb dataset (Penedo et al., 2024). For the from-scratch setting, we adopt the same model architecture as the SMDM framework (Nie et al., 2025a). In this case, the model has 0.5B parameters and is trained on the SlimPajama dataset (Soboleva et al., 2023). During supervised fine-tuning (SFT), we fine-tune the DLM for 10 epochs on the augmented training data (Deng et al., 2023), with a context length of 2048. Additional training details are provided in Appendix A.1.

4.2 Experiment Results

Under the setting of training Diffusion Language Models from autoregressive LLMs, we compare three DLM modeling strategies at the 0.5B and

1.5B parameter scales: (1) the vanilla DLM, (2) DLM equipped with element-wise gated attention (Qiu et al., 2025), and (3) DLM with an additional sink token. The experimental results are reported in Table 1. Implementation details for gated attention are provided in Appendix A.2.

At the 0.5B scale, both gated attention and the introduction of an additional sink token lead to consistent performance improvements over the vanilla DLM, demonstrating the effectiveness of these modifications. However, at the 1.5B scale, the DLM with gated attention exhibits degraded performance compared to the vanilla baseline. This suggests that, when training DLMs from larger autoregressive LLMs, forcibly introducing additional parameters to implement attention gating may harm model stability. In contrast, our approach of adding an extra sink token continues to yield performance gains at the 1.5B scale, highlighting its robustness and effectiveness under larger model settings.

To further substantiate the effectiveness of our method, we extend our evaluation to training 0.5B-scale diffusion language models from scratch, maintaining the same baseline configurations as in the previous experiments. The results presented in Table 2 demonstrate that our approach yields consistent improvements across different training paradigms, highlighting its robustness independent

¹<https://huggingface.co/Qwen/Qwen2.5-0.5B>

²<https://huggingface.co/Qwen/Qwen2.5-1.5B>

Benchmark	Front	End
ARC-e	55.47	56.14
ARC-c	32.17	32.51
RACE	33.79	31.77
HellaSwag	48.91	48.16
PIQA	63.44	64.31

Table 3: Evaluation of different sink token positions.

of the initialization setting.

4.3 Ablation Studies

In this part, we conduct ablation studies to analyze the contributions of different components. For efficiency, we utilize the Qwen2.5-0.5B checkpoint as the initialization and train on a subset of 30B tokens from FineWeb. Based on this setup, we examine the impact of sink token placement and quantity, followed by an in-depth analysis of how these tokens influence attention allocation.

4.3.1 The position of sink token

Unlike autoregressive language models where the initial token naturally serves as an attention sink due to causal masking, diffusion language models utilize bidirectional attention with global context visibility. This distinction raises a critical question for DLMS: does the explicit sink token need to be placed at a specific position to function effectively, or is its efficacy position-agnostic? To investigate this, we conduct ablation studies by varying the placement of the sink token within the input sequence, comparing two distinct configurations: prepending it at the beginning versus appending it at the end.

Formally, for the configuration where the sink token is placed at the beginning of the sequence, given an input sequence $X \in \mathbb{R}^{n \times d_{\text{model}}}$, we prepend a sink token $s \in \mathbb{R}^{d_{\text{model}}}$ to obtain the final input sequence, as defined in Equation 5. For the configuration where the sink token is placed at the end of the sequence, given the same input sequence $X \in \mathbb{R}^{n \times d_{\text{model}}}$, we append the extra vector $s \in \mathbb{R}^{d_{\text{model}}}$ to form the final input, which can be expressed as:

$$\tilde{X} = [X; s]. \quad (7)$$

The sink token is added at the corresponding position during both training and inference.

The results, reported in Table 3, indicate stable improvements across both configurations. This confirms that our method is robust to positional

Benchmark	0	1	2	4
ARC-e	44.70	55.47	54.76	55.39
ARC-c	27.13	32.17	32.42	31.47
RACE	31.29	33.79	32.89	33.06
HellaSwag	35.21	48.91	48.94	48.45
PIQA	55.28	63.44	63.66	62.62

Table 4: Performance sensitivity to the number of reintroduced tokens on our approach.

Benchmark	Stable sink token	Zero-value token
ARC-e	55.47	55.68
ARC-c	32.17	32.94
RACE	33.79	32.87
HellaSwag	48.91	48.81
PIQA	63.44	64.09

Table 5: Evaluation for setting sink token to zero-vector in value space.


variations. It further validates that the sink token functions globally as an attention attractor due to its properties, rather than relying on the positional bias typically seen in causal models.

