

# 000 BEYOND SURFACE REASONING: UNVEILING THE TRUE 001 LONG CHAIN-OF-THOUGHT CAPACITY OF DIFFUSION 002 LARGE LANGUAGE MODELS 003

004 **Anonymous authors**  
005  
006

007 Paper under double-blind review  
008  
009

## 010 ABSTRACT 011

012 Recently, Diffusion Large Language Models (DLLMs) have offered high through-  
013 put and effective sequential reasoning, making them a competitive alternative  
014 to autoregressive LLMs (ALLMs). However, parallel decoding, which enables  
015 simultaneous token updates, conflicts with the causal order often required for  
016 rigorous reasoning. We first identify this conflict as the core Parallel–Sequential  
017 Contradiction (PSC). Behavioral analyses in both simple and complex reasoning  
018 tasks show that DLLMs exhibit genuine parallelism only for directly decidable  
019 outputs. As task difficulty increases, they revert to autoregressive-like behavior,  
020 a limitation exacerbated by autoregressive prompting, which nearly doubles the  
021 number of decoding steps with remasking without improving quality. Moreover,  
022 PSC restricts DLLMs’ self-reflection, reasoning depth, and exploratory breadth.  
023 To further characterize PSC, we introduce three scaling dimensions for DLLMs:  
024 parallel, diffusion, and sequential. Empirically, while parallel scaling yields con-  
025 sistent improvements, diffusion and sequential scaling are constrained by PSC.  
026 Based on these findings, we propose several practical mitigations, parallel-oriented  
027 prompting, diffusion early stopping, and parallel scaling, to reduce PSC-induced  
028 ineffectiveness and inefficiencies.  
029

## 030 1 INTRODUCTION 031

032 In recent years, diffusion large language models (DLLMs) have emerged as a novel generative  
033 paradigm, attracting increasing research attention (Li et al., 2025; Yang et al., 2025). Repre-  
034 sentative works such as LLaDA (Nie et al., 2025b) and Dream (Ye et al., 2025) adopt a two-stage  
035 mask-denoising training strategy combined with parallel decoding for masked token prediction, effec-  
036 tively mitigating the “reversal curse” in traditional autoregressive large language models (ALLMs).  
037 Mercury (Inception Labs, 2025) and Fast-DLLM (Wu et al., 2025) further demonstrate the parallel  
038 efficiency of DLLMs, achieving an impressive generation speed in code tasks.

039 Meanwhile, the rapid development of the Long Chain-of-Thought (Long CoT) (Guo et al., 2025;  
040 Chen et al., 2024; 2025) has spurred increasing research on applying DLLMs to extended reasoning  
041 tasks (Wang et al., 2025b;a). Zhao et al. (2025) and Tang et al. (2025) employ diffusion-augmented  
042 SFT and GRPO to further improve reasoning (Gong et al., 2025). Moreover, Trado (Wang et al.,  
043 2025b) exploits overlooked information in sampling trajectories, achieving substantial gains.

044 As shown in Figure 1 (a), DLLMs generate text in parallel, producing a few non-sequential words in a  
045 single diffusion step. In sequential reasoning scenarios (Figure 1 (b)), the generation of  $step_i$  requires  
046 the completion of  $step_{i-1}$ , leading to lower entropy (Cui et al., 2025; Agarwal et al., 2025). In  
047 contrast, Figure 1 (c) shows DLLMs to parallel-decode by generating  $step_{i+1}$  before  $step_i$ , resulting  
048 in high entropy. Nevertheless, these parallel and sequential processes are inherently contradictory:  
049 **parallelism involves simultaneous processing, while sequential reasoning requires ordered steps.**  
050 To address this, we introduce the **Parallel–Sequential Contradiction (PSC)**, which explores the  
051 underlying mechanisms and practical implications of diffusion-based reasoning.

052 To investigate this issue systematically, as shown in Figure 1 (d, e), we focus on two central  
053 research questions: (1) **Do DLLMs truly perform parallel reasoning that avoids PSC?** (2) **What**  
**challenges do DLLMs meet in Long CoT based on PSC?** To address the first question, we analyze

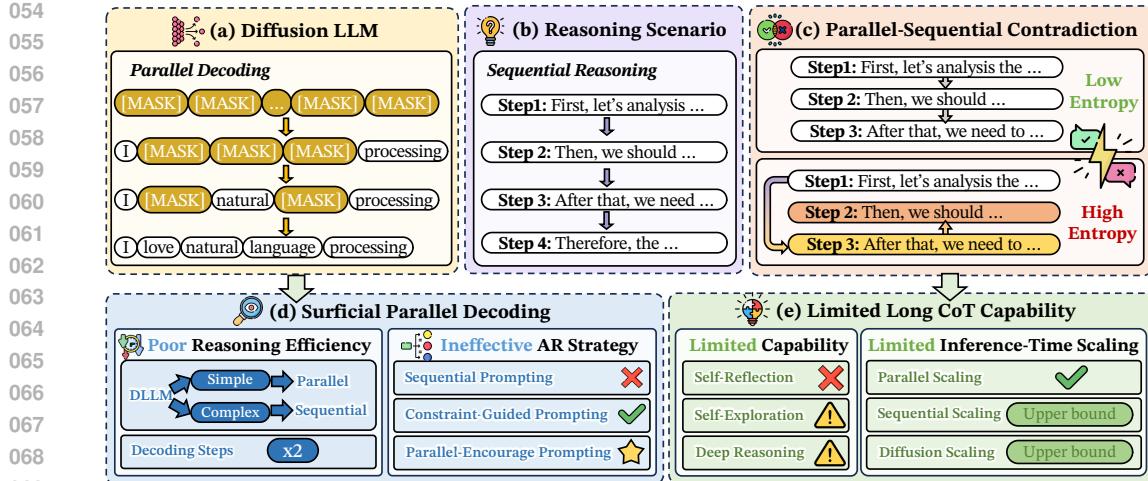


Figure 1: Overview of our work. Applying DLLMs to reasoning scenarios reveals an inherent contradiction between parallel processing and sequential reasoning, leading to high entropy, **surficial** parallel decoding, and limited Long CoT capabilities.

the decoding behavior of DLLMs in both simple and complex reasoning scenarios. Our findings show that DLLMs fail to achieve genuine parallel reasoning due to the PSC. They perform superficial parallel computation when outputs can be directly produced, but revert to an autoregressive mode under higher reasoning demands. This reliance on autoregression affects computational efficiency, which nearly doubles the computational cost with low confidence remasking. Furthermore, while autoregressive prompting is effective in ALLMs, it conflicts with DLLMs’ parallel decoding design, amplifying the PSC of DLLMs. In contrast, strategies that reduce contradiction, such as conditional prompting or prompts that encourage parallel generation, effectively enhance prompting performance.

To understand the second question, we examine the core capabilities of Long CoT in DLLMs. Our analysis reveals that, when faced with PSC, DLLMs often demonstrate limited self-reflection, shallow reasoning depth, and constrained exploratory behavior. Furthermore, we introduce three scaling dimensions for inference time, specifically designed for DLLMs: parallel, diffusion, and sequential scaling. Our findings show that both diffusion and sequential scaling are significantly constrained by PSC, while the parallel scaling law remains unaffected due to its vertical relationship with PSC.

In summary, our key contributions are as follows:

- **Identification of Parallel-Sequential Contradictions:** To our knowledge, we first identify the Parallel-Sequential Contradiction (PSC) in DLLMs for Long CoT. We demonstrate that PSC leads to superficial parallel reasoning and reduced efficiency, requiring twice the decoding steps.
- **Systematic Exploration of DLLM Reasoning Limitation:** We conduct a systematic evaluation of DLLM reasoning, identifying the degradation of three core Long CoT capabilities, confirming the ineffectiveness of traditional autoregressive prompting methods, and demonstrating that diffusion scaling and sequential scaling are upper-bounded by PSC limitations.
- **Novel Mitigation Strategies:** We propose novel strategies to mitigate these issues and enhance DLLM reasoning. Our methods include parallel-encouraging prompting, diffusion early stopping, and parallel scaling, which substantially alleviate the constraints imposed by PSC.

