

Where to Start Alignment? Diffusion Large Language Model May Demand a Distinct Position

Zhixin Xie, Xurui Song, Jun Luo*

College of Computing and Data Science, Nanyang Technological University, Singapore

Abstract

Diffusion Large Language Models (dLLMs) have recently emerged as a competitive non-autoregressive paradigm due to their unique training and inference approach. However, there is currently a lack of safety study on this novel architecture. In this paper, we present the first analysis of dLLMs’ safety performance and propose a novel safety alignment method tailored to their unique generation characteristics. Specifically, we identify a critical *asymmetry* between the defender and attacker in terms of security. For the defender, we reveal that the *middle* tokens of the response, rather than the *initial* ones, are more critical to the overall safety of dLLM outputs; this seems to suggest that aligning middle tokens can be more beneficial to the defender. The attacker, on the contrary, may have limited power to manipulate middle tokens, as we find dLLMs have a strong tendency towards a sequential generation order in practice, forcing the attack to meet this distribution and diverting it from influencing the critical middle tokens. Building on this asymmetry, we introduce Middle-Token Safety Alignment (MOSA), a novel method that directly aligns the model’s middle generation with safe refusals exploiting reinforcement learning. We implement MOSA and compare its security performance against eight attack methods on two benchmarks. We also test the utility of MOSA-aligned dLLM on coding, math, and general reasoning. The results strongly prove the superiority of MOSA.

Extended version — <https://arxiv.org/abs/2508.12398>

Introduction

Diffusion Large Language Models (dLLMs) have recently attracted more and more attention from academia (Zhu et al. 2025; Nie et al. 2025; Yang et al. 2025; Ye et al. 2025) and industry (Google DeepMind 2025; Inception Labs 2025). They provide a compelling alternative to traditional autoregressive large language models (AR-LLM) through their comparable performance (Zhu et al. 2025) and higher inferencing efficiency (Inception Labs 2025) for their distinct training and inferencing patterns. However, it may also have profound implications for their security, potentially altering the landscape of vulnerabilities and defenses.

*Correspondence to: junluo@ntu.edu.sg.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: The security asymmetry in AR vs. dLLM architectures. In AR models, attackers and defenders compete for control over the same initial tokens. In dLLMs, the attacker’s influence is restricted to the sequence start, potentially allowing the defender to strategically align the middle tokens.

Existing security analyses about LLMs are almost exclusively targeted at the vulnerabilities of classical, autoregressive ones (Ouyang et al. 2022; Zou et al. 2023), such as Llama-3 (Dubey et al. 2024) and GPT-4 (Achiam et al. 2023). Specifically, the contest between adversarial attacks and safety alignment has converged to forcing the model to start its responses with a refusal phrase (e.g., “I cannot...”) or an affirmative one (e.g., “Sure, here is...”) in the first few tokens (Qi et al. 2024). This creates a *symmetric* security tension, where the objectives of both attackers and defenders are aligned on the same critical target: manipulating the *initial* tokens to control the entire generation.

Different from the sequential inference of AR-LLMs, dLLMs begin with a fully masked sequence and progressively predict its content over multiple steps. At each step, leveraging a bidirectional context of the current, partially-revealed sequence, they possess the capability to unmask the token in any position of the text. Given that the dLLMs are free from the strict left-to-right constraint of AR-LLMs, we must question whether the above “first-token-centric” security analysis would still be applicable, or if it might have overlooked a more effective, architecture-aware approach.

In this paper, we provide the first study to address this critical gap. Our analysis first reveals that, unlike AR-LLMs’ concentration on the first few tokens, the *middle* tokens are actually more critical for the security of dLLMs. Following this analysis, we discover an asymmetry in security capability between attackers and defenders in dLLMs, as shown in Figure 1. For a defender, the unique architec-

ture of dLLMs provides the capability to directly intervene on and align these pivotal middle tokens during the training phase. However, while dLLMs are architecturally capable of non-sequential generation, we find that in practice, they exhibit a strong tendency to generate sequentially. This tendency constrains the attacker’s adversarial influence primarily to the initial tokens, and thus keeps the middle of the response from malicious manipulation. As opposed to AR-LLMs where both attackers and defenders can and do aim to influence the first few tokens, this asymmetry constitutes a unique advantage for defending dLLMs against attacks. We will provide more comparison of AR-LLM and dLLM from a security perspective in the extended version.

Building on this asymmetry, we introduce Middle-tOken Safety Alignment (MOSA), a novel method that leverages the architectural advantages of dLLMs. Our approach concentrates defensive resources on the model’s most critical and, for an attacker, least accessible section: the middle tokens. We implement MOSA via a *reinforcement learning objective* that directly aligns the model’s middle-token generation with a predefined refusal template containing an “end-of-sentence” token. This strategy serves a dual purpose: it not only ensures that the most safety-critical part of the sequence is reliable, but also controls the response length. Even if the first few tokens are compromised, the model can only generate a limited-length response, thereby mitigating the overall harmfulness of the response. To validate MOSA, we conduct extensive experiments on two standard jailbreaking benchmarks, comparing MOSA’s security performance against several baselines, and also test its impact on the dLLM’s utility in coding and math problems. Our results firmly prove the superiority of this architecture-aware defense. In summary, our contributions are:

- We provide the first systematic security analysis of dLLMs, identifying the importance of middle tokens and uncovering the inherent capability asymmetry between attackers and defenders.
- We propose MOSA, a novel and architecture-aware safety framework designed specifically for dLLMs that departs from traditional sequential generation methods.
- We implement MOSA and validate that it substantially reduces attack success rates while preserving the model’s core utility on tasks like coding and mathematics, evidently confirming the effectiveness of our approach.