4.3.2 The number of sink token


Next, we investigate the impact of sink token quantity to determine if the mechanism’s effectiveness is constrained by the capacity of a single token. To this end, we conduct an ablation study by varying the number of added tokens, specifically comparing configurations with 1, 2, and 4 sink tokens. The results, summarized in Table 4, reveal that adding a single sink token yields the most substantial improvement, while further increasing the quantity leads to negligible marginal gains. This saturation phenomenon indicates that the sink token functions purely as a structural anchor for attention offloading rather than a carrier of semantic information. Consequently, a single token provides sufficient capacity to fulfill this role.

4.4 Internal Analysis of Sink Tokens and Attention Allocation


To validate the effectiveness of our proposed theory beyond standard benchmarks, we conduct a statistical analysis of internal token representations and attention behaviors. Specifically, we examine DLMS initialized from Qwen2.5-0.5B and Qwen2.5-1.5B. Following supervised fine-tuning, we sample token representations in the value space and their corresponding attention maps across Transformer layers during inference on the GSM8K dataset. For each model, we sample 100K inference steps and




(a) Mean L_2 Norm per Layer: Sink vs Others (Model: 0.5B)



(b) Sink Token Attention Received per Layer (Model: 0.5B)



(c) Mean L_2 Norm per Layer: Sink vs Others (Model: 1.5B)



(d) Sink Token Attention Received per Layer (Model: 1.5B)

Figure 3: Sink token analysis across transformer layers for DLMs initialized from Qwen2.5-0.5B and Qwen2.5-1.5B. (a) and (c): mean value-space L_2 norm per transformer layer, comparing the extra sink token against all other tokens. (b) and (d): the proportion of attention mass received by the sink token per layer. After the initial few layers and before the final layers, across the intermediate transformer layers that are most closely related to information integration and transformation, the sink token exhibits substantially smaller value-space L_2 norms than other tokens while receiving a large share of global attention.

quantify (i) the L_2 norm statistics of token representations in the value space and (ii) the attention mass allocated to sink tokens.

As illustrated in Figure 3, for each layer, we measure the proportion of global attention directed toward the introduced sink token and analyze the distribution of token L_2 norms. The results reveal that the sink token maintains a significantly lower L_2 norm compared to the average of other tokens while absorbing redundant attention. This indicates that the model implicitly learns to utilize a token with minimal magnitude to offload redundant attention, thereby minimizing interference with the semantic information in the residual stream.

Based on the finding that DLMs benefit from directing attention to tokens with minimal norms,

we conduct a further validation study where we explicitly force the sink token’s value states to be zero vectors at every Transformer layer under the setting of training DLM from qwen2.5-0.5B. The results are reported in Table 5. We find that even when the sink token is set to a zero vector and contains no semantic information, it still yields consistent gains for the DLM. This confirms our hypothesis that the presence of a low-norm state is the critical factor for mitigating excessive information mixing.

5 Conclusion

In this work, we identify the critical role of attention sinks in DLMs to alleviate information over-mixing, while highlighting the positional shifts arising from their distinct attention mechanisms dis-

tinct from auto-regressive models. To resolve this, we introduce a position-stable sink token to anchor attention, yielding consistent performance gains across diverse settings. Beyond these improvements, we hope our work offers valuable insights into DLM attention mechanics.

Limitations

We conducted extensive experiments on Diffusion Language Models at the 0.5B and 1.5B parameter scales. However, we did not scale our experiments to larger Diffusion Language Models.

Ethics Statement

We trained and evaluated our model using publicly available datasets that have been verified and widely used. During the preparation of this manuscript, we used generative AI tools for grammar and language proofreading.