## 2 PARALLEL-SEQUENTIAL CONTRADICTION

### 2.1 PARALLEL MASKED DIFFUSION LANGUAGE MODELS

In Diffusion Large Language Models (DLLMs), inference reconstructs missing spans by predicting masked tokens conditioned on a partially masked input. Its goal is modelling the conditional likelihood  $p_\theta(x_0^i|x_l)$  for masked positions:

$$-\mathbb{E}_{l,x_0,x_l} \left[ \frac{L}{l} \sum_{i=1}^L \mathbf{1}[x_l^i \in \mathcal{M}] \log p_\theta(x_0^i|x_l) \right], \quad (1)$$

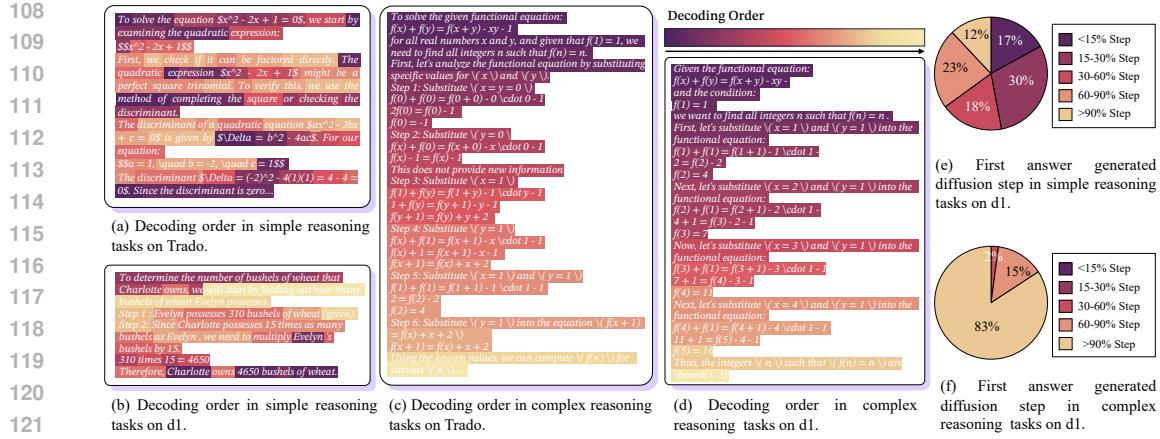


Figure 2: Diffusion order analysis with d1 (Zhao et al., 2025) and Trado (Wang et al., 2025b), where later decoding orders are indicated by shallower colors.

where  $L$  denotes the total number of tokens;  $l$  is the number of masked tokens, uniformly sampled from  $\{1, 2, \dots, L\}$ ;  $x_0$  is the complete original sequence.  $x_l$  is the partially masked sequence obtained by replacing those  $l$  positions in  $x_0$  with mask tokens, which serves as the conditional input. The indicator  $\mathbf{1}[x_l^i \in \mathcal{M}]$  equals 1 if position  $i$  is masked and 0 otherwise.

## 2.2 SEQUENTIAL LONG CHAIN-OF-THOUGHT REASONING

Long Chain-of-Thought (Long CoT) allows LLMs to tackle complex problems by generating a sequence of reasoning steps. This method solves a problem  $P$  by following an ordered series of steps  $S_1, S_2, \dots, S_n$ , leading to the final answer  $A$ . Formally, it can be defined as:

$$p_\theta(A|P) = \prod_{t=1}^{n+1} p_\theta(S_t|P, S_{<t}). \quad (2)$$

Here,  $S_{n+1} = A$ , meaning the final answer is treated as the last step of the reasoning sequence. When generating each step  $S_t$ , the model computes the conditional probability based on the problem  $P$  and all previously generated steps  $S_{<t}$ .

## 2.3 PARALLEL-SEQUENTIAL CONTRADICTION

We propose the term "Parallel-Sequential Contradiction (PSC)" to capture the fundamental tension faced by diffusion models in reasoning tasks. This contradiction manifests at two levels:

- **Mechanism Level:** The inherent parallel decoding nature of diffusion models directly conflicts with the sequential dependency logic required for CoT reasoning at the computational level.
- **Behavioral Level:** This conflict causes the model's generation process to oscillate between "following low-entropy sequential logic" and "falling into high-entropy parallel guessing", reflecting an inherent inconsistency.

For **tasks with high parallelism**, downstream states typically yield predictable, high-probability outcomes, resulting in low predictive entropy. In these cases, optimizing the conditional probability  $p_\theta(S_k | S_1)$  is efficient, making non-autoregressive or semi-parallel generation methods advantageous. In contrast, **tasks with strong sequential dependencies** exhibit high entropy when predicting distant future states in parallel. This uncertainty leads to significant predictive loss. To reduce this loss, the model is encouraged to break down the generation process into a sequence of low-entropy, step-by-step predictions. As a result, parallel generation conflicts with the model's objective of identifying a low-loss, high-probability sequential path. The formal proof is provided in Appendix B.

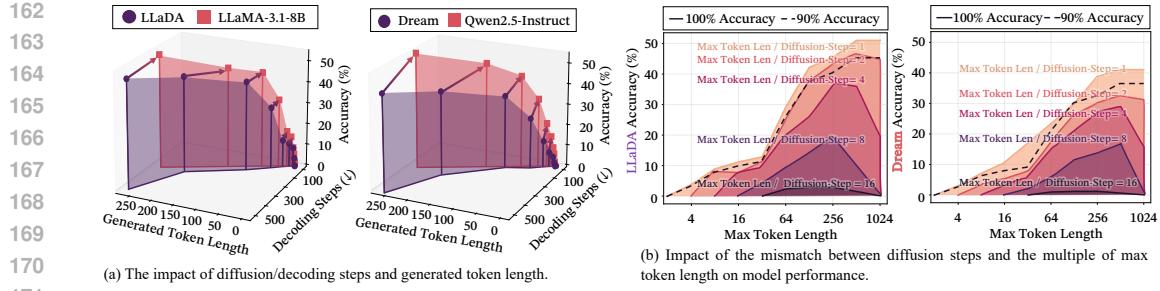


Figure 3: Diffusion speed analysis in Long-CoT-needed tasks with LLaDA-8B-Instruct (Nie et al., 2025a) and Dream-7B-Instruct (Ye et al., 2025) on BigGSM benchmark (Chen et al., 2024).

### 3 DO DLLMS TRUELY PERFORM PARALLEL REASONING THAT AVOIDS PSC?

#### 3.1 PARALLEL-SEQUENTIAL CONTRADICTIONS CAUSE SUPERFICIAL PARALLEL REASONING.

To examine whether DLLMs can genuinely perform parallel reasoning, we analyze their decoding behavior in both simple and complex scenarios. We have two key observations:

**Parallel-Sequential Contradictions cause superficial parallel reasoning in simple scenarios** As shown in Figure 2 (a, b), DLLMs demonstrate parallel reasoning in simple cases where the model can direct output results without reasoning. For instance, when solving " $x^2 + 2x + 1 = 0$ ", the model may simultaneously generate the Quadratic Formula " $\Delta = b^2 - 4ac$ " and the final solution " $x = -1$ " within a few diffusion steps. Following this, DLLMs complete the remaining reasoning steps in parallel, demonstrating the ability to leverage diffusion-based decoding to arrive at direct solutions without relying heavily on sequential reasoning. To further explore this, we analyze the distribution of answers generated in the initial diffusion steps. As shown in Figure 2 (e), over 47% of answers are produced within the first 30% of diffusion steps. This suggests that in simple scenarios, DLLMs are capable of performing parallel reasoning, even though the underlying thought process, such as applying the Quadratic Formula before deriving the result, is inherently sequential.

**For complex reasoning, DLLMs converge toward autoregressive-like behavior to avoid PSC** To examine DLLM behavior in complex reasoning tasks, we validate PSC where the model cannot directly output the correct answer. Figure 2 (c, d) shows that DLLMs increasingly resemble autoregressive models. For example, when addressing tasks beyond direct generation, the model defaults to an autoregressive process. This suggests difficulty in sustaining parallel reasoning, which shifts to step-by-step processing. Figure 2 (f) further confirms this observation: in complex tasks, answers emerge later in the diffusion steps, reflecting a stronger reliance on ordered reasoning. These findings indicate that DLLMs face inherent PSC challenges in balancing parallel generation with sequential reasoning, ultimately converging toward autoregressive-like processing in complex scenarios.