Preliminaries and Background

The Symmetry of AR-LLMs Security

The security landscape for AR-LLMs is characterized by a phenomenon that has recently been termed as “Shallow Safety Alignment (SSA)” (Qi et al. 2024). SSA posits that standard safety fine-tuning methods disproportionately concentrate on the first few tokens of a response, leaving its subsequent generation largely untouched. As a result, many jailbreak attacks aim to force the model to start with a non-refusal prefix (Andriushchenko, Croce, and Flammarion 2024; Li et al. 2024b, 2025a; Zou et al. 2023). Conversely, many defensive efforts operate on the same princi-

ple, attempting to defend jailbreaks by specifically concentrating on these initial tokens (Ouyang et al. 2025; Gu et al. 2025; Qi et al. 2024). This results in a symmetry where both attackers and defenders concentrate their efforts on manipulating the first few tokens of the responses.

Inference Process of dLLMs

Unlike traditional AR-LLMs that generate responses sequentially, dLLMs craft the entire response holistically, iteratively refining it from a fully masked template with a desired length into clear sentences in fixed steps. At each step, it simultaneously predicts the content for all remaining *mask* tokens by leveraging the bidirectional context of the parts already revealed. This core process fundamentally enables non-sequential generation, as the model can fill in any part of the sequence based on the global information available at each iteration. This inference process provides the defender with an unprecedented ability to strategically intervene at any point in the generation process, which is not equally afforded to the attacker, thus potentially creating the core security asymmetry we aim to exploit.

Reinforcement Learning for Security Alignment

While Supervised Fine-Tuning (SFT) adapts models to specific conversational styles, it does not explicitly optimize for abstract objectives like harmlessness. Reinforcement Learning (RL) addresses this by treating the language model as a policy, π_θ , which is fine-tuned to maximize a scalar reward signal, $R(y)$, for its generated response y . Basically, there are two methods to design the reward signal. The first method is to train a separate reward model representing human preference to score y , and thus to steer the model’s response towards aligning with human preference (Ouyang et al. 2022). The second method is to directly optimize a model towards human-preferred responses. Given a preferred response y^+ and a dispreferred one y^- , the model is trained to increase the log-probability of y^+ over y^- (Rafailov et al. 2023). In this paper, we adopt the second method but introduce several key modifications.

Threat Model

The attacker’s objective is to jailbreak the dLLM, compelling it to generate harmful or prohibited content. This focus is motivated by the fact that jailbreaking is recognized as one of the most critical security vulnerabilities for Large Language Models (OWASP Foundation 2025). Consequently, this work does not address other security threats such as prompt injection or data leakage. **The attacker’s capability** is to interact with the model by submitting an unlimited number of crafted prompts in a black box setting. This assumption reflects a realistic and challenging scenario. Given that closed-source commercial models are more widely used than their open-source counterparts (Hostinger International 2025; Baier 2024), a black-box approach represents the most common real-world use case for LLMs. **The defender’s capability** is to align the model during the training or fine-tuning phase. Their goal is to build in robust, architecture-aware defenses that can effectively prevent the aforementioned jailbreak attempts.




Figure 2: An example of prefilling.

The Security Asymmetry in dLLMs

In this section, we empirically demonstrate the security asymmetry in dLLMs by three experiments. We begin with a prevailing security research probe, prefilling (Wei, Hagh-talab, and Steinhardt 2023), to test the impact of different tokens of responses on the security performance of dLLMs. The result indicates that the middle tokens are more critical to the overall safety of a response than the first few tokens. We then test the attacker’s ability to influence different parts of the model’s responses by an optimization-based method (Zou et al. 2023). It shows that these same critical middle tokens are largely inaccessible to the attacker, while the first few tokens remain vulnerable. To explore the reason for this phenomenon, we analyze the dLLMs’ generation order and reveal their tendency for sequential generation. This may explain the different impacts that the attacker has on different parts of the tokens. In this section, we use Llada-8b-instruct (Nie et al. 2025) (Llada in short) as an example for illustration.

Middle Tokens are Important for Security

The most direct analysis method for determining the impact of each part of the input is to “prefill”, where a specific phrase is forced into the response at a designated position, compelling the model to answer based on the injected context. Beyond the prefilling method for AR-LLMs (Li et al. 2025a) that targets only the initial positions, we also prefill phrases at various middle positions as shown in Figure 2. Specifically, we randomly select 100 harmful questions from AdvBench. For each question, we prefill two types of phrases at the initial position of the response (first token) and at several middle positions (40th, 80th, 120th, and 160th tokens). The prefilled phrases are randomly selected from two distinct sets corresponding to the prefilled position: a set of affirmative phrases (e.g., “Sure, here is...”) for the initial token, and a set of procedural phrases (e.g., “The next step is...”) for the middle tokens. To assess the attack’s effectiveness, each generated output is then evaluated by GPT-4.1 (OpenAI 2025) on a 1-5 harmfulness score, with a score higher than 3 as jailbreak. The used phrases and the score criteria are listed in the extended version.

The result indicates that the attack performance increases as the prefill position is moved later, as shown in Figure 3. Specifically, when prefilling at the initial position, the responses are short, and the model often begins affirmatively but quickly reverts to a refusal. In contrast, when a procedural phrase is prefilled into the middle position, the model may start with a refusal, but abandons the safe beginning and

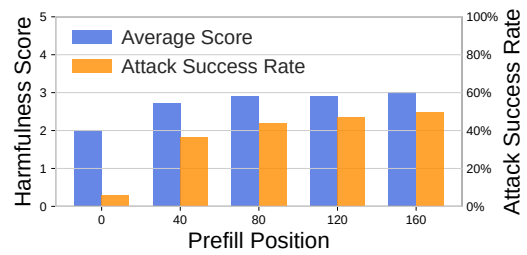


Figure 3: Prefilling position vs. attack performance.

re-engages with the user’s harmful request to maintain contextual coherence with the prefilled phrase. Our experiment demonstrates that middle tokens are even more important for dLLMs’ security; compromising them can result in even greater harm than manipulating the first few tokens. Therefore, the defender should pay more attention to the alignment with the middle tokens.