References

- Federico Barbero, Alvaro Arroyo, Xiangming Gu, Christos Perivolaropoulos, Michael Bronstein, Petar Veličković, and Razvan Pascanu. 2025. Why do llms attend to the first token? *arXiv 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 *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439.
- 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 arXiv:2107.03374*.
- En Ci, Shanyan Guan, Yanhao Ge, Yilin Zhang, Wei Li, Zhenyu Zhang, Jian Yang, and Ying Tai. 2025. Describe, don't dictate: Semantic image editing with natural language intent. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19185–19194.
- 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. *arXiv 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. *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*.
- Yuntian Deng, Kiran Prasad, Roland Fernandez, Paul Smolensky, Vishrav Chaudhary, and Stuart Shieber. 2023. Implicit chain of thought reasoning via knowledge distillation. *arXiv preprint arXiv:2311.01460*.
- Suyu Ge, Yunan Zhang, Liyuan Liu, Minjia Zhang, Jiawei Han, and Jianfeng Gao. 2024. Model tells you what to discard: Adaptive kv cache compression for llms. In *International Conference on Learning Representations (ICLR)*.
- Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An, Peilin Zhao, Wei Bi, Jiawei Han, Peng Hao, and Kong Lingpeng. 2025. Scaling diffusion language models via adaptation from autoregressive models. In *The Thirteenth 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*.
- Ishaan Gulrajani and Tatsunori B Hashimoto. 2023. Likelihood-based diffusion language models. *Advances in Neural Information Processing Systems*, 36:16693–16715.
- 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 preprint arXiv:2506.17298*.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. Race: Large-scale reading comprehension dataset from examinations. *arXiv preprint arXiv:1704.04683*.
- Nianyi Lin, Jiajie Zhang, Lei Hou, and Juanzi Li. 2025. Boundary-guided policy optimization for memory-efficient rl of diffusion large language models. *arXiv preprint arXiv:2510.11683*.
- Ruikang Liu, Haoli Bai, Haokun Lin, Yuening Li, Han Gao, Zhengzhuo Xu, Lu Hou, Jun Yao, and Chun Yuan. 2024. IntactKV: Improving large language model quantization by keeping pivot tokens intact. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 7716–7741, Bangkok, Thailand. Association for Computational Linguistics.

- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Shen Nie, Fengqi Zhu, Chao Du, Tianyu Pang, Qian Liu, Guangtao Zeng, Min Lin, and Chongxuan Li. 2025a. Scaling up masked diffusion models on text. In *The Thirteenth International Conference on Learning Representations*.
- Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. 2025b. Large language diffusion models. *arXiv preprint arXiv:2502.09992*.
- Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc-Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. 2016. The lambada dataset: Word prediction requiring a broad discourse context. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 1525–1534.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben Allal, Anton Lozhkov, Margaret Mitchell, Colin A. Raffel, Leandro von Werra, and Thomas Wolf. 2024. The fineweb datasets: Decanting the web for the finest text data at scale. *Advances in Neural Information Processing Systems*, 37:30811–30849.
- Zihan Qiu, Zekun Wang, Bo Zheng, Zeyu Huang, Kaiyue Wen, Songlin Yang, Rui Men, Le Yu, Fei Huang, Suozhi Huang, Liu Dayiheng, Zhou Jingren, and Lin Junyang. 2025. Gated attention for large language models: Non-linearity, sparsity, and attention-sink-free. *arXiv preprint arXiv:2505.06708*.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, and 24 others. 2025. Qwen2.5 Technical Report. [10.48550/arXiv.2412.15115](https://arxiv.org/abs/2412.15115).
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.
- Maximo Eduardo Rulli, Simone Petrucci, Edoardo Michielon, Fabrizio Silvestri, Simone Scardapane, and Alessio Devoto. 2025. Attention sinks in diffusion language models. *arXiv preprint arXiv:2510.15731*.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. Socialliqa: Commonsense reasoning about social interactions. *arXiv preprint arXiv:1904.09728*.
- Daria Soboleva, Faisal Al-Khateeb, Robert Myers, Jacob R Steeves, Joel Hestness, and Nolan Dey. 2023. SlimPajama: A 627B token cleaned and deduplicated version of RedPajama.
- Yuerong Song, Xiaoran Liu, Ruixiao Li, Zhigeng Liu, Zengfeng Huang, Qipeng Guo, Ziwei He, and Xipeng Qiu. 2025. Sparse-dllm: Accelerating diffusion llms with dynamic cache eviction. *arXiv preprint arXiv:2508.02558*.
- Llama Team. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Zeqing Wang, Gongfan Fang, Xinyin Ma, Xingyi Yang, and Xinchao Wang. 2025. Sparsed: Sparse attention for diffusion language models. *arXiv preprint arXiv:2509.24014*.
- Chengyue Wu, Hao Zhang, Shuchen Xue, Zhijian Liu, Shizhe Diao, Ligeng Zhu, Ping Luo, Song Han, and Enze Xie. 2025. 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. *arXiv preprint arXiv:2309.17453*.
- Yi Xin, Qi Qin, Siqi Luo, Kaiwen Zhu, Juncheng Yan, Yan Tai, Jiayi Lei, Yuewen Cao, Keqi Wang, Yibin Wang, Jinbin Bai, Qian Yu, Dengyang Jiang, Yuandong Pu, Haoxing Chen, Le Zhuo, Junjun He, Gen Luo, Tianbin Li, and 13 others. 2025. Lumina-dimoo: An omni diffusion large language model for multimodal generation and understanding. *arXiv preprint arXiv:2510.06308*.
- Ling Yang, Ye Tian, Bowen Li, Xincheng Zhang, Ke Shen, Yunhai Tong, and Mengdi Wang. 2025. 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 preprint arXiv:2508.15487*.
- Zebin You, Shen Nie, Xiaolu Zhang, Jun Hu, Jun Zhou, Zhiwu Lu, Ji-Rong Wen, and Chongxuan Li. 2025. Llada-v: Large language diffusion models with visual instruction tuning. *arXiv preprint arXiv:2505.16933*.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*.
- Siyue Zhang, Yilun Zhao, Liyuan Geng, Arman Cohan, Anh Tuan Luu, and Chen Zhao. 2025. Diffusion vs. autoregressive language models: A text embedding perspective. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 4273–4303, Suzhou, China. Association for Computational Linguistics.