**PSC are relevant to stronger context dependencies, which even amplified by RL.** Through an information-theoretic analysis of conditional entropy and a dependency index, we show that stronger contextual dependencies intensify the PSC and cause marked performance degradation in DLLMs (see Appendix J). This behavior follows from inherent limitations of the parallel generation architecture, providing additional support for our theoretical interpretation of the PSC phenomenon. Additionally, reinforcement learning-based fine-tuning further amplifies the PSC in DLLMs (see Appendix I), suggesting that common post-training methods may worsen underlying architectural mismatches. Our analysis reveals that an increase in dependency strength (i.e., lower conditional entropy or higher DI) based on RL has a significant negative impact on DLLM performance, particularly when modeling long-range dependencies is required.

#### 3.2 DIFFUSION-STEP DILEMMA: SACRIFICING EFFICIENCY UNDER PSC

To investigate the reasoning efficiency of current DLLMs, we systematically categorize questions in BigGSM (Chen et al., 2024) into different sampling lengths and diffusion steps (with low-confidence remasking). We evaluate two representative DLLMs under exponentially increasing diffusion steps and max token lengths (ranging from 1 to 1024). See Appendix D for more details.

**When complex reasoning, DLLMs require significantly more diffusion steps than ALLMs.** As shown in Figure 3 (a), achieving comparable length and accuracy to ALLMs demands over 25% more

Model Name	BigGSM (Acc.)	GSM8K (Acc.)	Math-500 (Acc.)	HumanEval (Pass@1)	Average
Dream-7B-Instruct	41.15 (+0.00)	80.52 (+0.00)	37.00 (+0.00)	51.22 (+0.00)	52.47 (+0.00)
+Zero-CoT	41.15 (-0.00)	77.26 (-3.26)	34.00 (-3.00)	48.78 (-2.44)	50.30 (-2.18)
+Plan-and-Solve	34.75 (-6.40)	78.85 (-1.67)	20.20 (-16.80)	48.78 (-2.44)	45.65 (-6.83)
+Least-to-Most	35.25 (-5.90)	77.03 (-3.49)	16.20 (-20.80)	43.29 (-7.93)	42.94 (-9.53)
+Complex-CoT	41.80 (+0.65)	80.95 (+0.43)	37.40 (+0.40)	<b>52.44</b> (+1.22)	53.15 (+0.68)
+MARP	42.95 (+1.80)	80.52 (+0.00)	37.20 (+0.20)	51.22 (+0.00)	52.97 (+0.50)
+Diff-MARP	<b>47.21</b> (+6.06)	<b>82.64</b> (+2.21)	<b>43.60</b> (+6.60)	<b>52.44</b> (+1.22)	<b>56.47</b> (+4.00)
LLaDA-8B-Instruct	48.03 (+0.00)	75.36 (+0.00)	34.80 (+0.00)	32.32 (+0.00)	47.63 (+0.00)
+ Zero-CoT	35.57 (-12.46)	73.46 (-1.90)	32.40 (-2.40)	28.05 (-4.27)	42.37 (-5.26)
+ Plan-and-Solve	31.64 (-16.39)	72.33 (-3.03)	29.00 (-5.80)	27.44 (-4.88)	40.10 (-7.53)
+ Least-to-Most	34.75 (-13.28)	73.31 (-2.05)	30.80 (-4.00)	27.44 (-4.88)	41.58 (-6.05)
+ Complex-CoT	48.03 (+0.00)	76.50 (+1.14)	36.20 (+1.40)	36.59 (-4.27)	49.33 (+1.70)
+ MARP	48.20 (+0.17)	76.35 (+0.99)	34.40 (-0.40)	35.37 (-3.05)	48.58 (+0.95)
+Diff-MARP	<b>55.74</b> (+7.71)	<b>76.80</b> (+1.44)	<b>38.20</b> (+3.40)	<b>38.41</b> (+6.09)	<b>52.29</b> (+4.66)
LLaDA-v1.5	41.80 (+0.00)	74.98 (+0.00)	38.00 (+0.00)	36.59 (+0.00)	47.84 (+0.00)
+ Zero-CoT	36.39 (-5.41)	71.87 (-3.11)	37.20 (-0.80)	35.98 (-0.61)	45.36 (-2.48)
+ Plan-and-Solve	30.16 (-11.54)	74.37 (-0.61)	34.40 (-3.60)	35.98 (-0.61)	43.73 (-4.12)
+ Least-to-Most	35.90 (-5.90)	73.69 (-1.29)	34.60 (-3.40)	31.71 (-4.88)	43.98 (-3.87)
+ Complex-CoT	50.16 (+8.41)	75.51 (+0.53)	<b>39.40</b> (+1.40)	<b>39.02</b> (+2.43)	51.04 (+3.19)
+ MARP	42.13 (+0.33)	74.37 (0.61)	38.20 (+0.20)	37.20 (+0.61)	47.98 (+0.13)
+Diff-MARP	<b>54.49</b> (+12.79)	<b>76.50</b> (+1.52)	42.80 (+4.80)	38.41 (+1.82)	<b>53.08</b> (+5.23)
LLaDOU-Math	42.13 (+0.00)	81.88 (+0.00)	45.80 (+0.00)	39.02 (+0.00)	52.21 (+0.00)
+Zero-CoT	38.52 (-3.61)	80.95 (-0.93)	45.80 (-0.00)	37.80 (-1.22)	50.77 (-1.44)
+Plan-and-Solve	40.82 (-1.31)	81.12 (-0.76)	43.20 (-2.60)	38.41 (-0.61)	50.89 (-1.32)
+Least-to-Most	40.16 (-1.97)	79.08 (-2.80)	43.00 (-2.80)	36.59 (-2.43)	49.71 (-2.50)
+ Complex-CoT	43.77 (+1.64)	83.70 (+1.82)	45.80 (+0.00)	<b>42.07</b> (+3.05)	52.47 (+0.26)
+ MARP	41.15 (-0.98)	82.18 (+0.30)	45.60 (-0.20)	40.26 (+1.24)	52.30 (+0.09)
+Diff-MARP	<b>54.26</b> (+12.13)	<b>84.76</b> (+2.88)	<b>49.00</b> (+3.20)	40.85 (+1.83)	<b>57.22</b> (+5.01)

Table 1: Performance comparison across 4 benchmarks. **Bold** marks the best baseline score per metric. For each method, we report its most token-efficient variant. Here, “”: prompting strategies, “”: offline strategies, “”: online strategies. **See results in Table 4 with different decoding methods.**

diffusion or decoding steps with remasking. In extreme cases, DLLMs require up to twice the token length in diffusion steps to match the performance and output length of autoregressive models when generating over 256 tokens. It indicates that effective reasoning entails roughly double the diffusion steps relative to the answer length, underscoring a notable efficiency challenge in reasoning tasks.

**In reasoning scenarios, a large number of diffusion steps for autoregressive reasoning is unavoidable for acceptable accuracy.** Each generated token requires a sufficient number of diffusion iterations to allow the model to reason effectively and produce high-quality outputs. As illustrated in Figure 3 (b), performance sharply declines when diffusion steps fall below the target token length. For example, generating 80 tokens with a maximum length of 128 but only 64 diffusion steps results in over a 10% accuracy drop; with 32 steps, accuracy decreases by about 40%. This demonstrates that inadequate diffusion severely impairs reasoning, as the model lacks enough refinement iterations. Thus, diffusion steps should at least match the planned token length to maintain reasoning quality. Nonetheless, excessive diffusion can significantly reduce efficiency.

### 3.3 RETHINKING THE PROMPTING STRATEGIES IN DLLMs FROM PSC PERSPECTIVE

In general, traditional autoregressive inference methods are typically categorized into two types: pipeline-guided approaches and condition-following approaches (see Appendix F for details). In this section, we will begin by reviewing the theoretical foundations and representative implementations of these two categories. We will then examine their practical limitations and challenges. Furthermore, we introduce a parallel-encouraging prompting to improve DLLM effectiveness.