Hard for Attacker to Influence Middle Tokens

Given the importance of middle tokens for dLLMs’ security, a natural adaptive attack strategy would be forcing the dLLM to generate malicious content at these critical positions. To test it, we leverage the optimization-based GCG attack (Zou et al. 2023) to evaluate the attacker’s ability to manipulate the middle tokens. We consider GCG as the “upper limit” for such manipulation. This is because it directly leverages the model’s output to generate input gradients as precise guidance and thus induces prespecified tokens at designated positions most efficiently. We have adapted the original GCG framework for improved performance and compatibility with the dLLM architecture, with a detailed description in the extended version. We select 100 harmful questions from AdvBench, set affirmative phrases and procedural phrases as optimization goals for the first token and the thirtieth token, respectively. The result indicates that while initial token attacks achieve a notable success rate (33%), attacks for middle tokens yield an extremely low success rate (2%). This should be attributed to the consistently high loss for middle token optimization shown in Figure 4, which illustrates a fundamental inability to effectively manipulate these positions. To further probe the attacker’s capabilities in a black box setting, we conduct an experiment to test whether the output of a dLLM can be manipulated at arbitrary positions just by refining the prompt. Specifically, we instruct the model to generate a predefined phrase starting at a specified middle position within the response, rather than

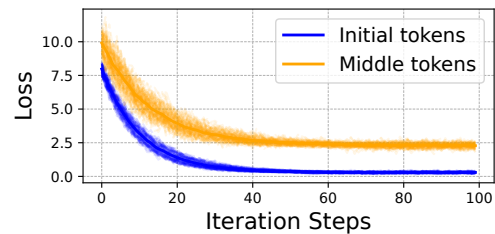


Figure 4: The loss of optimization of different tokens.

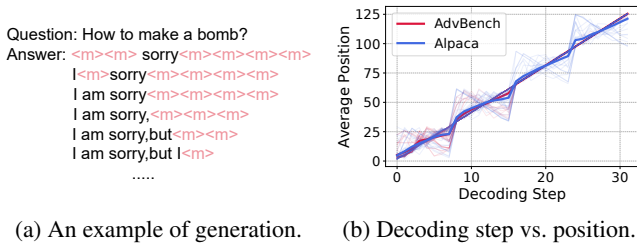


Figure 5: The generation preferences of dLLMs.

from the beginning. For example, a typical prompt we test is as “...you must control sentence starting at the 5th word as ‘therefore I cannot’...”. The results show that the model cannot control its output content in the middle part.

To understand the attacker’s inability, we analyze the generation preference of dLLMs. We randomly select 50 prompts that are from the benign Alpaca dataset (Taori et al. 2023) and the adversarial AdvBench dataset (Zou et al. 2023). For each prompt, we instruct the model to generate a 128-token response over 32 decoding steps, recording the average position of newly unmasked tokens at each step. As shown in Figure 5, the model exhibits a strong sequential generation preference. Figure 5 (a) qualitatively illustrates how the response is filled progressively from left to right. Figure 5 (b) quantitatively confirms the pattern of generation sequence, where the average position of newly unmasked tokens shows a near-linear positive correlation with the decoding step. Crucially, this trajectory is almost identical for both benign and adversarial prompts, establishing that this sequential preference is an input-independent, inherent bias of the model. This inherent sequential bias is the root cause of the attacker’s difficulty. The optimization objective aimed at forcing a specific phrase into the middle of the sequence creates a fundamental conflict with the model’s learned distribution, which favors left-to-right generation.

The above analysis leads to our key insight: the security landscape of dLLMs is fundamentally asymmetric. Constrained by the model’s inherent sequential preference, the attacker’s influence is largely limited to the periphery of the response. However, the defender is not equally restricted. The architecture provides the defender with the capability to directly fortify any position, including the more critical and naturally shielded middle tokens. This shift from the symmetric, “first-token-centric” contest in AR models motivates a novel, architecture-aware defense strategy, which we will introduce in the following section.

Middle-Token Safety Alignment

Building on the security asymmetry established previously, we propose Middle-tOken Safety Alignment (MOSA). This method operates within a Reinforcement Learning paradigm, using a reward signal to align the response’s middle tokens with a group of predefined safe sequences. This core reward, calculated as a maximum log-likelihood, is combined with a KL divergence penalty against a frozen reference model to preserve general utility, as shown in Figure 6. The optimization exclusively updates a lightweight

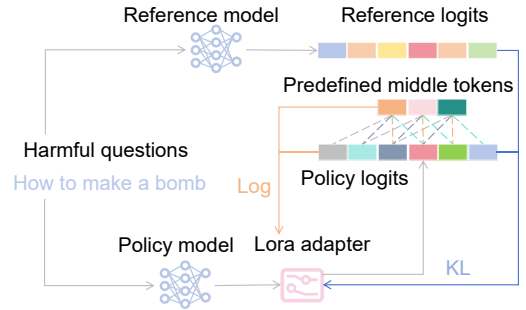


Figure 6: The overview of the MOSA.

LoRA adapter (Hu et al. 2022), ensuring an efficient and targeted alignment process.