Siyao Zhao, Devaansh Gupta, Qinqing Zheng, and Aditya Grover. 2025. d1: Scaling reasoning in diffusion large language models via reinforcement learning. *arXiv preprint arXiv:2504.12216*.

Fengqi Zhu, Rongzhen Wang, Shen Nie, Xiaolu Zhang, Chunwei Wu, Jun Hu, Jun Zhou, Jianfei Chen, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. 2025. Llada 1.5: Variance-reduced preference optimization for large language diffusion models. *arXiv preprint arXiv:2505.19223*.

A Appendix

A.1 Training details

Training DLM from autoregressive LLM: We optimize the model using AdamW (Loshchilov and Hutter, 2017), with $\beta_1 = 0.9$, $\beta_2 = 0.95$, and a weight decay of 0.1. We adopt a cosine learning-rate schedule, with the peak learning rate set to 1×10^{-4} and the minimum learning rate set to 1×10^{-5} , and use linear warmup for the first 1% of the training tokens. For the 0.5B model, we train on 30B tokens from the FineWeb (Penedo et al., 2024) dataset with a batch size of 512. For the 1.5B model, we train on 100B tokens from the FineWeb dataset with a batch size of 4096.

Training DLM from scratch: To be consistent with the SMDM (Nie et al., 2025a), we utilize the AdamW optimizer (Loshchilov and Hutter, 2017), setting $\beta_1 = 0.9$, $\beta_2 = 0.95$, and a weight decay of 0.1. Additionally, we apply a cosine learning rate schedule with a maximum learning rate of 2×10^{-4} and a minimum learning rate of 2×10^{-5} with 1% of the tokens for linear warmup. We train the model on 100B tokens from the SlimPajama dataset (Soboleva et al., 2023), using a batch size of 256.

A.2 Gated Attention for DLM

We augment the standard attention layer in DLMs with a gating mechanism applied after the scaled dot-product attention (SDPA). The overall Transformer architecture follows the same design as the Qwen model, and the gated attention layer is used as a drop-in replacement for the vanilla attention layer at every diffusion step.

Given the SDPA output

$$\begin{aligned} Y &= \text{Attention}(Q, K, V) \\ &= \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V. \end{aligned} \quad (8)$$

we apply an element-wise, input-dependent gate to modulate the attention output:

Dataset	Zero-shot		Many-shot	
	origin	w/ sink	origin	w/ sink
ARC-e	0.4470	0.5547	0.6279	0.6132
PIQA	0.5528	0.6344	0.6839	0.6915
LAMBADA	0.4409	0.4524	0.4578	0.4694

Table 6: In-context learning comparison between the original DLM (origin) and the DLM with a sink token (w/ sink). In the Many-shot setting, ARC-e uses 96 shots, while LAMBADA and PIQA use 32 shots.

$$Y' = g(Y, X) = Y \odot \sigma(XW_g). \quad (9)$$

where $X \in \mathbb{R}^{n \times d_{\text{model}}}$ is the input hidden representation to the attention layer, $W_g \in \mathbb{R}^{d_{\text{model}} \times d_k}$ is a learnable gating projection, $\sigma(\cdot)$ denotes the sigmoid function, and \odot represents element-wise multiplication.

The gated attention output is then passed to the output projection layer:

$$O = Y'W_O, \quad (10)$$

where $W_O \in \mathbb{R}^{d_k \times d_{\text{model}}}$.

The gating function introduces non-linearity into the attention layer by dynamically scaling the SDPA output based on the current hidden states. Since the gate is applied after attention weight computation, it preserves the original attention score normalization while enabling adaptive suppression or amplification of attended features. In the context of DLMs, the same gated attention layer is applied across diffusion timesteps, allowing the model to regulate information flow under varying noise conditions with minimal additional computational overhead.