**Sequential Reasoning Prompting will enlarge PSC’s negative impact for DLLMs.** Sequential prompting strategies, which facilitate sequential reasoning, have been shown to significantly improve the performance of ALLMs on complex tasks. However, as indicated in of Table 1, we observed a notable decline in performance as tasks required an increasing number of reasoning steps. We attribute this decline to the fact that sequential reasoning prompts exacerbate the negative impact of PSC, thereby impairing the reasoning performance in DLLMs. **Additionally, we conducted evaluations under high computational budgets. The results in Appendix F6 indicate that even with sufficient inference resources, performance remains limited.**

270 **Constraint-guided Reasoning Prompting enhances model performance by preventing the introduction of additional PSC.** By incorporating explicit constraints into the reasoning process, constraint-guided prompting effectively narrows the model’s search space, thereby preventing the emergence of additional PSC during the reasoning process in DLLMs. This focused approach results in more accurate and reliable solutions. As shown [blue rows](#) of in Table 1, methods based on this principle, such as Complex-CoT (Fu et al., 2022) and MARP (Chen et al., 2024), demonstrate superior reasoning capabilities in DLLMs compared to traditional sequential prompting methods.

271 **Parallel-encouraging Prompting reduces the sequential feature so that it further improves performance.** Parallel-encouraging prompting refers to the technique of presenting multiple related tasks or questions simultaneously. This approach reduces the impact of PSC and minimizes the sequential features in the prompting process. By encouraging the model to make connections across these tasks, as illustrated in [green rows](#) of Table 1, it effectively fosters DLLMs’ performance, leading to more efficient reasoning and information integration ([See more examples in Appendix F.4](#)). Leveraging the parallel processing capabilities of DLLMs, this method has the potential to significantly enhance performance, particularly in complex reasoning tasks, by promoting more comprehensive and coherent solutions ([See more analysis in Appendix F.5](#)). [To inspire concrete ideas for a DLLM-friendly dataset, we discuss this in Appendix F.8 to guide future work.](#)

### Takeaways

1. Due to PSC, DLLMs engage in superficial parallel reasoning and exhibit autoregressive behavior in complex scenarios, which compromises their reasoning efficiency.
2. Sequential prompts prove ineffective for DLLMs, requiring PSC-free or [PSC-reduced](#) approaches like constraint-guided and parallel-encouraging prompts to guide their operation.

## 4 WHAT CHALLENGES DO DLLMs MEET IN LONG COT BASED ON PSC?

295 Despite impressive empirical results, DLLMs’ genuine reasoning abilities and scalability under Parallel-Sequential Contradictions remain open questions. We systematically evaluate Long CoT to 296 assess these fundamental capabilities and scaling strategies.

### 4.1 DLLMs DO NOT HAVE SUFFICIENT BASIC CAPABILITIES TO SUPPORT LONG COT.

301 Long CoT is the primary innovation in recent reasoning large language models, leveraging inference-time 302 scaling for self-exploration, self-reflection, and deep reasoning (Chen et al., 2025). Evaluation 303 details are in Appendix G and Table 5.

306 **Traditional reflection strategies are Ineffective for DLLMs.** Long CoT models always employ a 307 self-reflection mechanism for iterative reasoning refinement. To assess its efficacy, we examine two 308 LLM paradigms: (1) Prompting Reflection and (2) Autoregressive Forcing Reflection. As shown in 309 Figure 4 (a, b), reflection paradigms yield no significant differences from vanilla reasoning chains in 310 semantic similarity, informativeness, or token-level entropy<sup>1</sup>. Though the reflection process increases 311 entropy and reduces informativeness, it maintains over 0.95 semantic similarity to original reasoning 312 chains. These findings suggest the reflection mechanism offers only limited surface-level optimization. 313 Figure 4 (c) further reveals a substantial token repetition ratio compared to the original path, resulting 314 in approximately 10% reflection-to-error responses.

315 **Limited Efficacy of traditional exploration strategies for novel reasoning path generation.** 316 Exploration, a fundamental competency for complex reasoning, involves a model’s ability to generate 317 diverse and innovative solutions. To assess this potential in DLLMs, we designed experiments utilizing 318 two strategies: (1) Prompting Exploration and (2) Autoregressive Forcing Exploration. Figure 5 319 (a, b) reveal that current exploration strategies offer several improvements in the novel semantic of 320 generated reasoning processes. However, these improvements remain superficial, evidenced by a 321 high similarity ( $> 0.84$ ) between explored paths and original results. Furthermore, as depicted in 322 Figure 5 (c), while the new path and explore-to-correct ratios are limited ( $\sim 5\%$ ), they nonetheless 323 indicate a positive, albeit constrained, effect.

<sup>1</sup>An effective reflection is generally expected to drive model toward lower entropy and higher certainty.

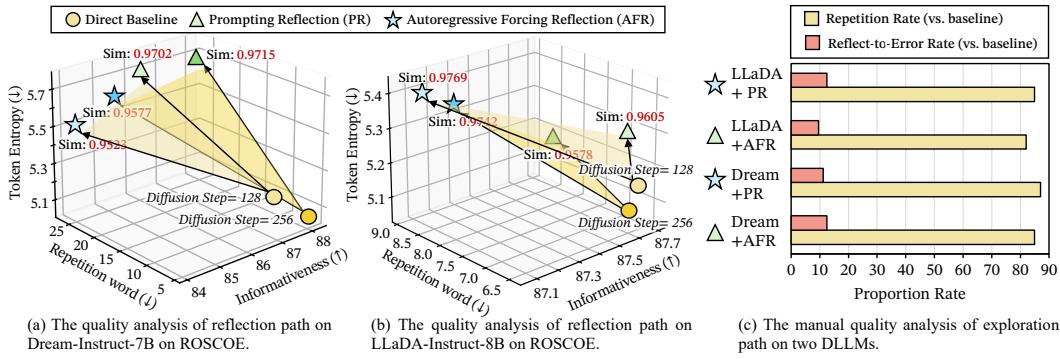


Figure 4: Self-reflection performance and rationale quality evaluation on DLLMs. In Figure (c), We conduct manual annotation for each reflection or exploration trajectory to determine whether it is repetitive compared to the preceding (baseline) trajectory (True/False) and whether its final outcome is correct (True/False), thereby obtaining the proportion of new paths or erroneous paths.

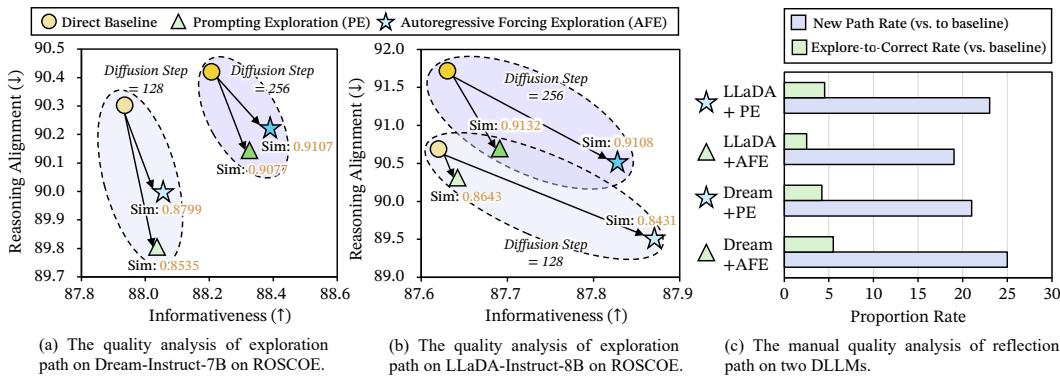


Figure 5: Self-exploration performance and rationale quality evaluation on DLLMs.