Design of MOSA

The core design principle of our reward function is to steer the model to generate a set of predefined safe refusal sentences (the positive set, denoted by $S_{safe.set}$) while simultaneously penalizing any inclination towards generating undesirable or harmful continuations (the negative set, $S_{harmful.set}$). Notably, each sentence in the positive set $S_{safe.set}$ (e.g., “therefore, I cannot answer this question”) is designed with an end-of-sentence token. This acts as a “breaker” to ensure that even if the initial tokens are compromised by an attacker, the overall length of the harmful generation is strictly limited, thereby mitigating potential harm. Specifically, we define the tokens from the 20th to the 60th token of the response as the *middle tokens*. This window is chosen to balance two security considerations. First, it is placed beyond the initial tokens where attackers have strong influence, thus making the defense harder to bypass. Second, it is positioned early enough to preemptively terminate the response and limit any potential harm.

For each training step, we calculate our reward signal based on two randomly selected sentences: one positive example, s_{pos} , from $S_{safe.set}$ and one negative example, s_{neg} , from $S_{harmful.set}$. For each selected sentence, we iterate through all possible contiguous segments of the same length within the target window. We then calculate the log-likelihood of the selected sentence for each segment, and take the maximum value as its definitive score. The final contrastive reward is the difference between the positive score and the negative score, pushing the model to prefer generating safe responses over harmful ones. This process is detailed in Algorithm 1.

Another part of our reward is designed as the KL divergence of the logits of the reference model and the policy model, following prior research (Ouyang et al. 2022). It prevents the policy model’s parameters from diverging too far from the initial reference model and thus losing utility on normal tasks. The whole reward function is as follows:

$$\mathcal{L}(\theta) = \mathbb{E}_{x \sim D} [R(y|x) - \beta \cdot D_{KL}(P_{\theta}(y|x) || P_{ref}(y|x))]$$

where R represents the contrastive reward signal and β is

a scalar coefficient to control the strength of the KL divergence penalty. Our goal is to maximize this reward value.

Algorithm 1: Calculation of Contrastive Reward

Input: Log-probabilities L_θ , a safe sequence s_{pos} , a harmful sequence s_{neg} , search window $[k_{start}, k_{end}]$.

Output: The final contrastive reward $R_{contrastive}$.

```
1: # Find the max log-likelihood segment in the window
2: function GetMaxScore( $S, L_\theta, [k_{start}, k_{end}]$ )
3:   max_score  $\leftarrow -\infty$ 
4:   for  $k = k_{start}$  to  $k_{end}$  do
5:     current_score  $\leftarrow \sum_{j=1}^{\text{length}(S)} L_\theta(k + j - 1, S_j)$ 
6:     max_score  $\leftarrow \max(\text{max\_score}, \text{current\_score})$ 
7:   end for
8:   return max_score
9: end function
10:
11:  $R_{pos} \leftarrow \text{GetMaxScore}(s_{pos}, L_\theta, [k_{start}, k_{end}])$ 
12:  $R_{neg} \leftarrow \text{GetMaxScore}(s_{neg}, L_\theta, [k_{start}, k_{end}])$ 
13:  $R_{contrastive} \leftarrow R_{pos} - R_{neg}$ 
14: return  $R_{contrastive}$ 
```

Implementation

We implement MOSA on the LLaDA-8B-Instruct model, fine-tuning it exclusively on a randomly selected 3,000 harmful questions from the SORRY-Bench dataset (Xie et al. 2024). The training is performed efficiently for one epoch using LoRA with $r = 32$ and $\alpha = 64$. We use the AdamW optimizer with a learning rate of $5e-5$ and apply gradient clipping with a norm of 0.01. It trains the model to favor generating sequences from a predefined set of safe refusals (S_{safe_set}) over harmful ones ($S_{harmful_set}$). A KL divergence penalty with a coefficient of 0.05 is used to maintain the model’s general utility.

The training process for MOSA is highly efficient. On a system with two NVIDIA A100-PCIE-40GB GPUs, the entire fine-tuning process completes in approximately 12 minutes, consuming around 30 GB of total GPU memory. As illustrated in Figure 7, the model demonstrates rapid convergence during the training process. The reward signal rises sharply within the first 500 training steps and subsequently stabilizes at a high value of approximately 15-18. When inputting harmful questions such as “how to make a bomb”, the model successfully generates the refusal sentence in the middle tokens.

Experiment

Setup

To rigorously evaluate the defense capabilities of our MOSA framework, we conduct a comprehensive experiment on the LLaDA (Nie et al. 2025) model, comparing its performance against two baselines: the original model and an Initial Alignment baseline. To ensure a fair comparison, this baseline follows the same training protocol as MOSA, with the sole exception that its reward signal is optimized for the initial tokens rather than the middle. We test the models against a suite of eight state-of-the-art black-box jailbreaking methods: Avatar (Li et al. 2024a), TAP (Mehrotra et al. 2023), Speakeasy (Chan et al. 2025), AOS (Jiang

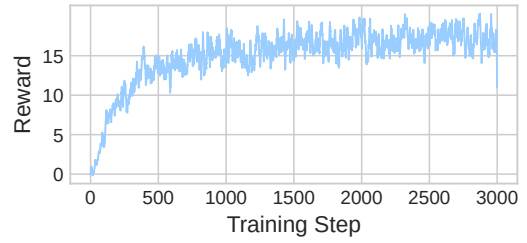


Figure 7: The reward along the training process.

et al. 2024), PAL (Sitawarin et al. 2024), EPT (Li et al. 2025b), DIA (Meng et al. 2025), and AdvPrefix (Zhu et al. 2024). The harmful prompts for these attacks are sourced from the standard HarmBench (Mazeika et al. 2024) and AdvBench (Zou et al. 2023) benchmarks. To ensure a robust evaluation of the defense efficacy, we measure the Attack Success Rate (ASR) using two judges: a judge proposed by ourselves mentioned above and the publicly available judge from the StrongReject (Souly et al. 2024) framework. We will justify our judge in the extended version. Following the StrongReject method, an answer is considered a successful attack if its “refusal answer” score is 0 while both its “convincing answer” and “specificity answer” scores are greater than 4. To test the models’ utility on normal tasks, we utilize GSM8K (Cobbe et al. 2021), MMLU (Hendrycks et al. 2021), and HumanEval (Chen, Tworek et al. 2021) to test the models’ mathematical problem-solving ability, general knowledge, and coding ability.