A.3 Value-Space Token Vectors Analysis

A.3.1 Value-space token vectors in a Transformer

Consider a Transformer layer ℓ with input hidden states $\mathbf{H}^{(\ell)} = [\mathbf{h}_1^{(\ell)}, \dots, \mathbf{h}_n^{(\ell)}]^\top \in \mathbb{R}^{n \times d}$, where n is the sequence length and d is the hidden size. For each attention head $m \in \{1, \dots, M\}$, the layer computes query, key, and value projections:

$$\mathbf{Q}_m = \mathbf{H}^{(\ell)} \mathbf{W}_m^Q, \quad (11)$$

$$\mathbf{K}_m = \mathbf{H}^{(\ell)} \mathbf{W}_m^K, \quad (12)$$

$$\mathbf{V}_m = \mathbf{H}^{(\ell)} \mathbf{W}_m^V, \quad (13)$$

where $\mathbf{W}_m^Q, \mathbf{W}_m^K, \mathbf{W}_m^V \in \mathbb{R}^{d \times d_h}$ and $d_h = d/M$. The *value-space token vector* of token j (in head m) is the j -th row of \mathbf{V}_m :

$$\mathbf{v}_j^{(\ell, m)} = \mathbf{h}_j^{(\ell)} \mathbf{W}_m^V \in \mathbb{R}^{d_h}. \quad (14)$$

Intuitively, $\mathbf{v}_j^{(\ell, m)}$ is the content that token j contributes to other tokens through attention mixing.

A.3.2 L_2 norm of a token in value space

For a head-specific value vector, the L_2 norm is computed as

$$\left\| \mathbf{v}_j^{(\ell, m)} \right\|_2 = \sqrt{\sum_{k=1}^{d_h} \left(\mathbf{v}_{j,k}^{(\ell, m)} \right)^2}. \quad (15)$$

To obtain a single magnitude per token across heads, we average per-head norms:

$$\bar{r}_j^{(\ell)} = \frac{1}{M} \sum_{m=1}^M \left\| \mathbf{v}_j^{(\ell, m)} \right\|_2. \quad (16)$$

A.3.3 Why attending to low-norm tokens approximates a “no-op”

In head m , the attention output for query token i is a convex combination of value vectors:

$$\mathbf{o}_i^{(\ell, m)} = \sum_{j=1}^n \alpha_{ij}^{(\ell, m)} \mathbf{v}_j^{(\ell, m)}, \quad (17)$$

where the attention weights satisfy

$$\sum_{j=1}^n \alpha_{ij}^{(\ell, m)} = 1, \quad (18)$$

$$\alpha_{ij}^{(\ell, m)} \geq 0. \quad (19)$$

By the triangle inequality, the output norm is upper bounded by the weighted sum of value norms:

$$\left\| \mathbf{o}_i^{(\ell, m)} \right\|_2 \leq \sum_{j=1}^n \alpha_{ij}^{(\ell, m)} \left\| \mathbf{v}_j^{(\ell, m)} \right\|_2. \quad (20)$$

Therefore, if a substantial portion of attention mass is assigned to tokens with very small $\left\| \mathbf{v}_j^{(\ell, m)} \right\|_2$, their contribution to $\mathbf{o}_i^{(\ell, m)}$ is correspondingly small.

After concatenating heads and applying the output projection, the multi-head attention update can be written as

$$\mathbf{z}_i^{(\ell)} = \mathbf{W}^O \left[\mathbf{o}_i^{(\ell, 1)}; \dots; \mathbf{o}_i^{(\ell, M)} \right], \quad (21)$$

and the residual connection updates the hidden state by

$$\mathbf{h}_i^{(\ell+1)} = \mathbf{h}_i^{(\ell)} + \mathbf{z}_i^{(\ell)}. \quad (22)$$

When attending mainly to low-norm value vectors, $\mathbf{z}_i^{(\ell)}$ becomes small in magnitude, making $\mathbf{h}_i^{(\ell+1)} \approx \mathbf{h}_i^{(\ell)}$. In this sense, focusing attention on low-norm tokens implements an approximate *no-op* update: the model can allocate attention probability mass while minimally perturbing representations, thereby reducing unnecessary information mixing among content-bearing tokens.

A.4 In-context Learning Evaluation

Recent work (Dai et al., 2026) investigated how the sink mechanism in dLLMs affects the model’s in-context learning ability. We choose a 0.5B DLM initialized from Qwen2.5 as the subject of our study. The experimental results, shown in Table 6, indicate that the position of the sink token indeed affects the model’s in-context learning performance. In the many-shot setting, the stable-sink DLM shows a smaller improvement than the moving-sink DLM. These findings support the conclusions of (Dai et al., 2026).