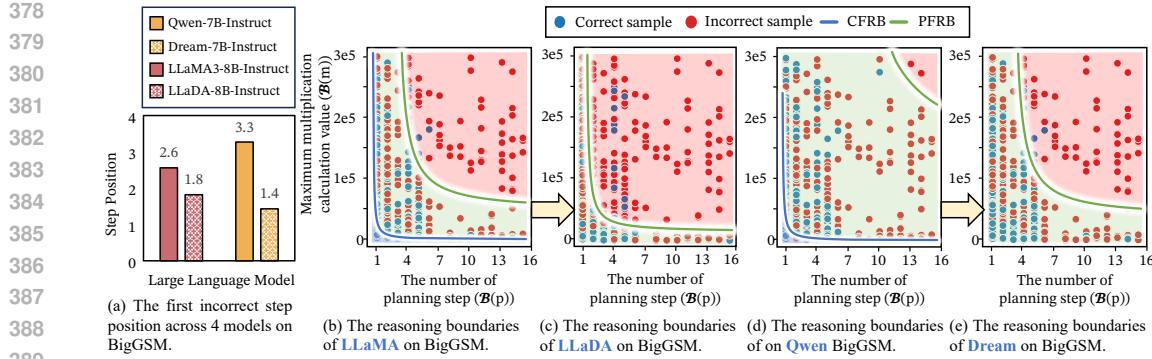
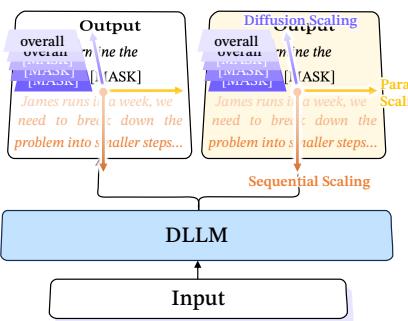
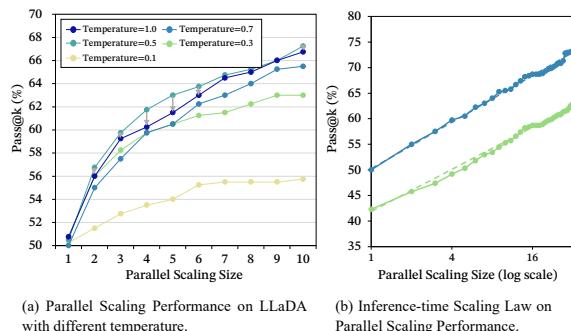
**DLLMs possess limited reasoning boundaries and, consequently, exhibit restricted deep reasoning abilities.** To examine the limitations of DLLMs on deep reasoning, we evaluate their capacity to sustain reasoning across sufficient depths. Figure 6 (a) demonstrates that error steps are all less than 2, which suggests that current DLLMs are unable to consistently sustain deep reasoning performance. Furthermore, following Chen et al. (2024), we define the 90% correctness step count as the models’ completely feasible reasoning boundary (CFRB), and the 10% correctness step count as the completely infeasible reasoning boundary (CIRB). As shown in Figure 6 (b), current DLLMs display similar CFRB values but lower CIRB values, indicating narrower feasible reasoning ranges.

#### 4.2 CURRENT DLLMs HAVE THREE-DIRECTIONAL BUT LIMITED INFERENCE-TIME SCALING

Given their denoising characteristics, we investigate a fundamental question: **Is there also Inference-Time Scaling Law in DLLM under such contradictions?** As shown in Figure 7, we examine this through three complementary perspectives:

- **Diffusion Scaling:** Increasing the number of denoising reasoning steps in the model. The stacked MASK blocks represent multiple, deeper iterations of denoising.
- **Parallel Scaling:** Generating multiple distinct output candidates simultaneously in parallel. The adjacent output blocks at the top represent the concurrent production of multiple output variants.
- **Sequential Scaling:** The visualization is designed to emphasize the step-by-step output generation process. The arrow direction explicitly illustrates this incremental reasoning progression.

These experiments determine whether DLLMs follow inference-time scaling laws and provide practical insights for optimizing reasoning performance. Implementation details can be seen in Appendix H.

390 Figure 6: Incorrect Step and Reasoning Boundaries Distribution of DLLMs on BigGSM.  
391403 Figure 7: Three primary scaling directions  
404 for DLLMs proposed in our work.  
405406 Figure 8: Performance analysis under Parallel Scaling.  
407

## 4.2.1 PARALLEL SCALING LAW HOLDS UNDER PSC

408 For DLLMs, a key question is whether their unique diffusion generation mechanism supports efficient  
409 parallel sampling and whether parallel sampling can effectively enhance reasoning performance.

410 **Higher temperatures do not always yield more diverse and effective parallel sampling.** The  
411 decoding temperature controls generation randomness, with higher values typically increasing output  
412 diversity in ALLM reasoning. We adjust the temperature during generation (0.1 to 1.0) to evaluate  
413 its impact on parallel sampling. Model accuracy improves steadily with increasing Pass@k values  
414 across all temperature settings before plateauing. As shown in Figure 8 (a), moderate temperatures  
415 (e.g.,  $T = 0.5$ ) achieve optimal performance, while both lower and higher temperatures yield  
416 diminished performance gains. **This result confirms that DLLMs also obey the widely recognized  
417 trade-off governed by temperature.**

418 **DLLM reasoning accuracy improves with increased parallel samples, following inference-time  
419 scaling patterns.** As shown in Figure 8 (b), when  $k$  increases from 1 to 32, accuracy demonstrates  
420 nearly linear improvement on a logarithmic scale. This indicates that DLLMs effectively utilize  
421 parallel sampling to enhance reasoning performance, as diverse outputs increase the probability of  
422 generating correct solutions. This pattern aligns with inference-time scaling laws observed in other  
423 advanced language models, where performance scales with computational effort during inference.

## 424 4.2.2 DIFFUSION SCALING IS CONSTRAINED BY PSC

426 **Diffusion Scaling of DLLMs ensures performance gains, with diffusion time positively corre-  
427 lated.** Model accuracy increases monotonically with the number of diffusion steps. We are the first  
428 to formalize *Diffusion Scaling* in DLLMs, proposing a positive correlation between model perfor-  
429 mance and diffusion iterations. To validate this claim, we benchmark two representative DLLMs,  
430 DREAM (Ye et al., 2025) and LLaDA (Nie et al., 2025a), under an exponential schedule of diffusion  
431 steps (1–1024). By tracking accuracy at each step, we observe how DLLMs address reasoning tasks of  
432 varying complexity across the diffusion process. As shown in Figure 9 (a), performance consistently

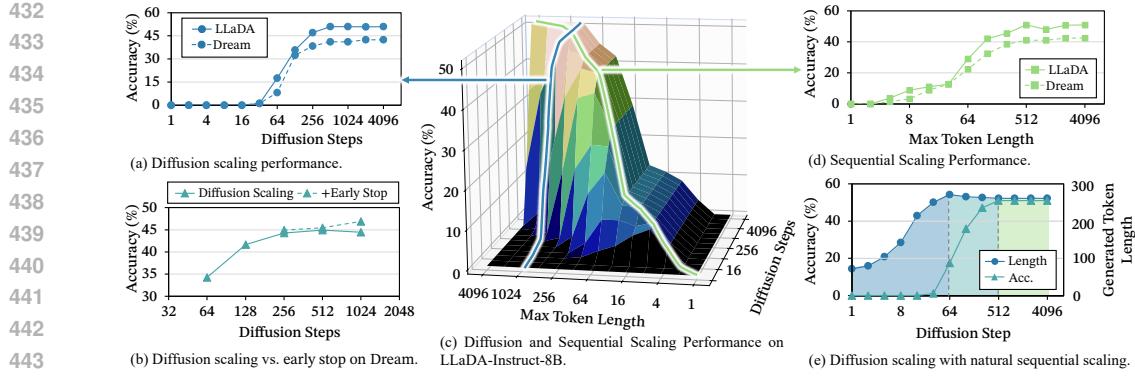


Figure 9: Diffusion Scaling analysis of reasoning accuracy across difficulty levels on BigGSM.

improves with deeper diffusion; however, the rate of improvement depends on task difficulty: simpler problems gain substantially, while tasks beyond the model’s capacity yield only limited benefits.

**Diffusion Scaling is effective and exhibits an upper bound, beyond which an over-diffusion phenomenon emerges due to PSC.** Consistent with classical scaling laws, the benefits of diffusion scaling are inherently capped. As shown in Figure 9 (a), increasing diffusion steps improves performance from 32 to 512 steps, after which gains plateau. More importantly, excessive diffusion reduces accuracy: Figure 9 (b) shows a drop from 44.92% to 44.43%. This decline illustrates *over-diffusion*, where extended denoising introduces excessive corrections that disrupt reasoning chains, akin to overfitting caused by training without early stopping.