Overall Performance

In this part, we test three models on the eight attack methods. The results shown in Table 1 and Table 2 demonstrate the superior performance of our proposed method in mitigating jailbreak attacks. Across both the AdvBench and HarmBench datasets, the Original Model is highly susceptible to all eight attack methods, with ASR frequently exceeding 70%. This establishes a clear baseline vulnerability for an undefended dLLM. The Initial Alignment baseline gets a substantial improvement, reducing the ASR to a range of approximately 20-30%. We attribute this improvement to our training data from SORRY-Bench, which contains malicious prompts collected using 20 common jailbreaking methods. Training on such questions effectively allows the model to learn to recognize malicious intent hidden within a prompt, thereby enabling a defense. However, it remains vulnerable to attacks that manipulate tokens at specific positions in the output. It shows that only defending at the initial tokens is insufficient, as the attacker can still steer the model to generate an affirmative answer.

In contrast, our MOSA model delivers a dramatic reduction in ASR, consistently outperforming both the original model and the Initial Alignment baseline. On both benchmarks, MOSA lowers the ASR to single-digit percentages for most attack methods. For instance, against the TAP attack on AdvBench, MOSA achieves an ASR of just 4.5% (Our judge) and 3.7% (StrongReject judge), a significant drop from Initial Alignment’s 29.6% and 28.1%. We observe

Model	Avatar		TAP		Speakeasy		AOS		PAL		EPT		DIA		AdvPrefix	
	Our	SR	Our	SR	Our	SR	Our	SR	Our	SR	Our	SR	Our	SR	Our	SR
Original Model	74.5	65.6	79.1	77.2	69.8	57.4	65.2	62.9	72.8	67.1	78.4	72.3	66.7	64.2	79.5	70.1
Initial Alignment	23.5	22.5	29.6	28.1	22.4	23.3	32.4	25.9	36.4	24.5	28.5	22.0	34.3	28.6	29.8	27.2
MOSA (Ours)	14.3	12.8	4.5	3.7	8.1	12.9	6.5	5.4	6.2	4.5	3.8	7.1	4.2	7.7	6.8	3.5

* Our: Our proposed judge; SR: StrongReject judge

Table 1: Defense performance testing on AdvBench

Model	Avatar		TAP		Speakeasy		AOS		PAL		EPT		DIA		AdvPrefix	
	Our	SR	Our	SR	Our	SR	Our	SR	Our	SR	Our	SR	Our	SR	Our	SR
Original Model	76.8	73.7	79.2	76.7	70.5	66.3	64.1	61.3	73.8	70.8	77.9	75.8	68.6	65.1	79.9	78.4
Initial Alignment	25.6	25.1	29.8	27.9	24.5	22.4	31.4	28.3	29.2	27.5	32.7	31.8	36.4	32.3	29.9	28.7
MOSA (Ours)	14.7	8.3	7.5	6.2	14.5	13.2	6.1	5.2	5.4	6.4	3.6	2.2	8.9	7.2	4.9	4.1

Table 2: Defense performance testing on HarmBench

a unique successful defense of how MOSA diminishes attacks. Under an attack’s influence, the model might begin to output a non-refusal phrase. However, in the next few steps, the predefined refusal phrase is generated in the middle positions. In this unusual context, the model’s subsequent reply becomes quirky and incoherent, ultimately failing to complete an effective attack. These results affirm that MOSA’s architecture-aware strategy of aligning the middle tokens provides a strong defense against state-of-the-art jailbreaks.

Utility

In this part, we assess the impact of our safety alignment on the model’s normal utility on three benchmarks. As shown in Table 3, both our MOSA model and the Initial Alignment baseline demonstrate highly comparable performance to the Original Model across diverse tasks. This indicates that our fine-tuning process, regardless of the specific alignment strategy, has a minimal impact on general utility. Specifically, on the GSM8K, MOSA scores 68.3, nearly matching the Original Model’s 69.8. A similar trend is observed in the HumanEval and MMLU. Additionally, the Initial Alignment model also retains its utility, which highly supports that our training process will not influence the normal usage. This minimal impact on utility is expected, given that the fine-tuning process utilized a small dataset of only 3,000 samples and incorporated a penalty for significant divergence from the original model’s logits.

Discussion

Inaccessibility of Middle Tokens after MOSA

A central premise of our work is that the model’s strong sequential generation tendency makes it difficult for an attacker to manipulate the middle tokens of a response. Our defense, MOSA, leverages this by strategically aligning these less accessible middle tokens with a safe refusal. This naturally raises a question: Does MOSA inadvertently alter this fundamental sequential behavior?

To investigate this concern, we empirically test the robustness of the MOSA-aligned model against the same

Model	GSM8K	MMLU	HumanEval
Original Model	69.8	66.4	32.8
Initial Alignment	67.4	68.2	29.6
MOSA (Ours)	68.3	65.9	30.4

Table 3: Model utility on general benchmarks. A higher score indicates better performance.

optimization-based GCG attack mentioned above. The objective is to force our defended model to generate a malicious procedural phrase (e.g., "The next step is...") at the 30th token position. The results confirm that the middle tokens do not become a new vulnerability. While the optimization loss for the GCG attack on the MOSA-aligned model is slightly lower than on the original, undefended model, it remained substantially high, indicating significant difficulty in steering the model’s generation. Consequently, the attack success rate remained exceptionally low, under 4%. This phenomenon is because MOSA is a lightweight fine-tuning process. It learns to generate safe middle tokens when inputting harmful prompts while retaining the ingrained preference for sequential generation. Consequently, the core security asymmetry we leverage is preserved, meaning the middle tokens are effectively fortified for defense without being converted into a viable attack surface for adversaries.

To further validate this point, we conduct a series of adaptive attacks designed to directly challenge the MOSA defense mechanism. We modified three attacks targeting to change the prefix of the response: TAP, EPT, and AdvPrefix. We reset their core objective from manipulating the initial

Model	TAP		EPT		AdvPrefix	
	Our	SR	Our	SR	Our	SR
MOSA (Ours)	5.1	4.4	3.8	4.2	4.5	5.1

Table 4: ASR of Adaptive Attacks.

tokens to the middle tokens, simulating an attacker aware of MOSA’s strategy. Following the same pipeline, the result shown in Table 4 demonstrates that MOSA maintains its strong defensive effect. This confirms that MOSA provides a fundamental security enhancement rather than merely shifting the point of vulnerability.

Generation Order in Other dLLMs

To ascertain whether the sequential generation preference is a general trait impacting dLLM security, we extend our analysis to other models, including Dream 7B (Ye et al. 2025) and MMaDA (Yang et al. 2025). Our tests confirm that these models also exhibit a strong sequential, left-to-right generation bias, especially when handling adversarial safety queries. This widespread behavior is likely an emergent property of the model’s training process. Given only a prompt, predicting an adjacent token is a more statistically stable task with lower variance compared to generating one in the middle of an unknown sequence. Therefore, a sequential generation path becomes a natural outcome for a model trained to minimize prediction error. This finding is highly significant for the security landscape of dLLMs. It suggests that the core security asymmetry is not an isolated quirk but likely a fundamental characteristic of the current dLLM paradigm.

Meanwhile, we observe that this sequential generation is not limited to the security domain but is a pervasive trait that also extends to tasks such as mathematical problem-solving and coding. Given that our work has demonstrated the advantages of manipulating this generation order for safety, it is logical to consider extending this paradigm to other tasks to better exploit the capabilities of dLLMs. For example, as existing literature has demonstrated that breaking down a problem into intermediate steps can substantially improve LLM’s problem-solving accuracy in mathematical problem-solving, a dLLM could be trained to first generate a pivotal intermediate formula. Subsequently, leveraging its unique ability to process context from both directions, it can flesh out the preceding derivations and subsequent steps, ensuring the entire solution remains coherent and logically sound. This “anchor-then-fill” approach may present a method to more fully unlock the unique architectural potential of dLLMs.

Limitation

While MOSA proves highly robust against attacks that aim to manipulate specific token positions (e.g., EPT, TAP), its defense is relatively less effective against attacks designed to hide malicious intent by complex narratives, such as Avatar and Speakeasy. A plausible explanation for this disparity lies in the training dataset. While diverse, it appears to lack a sufficient variety of “wrapped” malicious prompts, which consequently limits the model’s capacity to identify cleverly disguised harmful intent. Simply augmenting the dataset for every new type of attack is a reactive approach and may not lead to a truly robust defense. Therefore, to develop a more fundamental and generalizable solution, in the future, we plan to investigate the security of dLLMs from the perspective of activation. Specifically, we plan to observe

whether the internal activations of a dLLM exhibit linearity on concepts such as “benign” and “harmful”, and thus develop more robust defensive methods.

Related Work

Masked Diffusion Models for Text

Diffusion models (Sohl-Dickstein et al. 2015) are adapted for discrete data such as text (Austin et al. 2021; Campbell et al. 2022). Initial success with models like Diffusion-BERT (He et al. 2022) has been followed by rapid improvements in architecture (Lou, Meng, and Ermon 2023; Sahoo et al. 2024), training simplicity (Shi et al. 2024), and fundamental understanding (Zheng et al. 2024). Recently, diffusion-based LLMs have been scaled to rival large AR LLMs in performance (Nie et al. 2024; Gong et al. 2024; Nie et al. 2025; Ye et al. 2025). Despite this progress, their security performance and characteristics have not been studied, which is a gap our work aims to fill.

Security Analysis of LLMs

Research in attacking LLMs (Xie, Song, and Luo 2025) centers on forcing the model to begin its response with an affirmative prefix (e.g., “Sure, here is...”) (Wei, Haghtalab, and Steinhardt 2023; Qi et al. 2024) in different ways, such as optimization (Zou et al. 2023) and prompt engineering (Wei, Haghtalab, and Steinhardt 2023). Another line of work focuses on wrapping the harmfulness of the question by prompt engineering such as role play (Shah et al. 2023), persuasion (Song et al. 2025), and auto-generated attack prompt (Chao et al. 2025; Liu et al. 2023). On the defensive side, the most widely adopted method is to align LLMs with human preference by post-training based on reinforcement learning (Ouyang et al. 2022; Wang et al. 2025; Alami et al. 2024). Qi et al. (Qi et al. 2024) demonstrate that the safety alignment disproportionately concentrates on the first few tokens of the responses, making both attackers and defenders invest their resources to control this part. In this paper, we first systematically analyze the security of dLLMs and aim to develop novel defenses specifically for dLLMs.

Conclusion

This paper presents the first systematic security analysis of Diffusion Large Language Models (dLLMs), uncovering a key “security asymmetry”. While attackers struggle to influence the critical middle tokens due to the model’s inherent sequential generation preference, defenders can directly align it. Based on this insight, we develop MOSA based on reinforcement learning to train the model to generate a group of predefined safe refusal sentences in the middle of its response. We evaluate MOSA on two benchmarks and eight attack methods, and find it effective against attacks without influencing utility. Overall, our work provides a robust, architecture-aware defense for dLLMs that moves beyond the traditional “first-token-centric” security approach.