**Early stopping can effectively mitigate over-diffusion.** To address over-diffusion, we propose a Diffusion Early Stopping (DES) strategy that halts the process when generated tokens stabilize. The implementation comprises three components: (1) **Overlap Ratio Calculation**: computed as the proportion of identical tokens between consecutive steps. (2) **Convergence Detection**: potential convergence occurs when the overlap ratio meets or exceeds a predefined threshold. (3) **Activation Condition**: early stopping triggers only after three consecutive steps satisfy the threshold, preventing false positives from transient fluctuations. As shown in Figure 9 (b), we observe that beyond 256 steps, early stopping outperforms standard diffusion, with accuracy improving from 44.26% to 46.89% at 1024 steps. Early stopping captures convergence states and terminates upon stabilization, while preventing performance degradation from excessive diffusion.

#### 4.2.3 SEQUENTIAL SCALING IS EQUALLY CONSTRAINED BY PSC

**The inherent limitations of sequential scaling for DLLMs.** While sequential scaling has shown promise in enhancing reasoning capabilities, it remains constrained by the inherent characteristics of DLLMs. **Unlike traditional nearly unbounded diffusion scaling on pure diffusion models (Ma et al., 2025)**, as shown in Figure 9 (d), the performance improvements are at first increasing but eventually converge. This limitation arises from the fact that sequential scaling relies on the model’s ability to maintain context over extended reasoning chains, a challenge for current DLLMs.

**Sequential Scaling also meets over-thinking challenges.** Similar to diffusion scaling, as shown in Figure 9 (d), sequential scaling faces its own set of over-thinking challenges. As the model attempts to extend its reasoning across longer contexts, it may encounter diminishing returns or even performance degradation. This phenomenon is particularly evident in tasks that require intricate reasoning over extended text, where the model’s ability to track and integrate information can become strained.

**Diffusion Scaling can naturally yield Sequential Scaling benefits.** As shown in Figure 9 (e), diffusion scaling alleviates the limitations of sequential scaling. We identify three stages in DLLMs during diffusion: (1) sequential scaling, (2) compression, and (3) convergence. In the first stage, increasing diffusion steps leads to stable performance but longer solutions, indicating that DLLMs explore suitable lengths for reasoning. In the second stage, the model compresses its reasoning by eliminating redundancy, generating more efficient solutions. In the third stage, the model converges on an optimal strategy, achieving high performance while reducing computational cost.

486

487

488

489

490

491

492

493

494

495

496

**Takeaways**

1. DLLMs are deficient in three basic Long CoT capabilities, hindering their effectiveness.
2. DLLM can be optimized via parallel, diffusion, and sequential scaling. Diffusion scaling inherently encompasses the benefits of sequential scaling.
3. The performance of both diffusion and sequential scaling is ultimately upper-bounded by a parallel-sequential contradiction. But Parallel scaling law remains the most effective strategy, although it is also the most computationally expensive.

## 5 RELATED WORK

The application of diffusion models to text generation has emerged as an alternative to autoregressive methods. Early work by D3PM (Austin et al., 2021) proposed discrete denoising diffusion probabilistic models, and Diffusion-BERT (He et al., 2022) demonstrated scalability to BERT-style architectures. SEDD (Lou et al., 2023) achieved performance comparable to GPT-2. Recent progress has broadened the scope of Diffusion Large Language Models (DLLMs) (Yang et al., 2025; Wu et al., 2025; Gong et al., 2025). LLaDA (Nie et al., 2025a) and Dream (Ye et al., 2025) scaled to billion-parameter models with notable inference gains. The D2F strategy (Wang et al., 2025a) further enhanced inference by enabling block-level autoregression and parallel decoding, maintaining a balance between speed and accuracy. This direction aligns with the growing interest in applying DLLMs to extended reasoning (Wang et al., 2025b; Zhao et al., 2025). Diffusion-of-Thought (DoT) (Ye et al., 2024) combines diffusion with chain-of-thought reasoning. Building on this, Zhao et al. (2025) and Tang et al. (2025) applied diffusion-augmented SFT and GRPO to strengthen reasoning. Similarly, Trado (Wang et al., 2025b) exploits overlooked sampling signals, yielding further reasoning gains.

However, while DLLMs exhibit notable parallel decoding in text generation and consistently strong step-by-step reasoning, these features appear conceptually opposed: parallelism implies simultaneous processing, whereas sequential reasoning demands ordered progression. This apparent Parallel–Sequential Contradiction (PSC) suggests that both the underlying mechanisms and the practical effectiveness of DLLMs’ diffusion-based reasoning remain insufficiently understood.

## 6 CONCLUSION

In this work, we formalize the Parallel–Sequential Contradiction (PSC) to explain why DLLMs, though built for parallel decoding, revert to autoregression as reasoning difficulty rises. Empirically, DLLMs exploit parallelism only when tokens are locally decidable; otherwise, they fall back to sequential computation, reducing efficiency. Further, we first define three-dimensional scaling: parallel, diffusion, and sequential scaling, and show that PSC restricts the latter two while parallel scaling holds. We mitigate PSC through parallel-focused prompting, diffusion early stopping, and parallel scaling, improving both accuracy and throughput. Future work should align training and architectures with PSC-aware reasoning and design benchmarks, isolating its effects.

## REPRODUCIBILITY DISCUSSION

Since our paper is an analytical paper, we directly use the official DLLMs warehouse for inference. Except for the Scaling experiment, fast-dllm was added for time spent, and the others are used for inference using official inference codes, and only relevant parameters are adjusted. In addition, the relevant prompt experiments are all in Appendix F.

## REFERENCES

Shivam Agarwal, Zimin Zhang, Lifan Yuan, Jiawei Han, and Hao Peng. The unreasonable effectiveness of entropy minimization in llm reasoning. *arXiv preprint arXiv:2505.15134*, 2025.

Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in neural information processing systems*, 34:17981–17993, 2021.

540 Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge,  
 541 Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu,  
 542 Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanchi Tan,  
 543 Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin  
 544 Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng  
 545 Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou,  
 546 Xiaohuan Zhou, and Tianhang Zhu. Qwen technical report. *arXiv preprint arXiv:2309.16609*,  
 547 2023.

548 Qiguang Chen, Libo Qin, Jiaqi Wang, Jingxuan Zhou, and Wanxiang Che. Unlocking the capabilities  
 549 of thought: A reasoning boundary framework to quantify and optimize chain-of-thought. *Advances*  
 550 *in Neural Information Processing Systems*, 37:54872–54904, 2024.

551 Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu,  
 552 Yuhang Zhou, Te Gao, and Wanxiang Che. Towards reasoning era: A survey of long chain-of-  
 553 thought for reasoning large language models. *arXiv preprint arXiv:2503.09567*, 2025.

554 Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen  
 555 Fan, Huayu Chen, Weize Chen, et al. The entropy mechanism of reinforcement learning for  
 556 reasoning language models. *arXiv preprint arXiv:2505.22617*, 2025.

557 Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. Complexity-based prompting  
 558 for multi-step reasoning. *arXiv preprint arXiv:2210.00720*, 2022.

559 Olga Golovneva, Moya Chen, Spencer Poff, Martin Corredor, Luke Zettlemoyer, Maryam Fazel-  
 560 Zarandi, and Asli Celikyilmaz. ROSCOE: A suite of metrics for scoring step-by-step reasoning.  
 561 2022.

562 Shansan Gong, Ruixiang Zhang, Huangjie Zheng, Jiatao Gu, Navdeep Jaitly, Lingpeng Kong, and  
 563 Yizhe Zhang. Diffucoder: Understanding and improving masked diffusion models for code  
 564 generation. *arXiv preprint arXiv:2506.20639*, 2025.

565 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,  
 566 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms  
 567 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

568 Zhengfu He, Tianxiang Sun, Kuanning Wang, Xuanjing Huang, and Xipeng Qiu. Diffusion-  
 569 bert: Improving generative masked language models with diffusion models. *arXiv preprint*  
 570 *arXiv:2211.15029*, 2022.

571 Zemin Huang, Zhiyang Chen, Zijun Wang, Tiancheng Li, and Guo-Jun Qi. Reinforcing the diffusion  
 572 chain of lateral thought with diffusion language models. *arXiv preprint arXiv:2505.10446*, 2025.

573 Inception Labs. Mercury: A diffusion large language model. Technical report, Inception Labs, 2025.  
 574 URL <https://www.inception-labs.ai/mercury>. Commercial-grade diffusion LLM  
 575 for code generation. Achieves over 1000 tokens/second on NVIDIA H100.