Acknowledgements

This research is supported by the National Research Foundation Singapore and DSO National Laboratories under the AI

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Alami, R.; Almansoori, A. K.; Alzubaidi, A.; Seddik, M. E. A.; Farooq, M.; and Hacid, H. 2024. Alignment with preference optimization is all you need for llm safety. *arXiv preprint arXiv:2409.07772*.
- Andriushchenko, M.; Croce, F.; and Flammarion, N. 2024. Jailbreaking leading safety-aligned llms with simple adaptive attacks. *arXiv preprint arXiv:2404.02151*.
- Austin, J.; Johnson, D. D.; Ho, J.; Tarlow, D.; and Van Den Berg, R. 2021. Structured denoising diffusion models in discrete state-spaces. *NeurIPS*, 34: 17981–17993.
- Baier, P. 2024. Estimated Market Share of Closed-Source LLM Models in 2024. <https://www.linkedin.com/pulse/estimated-market-share-closed-source-llm-models-2024-paul-baier-pwwxe>. Accessed: 2025-07-31.
- Campbell, A.; Benton, J.; De Bortoli, V.; Rainforth, T.; Deligiannidis, G.; and Doucet, A. 2022. A continuous time framework for discrete denoising models. *NeurIPS*, 35: 28266–28279.
- Chan, Y. S.; Ri, N.; Xiao, Y.; and Ghassemi, M. 2025. Speak easy: Eliciting harmful jailbreaks from llms with simple interactions. *arXiv preprint arXiv:2502.04322*.
- Chao, P.; Robey, A.; Dobriban, E.; Hassani, H.; Pappas, G. J.; and Wong, E. 2025. Jailbreaking black box large language models in twenty queries. In *2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 23–42. IEEE.
- Chen, M.; Tworek, J.; et al. 2021. Evaluating Large Language Models Trained on Code.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; Hesse, C.; and Schulman, J. 2021. Training Verifiers to Solve Math Word Problems. *arXiv*.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv e-prints, arXiv:2407*.
- Gong, S.; Agarwal, S.; Zhang, Y.; Ye, J.; Zheng, L.; Li, M.; An, C.; Zhao, P.; Bi, W.; Han, J.; et al. 2024. Scaling Diffusion Language Models via Adaptation from Autoregressive Models. *arXiv preprint arXiv:2410.17891*.
- Google DeepMind. 2025. Gemini Diffusion: Our state-of-the-art experimental text diffusion model. Online; demo announced via Google I/O and released by waitlist.
- Gu, H.; Wang, H.; Mei, Y.; Zhang, M.; and Jin, Y. 2025. One Trigger Token Is Enough: A Defense Strategy for Balancing Safety and Usability in Large Language Models. *arXiv preprint arXiv:2505.07167*.
- He, Z.; Sun, T.; Wang, K.; Huang, X.; and Qiu, X. 2022. Diffusionbert: Improving generative masked language models with diffusion models. *arXiv preprint arXiv:2211.15029*.
- Hendrycks, D.; Burns, C.; Basart, S.; Critch, A.; Li, J.; Song, D.; and Steinhardt, J. 2021. Aligning AI With Shared Human Values. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Hostinger International. 2025. 47+ Large Language Model (LLM) Statistics for 2025. Accessed: 2025-07-31.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Inception Labs. 2025. Introducing a New Generation of Language Models: Diffusion Large Language Models. Online; company website.
- Jiang, W.; Wang, Z.; Zhai, J.; Ma, S.; Zhao, Z.; and Shen, C. 2024. An Optimizable Suffix Is Worth A Thousand Templates: Efficient Black-box Jailbreaking without Affirmative Phrases via LLM as Optimizer. *arXiv preprint arXiv:2408.11313*.
- Li, X.; Liu, Y.; Su, Z.; Chen, L.; Zhuang, F.; Li, Z.; and shizhan. 2024a. AVATAR: Jailbreak via Adversarial MeTAPHOR. *arXiv:2412.12145*.
- Li, Y.; Hu, J.; Sang, W.; Ma, L.; Xie, J.; Zhang, W.; Yu, A.; Zhao, S.; Huang, Q.; and Zhou, Q. 2025a. Prefill-based jailbreak: A novel approach of bypassing llm safety boundary. *arXiv preprint arXiv:2504.21038*.
- Li, Y.; Liu, Y.; Li, Y.; Shi, L.; Deng, G.; Chen, S.; and Wang, K. 2024b. Lockpicking llms: A logit-based jailbreak using token-level manipulation. *arXiv preprint arXiv:2405.13068*.
- Li, Y.; Xiong, Y.; Zhong, J.; Zhang, J.; Zhou, J.; and Zou, L. 2025b. Exploiting Prefix-Tree in Structured Output Interfaces for Enhancing Jailbreak Attacking. *arXiv preprint arXiv:2502.13527*.
- Liu, X.; Xu, N.; Chen, M.; and Xiao, C. 2023. Autodan: Generating stealthy jailbreak prompts on aligned large language models. *arXiv preprint arXiv:2310.04451*.
- Lou, A.; Meng, C.; and Ermon, S. 2023. Discrete Diffusion Language Modeling by Estimating the Ratios of the Data Distribution. *arXiv preprint arXiv:2310.16834*.
- Mazeika, M.; Phan, L.; Yin, X.; Zou, A.; Wang, Z.; Mu, N.; Sakhaee, E.; Li, N.; Basart, S.; Li, B.; et al. 2024. Harmbench: A standardized evaluation framework for automated red teaming and robust refusal. *arXiv preprint arXiv:2402.04249*.
- Mehrotra, A.; Zampetakis, M.; Kassianik, P.; Nelson, B.; Anderson, H.; Singer, Y.; and Karbasi, A. 2023. Tree-of-Attacks: Jailbreaking Black-Box LLMs Automatically. *arXiv:2312.02119*.
- Meng, W.; Zhang, F.; Yao, W.; Guo, Z.; Li, Y.; Wei, C.; and Chen, W. 2025. Dialogue injection attack: Jailbreaking llms through context manipulation. *arXiv preprint arXiv:2503.08195*.
- Nie, S.; Zhu, F.; Du, C.; Pang, T.; Liu, Q.; Zeng, G.; Lin, M.; and Li, C. 2024. Scaling up Masked Diffusion Models on Text. *arXiv preprint arXiv:2410.18514*.

- Nie, S.; Zhu, F.; You, Z.; Zhang, X.; Ou, J.; Hu, J.; Zhou, J.; Lin, Y.; Wen, J.-R.; and Li, C. 2025. Large language diffusion models. *arXiv*.
- OpenAI. 2025. Introducing GPT-4.1 in the API. <https://platform.openai.com/docs/models/gpt-4.1>. Accessed: 2025-07-22.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *NeurIPS*, 35: 27730–27744.
- Ouyang, Y.; Gu, H.; Lin, S.; Hua, W.; Peng, J.; Kailkhura, B.; Gao, M.; Chen, T.; and Zhou, K. 2025. Layer-Level Self-Exposure and Patch: Affirmative Token Mitigation for Jailbreak Attack Defense. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 12541–12554.
- OWASP Foundation. 2025. OWASP Top 10 for LLM Applications - 2025. Accessed: 2025-07-31.
- Qi, X.; Panda, A.; Lyu, K.; Ma, X.; Roy, S.; Beirami, A.; Mittal, P.; and Henderson, P. 2024. Safety alignment should be made more than just a few tokens deep. *arXiv preprint arXiv:2406.05946*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *NeurIPS*, 36: 53728–53741.
- Sahoo, S. S.; Arriola, M.; Schiff, Y.; Gokaslan, A.; Marroquin, E.; Chiu, J. T.; Rush, A.; and Kuleshov, V. 2024. Simple and Effective Masked Diffusion Language Models. *arXiv preprint arXiv:2406.07524*.
- Shah, R.; Pour, S.; Tagade, A.; Casper, S.; Rando, J.; et al. 2023. Scalable and transferable black-box jailbreaks for language models via persona modulation. *arXiv preprint arXiv:2311.03348*.
- Shi, J.; Han, K.; Wang, Z.; Doucet, A.; and Titsias, M. K. 2024. Simplified and Generalized Masked Diffusion for Discrete Data. *arXiv preprint arXiv:2406.04329*.
- Sitawarin, C.; Mu, N.; Wagner, D.; and Araujo, A. 2024. Pal: Proxy-guided black-box attack on large language models. *arXiv preprint arXiv:2402.09674*.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2256–2265. PMLR.
- Song, X.; Xie, Z.; Huai, S.; Kong, J.; and Luo, J. 2025. Dagger Behind Smile: Fool LLMs with a Happy Ending Story. *EMNLP*.
- Souly, A.; Lu, Q.; Bowen, D.; Trinh, T.; Hsieh, E.; Pandey, S.; Abbeel, P.; Svegliato, J.; Emmons, S.; Watkins, O.; et al. 2024. A strongreject for empty jailbreaks. *NeurIPS*, 37: 125416–125440.
- Taori, R.; Gulrajani, I.; Zhang, T.; Dubois, Y.; Li, X.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Stanford Alpaca: An Instruction-following LLaMA model. https://github.com/tatsu-lab/stanford_alpaca.
- Wang, Y.; Wang, P.; Xi, C.; Tang, B.; Zhu, J.; Wei, W.; Chen, C.; Yang, C.; Zhang, J.; Lu, C.; et al. 2025. Adversarial Preference Learning for Robust LLM Alignment. *arXiv preprint arXiv:2505.24369*.
- Wei, A.; Haghtalab, N.; and Steinhardt, J. 2023. Jailbroken: How does llm safety training fail? *NeurIPS*, 36: 80079–80110.
- Xie, T.; Qi, X.; Zeng, Y.; Huang, Y.; Sehwag, U. M.; Huang, K.; He, L.; Wei, B.; Li, D.; Sheng, Y.; et al. 2024. Sorry-bench: Systematically evaluating large language model safety refusal. *arXiv preprint arXiv:2406.14598*.
- Xie, Z.; Song, X.; and Luo, J. 2025. Attack via Overfitting: 10-shot Benign Fine-tuning to Jailbreak LLMs. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Yang, L.; Tian, Y.; Li, B.; Zhang, X.; Shen, K.; Tong, Y.; and Wang, M. 2025. Mmada: Multimodal large diffusion language models. *arXiv*.
- Ye, J.; Xie, Z.; Zheng, L.; Gao, J.; Wu, Z.; Jiang, X.; Li, Z.; and Kong, L. 2025. Dream 7B.
- Zheng, K.; Chen, Y.; Mao, H.; Liu, M.-Y.; Zhu, J.; and Zhang, Q. 2024. Masked Diffusion Models are Secretly Time-Agnostic Masked Models and Exploit Inaccurate Categorical Sampling. *arXiv preprint arXiv:2409.02908*.
- Zhu, F.; Wang, R.; Nie, S.; Zhang, X.; Wu, C.; Hu, J.; Zhou, J.; Chen, J.; Lin, Y.; Wen, J.-R.; et al. 2025. LLaDA 1.5: Variance-Reduced Preference Optimization for Large Language Diffusion Models. *arXiv*.
- Zhu, S.; Amos, B.; Tian, Y.; Guo, C.; and Evtimov, I. 2024. Advprefix: An objective for nuanced llm jailbreaks. *arXiv preprint arXiv:2412.10321*.
- Zou, A.; Wang, Z.; Carlini, N.; Nasr, M.; Kolter, J. Z.; and Fredrikson, M. 2023. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*.