576 Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large  
 577 language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems*,  
 578 volume 35, pp. 22199–22213, 2022.

579 Tianyi Li, Mingda Chen, Bowei Guo, and Zhiqiang Shen. A survey on diffusion language models.  
 580 *arXiv preprint arXiv:2508.10875*, 2025.

581 Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios  
 582 of the data distribution. *arXiv preprint arXiv:2310.16834*, 2023.

583 Nanye Ma, Shangyuan Tong, Haolin Jia, Hexiang Hu, Yu-Chuan Su, Mingda Zhang, Xuan Yang, Yan-  
 584 dong Li, Tommi Jaakkola, Xuhui Jia, and Saining Xie. Inference-time scaling for diffusion models  
 585 beyond scaling denoising steps, 2025. URL <https://arxiv.org/abs/2501.09732>.

586 Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-  
 587 Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv preprint arXiv:2502.09992*,  
 588 2025a.

594 Yatao Nie, Jie Chen, Yufan Zhang, et al. Large language diffusion with masking. *arXiv preprint*,  
 595 2025b. URL <https://arxiv.org/abs/2502.09992>.

596

597 Libo Qin, Qiguang Chen, Fuxuan Wei, Shijue Huang, and Wanxiang Che. Cross-lingual  
 598 prompting: Improving zero-shot chain-of-thought reasoning across languages. *arXiv preprint*  
 599 *arXiv:2310.14799*, 2023.

600 Xiaohang Tang, Rares Dolga, Sangwoong Yoon, and Ilija Bogunovic. wd1: Weighted policy  
 601 optimization for reasoning in diffusion language models. *arXiv preprint arXiv:2507.08838*, 2025.

602

603 Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim.  
 604 Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language  
 605 models. *arXiv preprint arXiv:2305.04091*, 2023.

606 Xu Wang, Chenkai Xu, Yijie Jin, Jiachun Jin, Hao Zhang, and Zhijie Deng. Diffusion llms can do  
 607 faster-than-ar inference via discrete diffusion forcing, aug 2025a. URL <https://arxiv.org/abs/2508.09192>. arXiv:2508.09192.

608

609 Yinjie Wang, Ling Yang, Bowen Li, Ye Tian, Ke Shen, and Mengdi Wang. Revolutionizing reinforce-  
 610 ment learning framework for diffusion large language models. *arXiv preprint arXiv:2509.06949*,  
 611 2025b.

612

613 Chengyue Wu, Hao Zhang, Shuchen Xue, Zhijian Liu, Shizhe Diao, Ligeng Zhu, Ping Luo, Song  
 614 Han, and Enze Xie. Fast-dllm: Training-free acceleration of diffusion llm by enabling kv cache  
 615 and parallel decoding. *arXiv preprint arXiv:2505.22618*, 2025.

616

617 Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. C-pack: Packaged resources to  
 618 advance general chinese embedding, 2023.

619

620 Ling Yang, Ye Tian, Bowen Li, Xinchen Zhang, Ke Shen, Yunhai Tong, and Mengdi Wang. Mmada:  
 621 Multimodal large diffusion language models. *arXiv preprint arXiv:2505.15809*, 2025.

622

623 Jiacheng Ye, Shansan Gong, Liheng Chen, Lin Zheng, Jiahui Gao, Han Shi, Chuan Wu, Xin Jiang,  
 624 Zhenguo Li, Wei Bi, et al. Diffusion of thought: Chain-of-thought reasoning in diffusion language  
 625 models. *Advances in Neural Information Processing Systems*, 37:105345–105374, 2024.

626

627 Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng  
 628 Kong. Dream 7b: Diffusion large language models. *arXiv preprint arXiv:2508.15487*, 2025.

629

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

632

633 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans,  
 634 Claire Cui, Olivier Bousquet, Quoc Le, et al. Least-to-most prompting enables complex reasoning  
 635 in large language models. *arXiv preprint arXiv:2205.10625*, 2022.

636

637

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

641

642

643

644

645

646

647

# 648 649 650 651 652 653 654 655 656 657 658 659 660 661 662 663 664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680 681 682 683 684 685 686 687 688 689 690 691 692 693 694 695 696 697 698 699 700 701 Appendix

## A THE USE OF LARGE LANGUAGE MODELS

In preparing this manuscript, large language models were only utilized as general-purpose writing assistants. Their role was confined to improving the clarity and refining the phrasing of the text. All scientific content, analyses, and core arguments were developed by the human authors, who take full responsibility for the final version of the paper.

## B MATHEMATICAL PROOF OF DLLM DEGRADING TO AUTOREGRESSIVE

**Goal.** We rigorously show, via information theory and optimization, that the intrinsic statistical property of a generative task, namely its sensitivity to perturbations of initial conditions, fundamentally determines its optimal (lowest-loss) generation strategy. Concretely, we axiomatize two classes of tasks: serial tasks (step-by-step reasoning) exhibiting cascading sensitivity to initial conditions, and parallel tasks exhibiting partial invariance, and we prove that serial tasks induce significantly higher conditional entropy for “skip-step” parallel predictions  $S_k \mid S_1$  than parallel tasks, forcing any loss-minimizing learner to converge to an autoregressive strategy.

### B.1 PROBLEM SETUP AND NOTATION

Let all step states  $S_i$  take values in a metric space  $(\Omega, d)$ .

Specifically, the true data distribution  $p_\theta^S$  of a task is considered serial if, for any given  $s_1 \in \Omega$ , the mapping from  $s'_1$  to a subsequent state  $s'_k$  is highly divergent within a sufficiently small neighborhood  $N(s_1, \epsilon) = \{s'_1 \mid d(s_1, s'_1) < \epsilon\}$ . Conversely, a task’s true data distribution  $p_\theta^P$  is considered parallel if, for any given  $s_1 \in \Omega$ , there exists at least one subsequent state  $S_k$  (where  $k > 1$ ) that is insensitive to perturbations within its neighborhood  $N(s_1, \epsilon)$ . Formally, this leads to the following definitions for the two tasks.

**Definition 1** (Serial tasks: cascading sensitivity with locally continuous transitions). *A data-generating distribution  $p_\theta^S$  is serial if it satisfies the following properties:*

**Local continuity and learnable short-term dynamics:** *For each time step  $t$ , the conditional distribution  $p_\theta^S(S_{t+1} \mid S_t = s_t)$  is locally continuous: there exists  $\delta_t > 0$  such that for all  $s'_t \in N(s_t, \delta_t)$ , the distribution  $p_\theta^S(S_{t+1} \mid S_t = s'_t)$  has low entropy and is learnable via autoregression.*

**Sensitive long-term dependence:** *For any  $s_1 \in \Omega$  and any  $k > 1$ , the compounded long-range mapping  $s_1 \rightarrow s_k$  exhibits high sensitivity, satisfying:*

$$\lim_{\epsilon \rightarrow 0} \mathbb{E}_{s'_1 \sim U(N(s_1, \epsilon))} [p_\theta^S(S_k = s_k \mid S_1 = s'_1)] = 0, \quad (3)$$

*where  $s_k$  is the reference outcome drawn from the true conditional  $p_\theta^S(S_k \mid S_1 = s_1)$ , and  $U(N(s_1, \epsilon))$  is the uniform distribution over the  $\epsilon$ -neighborhood of  $s_1$ .*

*This captures the dichotomy between locally continuous and learnable short-term transitions versus sensitive long-term dependence on initial conditions.*

**Definition 2** (Parallel tasks: partial invariance). *A data-generating distribution  $p_\theta^P$  is parallel if there exists some  $k > 1$  and a constant  $C$  such that for any  $s_1 \in \Omega$ , the following condition satisfies:*

$$\lim_{\epsilon \rightarrow 0} \mathbb{E}_{s'_1 \sim U(N(s_1, \epsilon))} [p_\theta^P(S_k = s_k \mid S_1 = s'_1)] = C, \quad (4)$$

*where  $s_k$  is the reference outcome drawn from  $p_\theta^P(S_k \mid S_1 = s_1)$ , and the constant  $C$  satisfies  $0 < C \leq 1$ . Thus, a structurally stable downstream state persists with significant probability despite infinitesimal perturbations of the initial condition.*

### B.2 LEARNING PROBLEM

Let  $p_\theta$  be a parametric generative model trained by minimizing cross-entropy with respect to the true data distribution  $\hat{p}$ , i.e.,

$$L(p_\theta, \hat{p}) = \mathbb{E}_{x \sim \hat{p}} [-\log p_\theta(x)] = H(\hat{p}) + D_{\text{KL}}(\hat{p} \parallel p_\theta), \quad (5)$$

so minimizing cross-entropy is equivalent to minimizing  $D_{\text{KL}}(\hat{p} \| p_{\theta})$  and to maximum likelihood. For any conditional subproblem, the optimum satisfies  $p_{\theta}^*(S_k | S_1) = \hat{p}(S_k | S_1)$ , and the minimal expected negative log-likelihood equals the conditional entropy,

$$L^* = \mathbb{E}_{s_1 \sim \hat{p}(S_1)} \left[ H(\hat{p}(S_k | S_1 = s_1)) \right], \text{ where } H(Y | X) \equiv - \sum_{x \in X, y \in Y} p(x, y) \log \frac{p(x, y)}{p(x)}. \quad (6)$$

**Autoregression and chain rule.** A step-by-step reasoning strategy factorizes a joint distribution as a product of conditionals via the chain rule, thereby replacing a high-entropy “skip” conditional  $p_{\theta}(S_k | S_1)$  by a sequence of typically lower-entropy one-step conditionals  $p_{\theta}(S_{t+1} | S_t, \dots)$ . It is the standard rationale behind likelihood-based training of ALMs under teacher forcing.

### B.3 DISCRETIZATION

To compare entropies on a general metric space, consider a finite measurable partition  $\Pi_{\varepsilon}$  of  $\Omega$  with mesh size at most  $\varepsilon$ , and define the discretization and quantization map  $\phi_{\varepsilon} : \Omega \rightarrow [m_{\varepsilon}]$  that assigns each  $s \in \Omega$  to its cell index, where  $m_{\varepsilon} = |\Pi_{\varepsilon}|$ . Let  $\tilde{S}_i^{(\varepsilon)} = \phi_{\varepsilon}(S_i)$  and write  $p_{\theta}^{(\varepsilon)}(\cdot)$  for the induced discrete laws; we analyze  $H(\tilde{S}_k^{(\varepsilon)} | \tilde{S}_1^{(\varepsilon)} = \tilde{s}_1)$ , which is well-defined, and relate back to the original problem by taking  $\varepsilon \rightarrow 0$ . Two standard facts underpin the analysis: (i) for a fixed finite support, entropy is maximized by the uniform distribution; (ii) the Shannon entropy is bounded below by the min-entropy  $-\log p_{\max}$ , and admits tighter lower bounds in terms of the binary entropy function  $H_b$  and the support size.

**Lemma 1** (Pointwise probability caps). *Fix  $s_1 \in \Omega$ . Under Definition 1, for any  $k > 1$  and any  $\eta > 0$  there exists  $\delta_0 > 0$  such that for all discretization mesh sizes  $\delta < \delta_0$  and for all perturbation radii  $\varepsilon < \delta$ ,*

$$\max_{s'_1 \in N(s_1, \varepsilon)} p_{\theta, S}^{(\delta)}(\tilde{S}_k^{(\delta)} = s_k | \tilde{S}_1^{(\delta)} = \phi_{\delta}(s'_1)) \leq \eta, \quad (7)$$

where  $N(s_1, \varepsilon) = \{s'_1 | d(s_1, s'_1) < \varepsilon\}$  and  $\phi_{\delta}$  is the discretization map with mesh size  $\delta$ .

*Under Definition 2, there exist some  $k > 1$ ,  $C > 0$  and  $\delta_0 > 0$  such that for all  $\delta < \delta_0$  and  $\varepsilon < \delta$ ,*

$$\max_{s'_1 \in N(s_1, \varepsilon)} p_{\theta, P}^{(\delta)}(\tilde{S}_k^{(\delta)} = s_k | \tilde{S}_1^{(\delta)} = \phi_{\delta}(s'_1)) \geq C. \quad (8)$$

*Proof sketch.* By Definition 1, for serial tasks, the conditional probability of a reference outcome, averaged over shrinking neighborhoods of  $s'_1$ , vanishes. This forces any mass that could be concentrated on a particular cell containing  $s_k$  to diminish as the mesh refines. In contrast, for parallel tasks, Definition 2 guarantees a persistent mass  $C \in (0, 1]$  associated with a stable outcome across neighborhoods, which uniformly lower-bounds the maximum conditional atom in  $\varepsilon$ .

### B.4 MAIN PROPOSITION AND QUANTITATIVE BOUNDS

**Proposition 1** (Skip-step parallel predictions on serial vs. parallel tasks). *For any  $k > 1$ , the optimal expected skip-prediction loss on serial-task data strictly exceeds that on parallel-task data:*

$$L_S^*(p_{\theta}(S_k | S_1 = s'_1)) > L_P^*(p_{\theta}(S_k | S_1 = s'_1)), \quad (9)$$

Let  $H_S$  and  $H_P$  denote the conditional entropies under the Serial and Parallel data distributions, respectively. We have,

$$\mathbb{E}_{s'_1 \sim U(N(s_1, \varepsilon))} [H_S(S_k | S_1 = s'_1)] > \mathbb{E}_{s'_1 \sim U(N(s_1, \varepsilon))} [H_P(S_k | S_1 = s'_1)], \quad (10)$$

In the discrete case, this reduces to showing

$$\mathbb{E}_{s'_1 \sim N'(s_1)} [H_S(\tilde{S}_k^{(\varepsilon)} | \tilde{S}_1^{(\varepsilon)} = \phi_{\varepsilon}(s_1))] > \mathbb{E}_{s'_1 \sim N'(s_1)} [H_P(\tilde{S}_k^{(\varepsilon)} | \tilde{S}_1^{(\varepsilon)} = \phi_{\varepsilon}(s_1))] \quad (11)$$

with a strictly positive gap that can be quantified through discretization and classical entropy bounds.

*Proof.* It suffices to compare the conditional entropies pointwise and then take expectations. Fix  $s_1$  and a partition  $\Pi_{\varepsilon}$ . We first define the maximum of generation probability of serial tasks:

756

$$757 \quad p_{\max}^S(\delta; s_1, k) := \max_{s'_1 \in N(s_1, \varepsilon)} p_{\theta, S}^{(\delta)}(\tilde{S}_k^{(\delta)} = s_k \mid \tilde{S}_1^{(\delta)} = \phi_\delta(s'_1)) \quad (12)$$

758

759 where the maximization is performed over the neighborhood of  $s_1$  for a fixed target step  $k$ , and  
 760 analogously  $p_{\max}^P(\delta; s'_1)$  under  $p_{\theta, P}^{(\delta)}$ .

761 By the lemma,  $p_{\max}^S(\delta; s'_1) \rightarrow 0$  as  $\delta \rightarrow 0$ , while for parallel tasks one has  $p_{\max}^P(\delta; s'_1) \geq C$  for all  
 762 sufficiently small  $\delta$ . For any discrete distribution over  $m_\delta$  points with maximal atom  $p_{\max}$ , Fano's  
 763 inequality implies

764

$$H \leq H_b(p_{\max}) + (1 - p_{\max}) \log(m_\delta - 1), \quad (13)$$

765

766 where  $H_b$  is the binary entropy.

767 Thus, for parallel tasks,

768

$$H_P(\tilde{S}_k^{(\delta)} \mid \tilde{S}_1^{(\delta)} = \phi_\delta(s'_1)) \leq H_b(p_{\max}^P(\delta; s'_1)) + (1 - p_{\max}^P(\delta; s'_1)) \log(m_\delta - 1), \quad (14)$$

769

770 and since  $p_{\max}^P(\delta; s'_1) \geq C > 0$ , the entropy is uniformly bounded away from the maximal value  
 771  $\log(m_\delta)$  by a constant determined by  $C$ .

772 For serial tasks, since  $p_{\max}^S(\varepsilon; s'_1) \rightarrow 0$ , we have error probability  $p_e \rightarrow 0$ . Now, we should apply the  
 773 contrapositive of Fano's inequality. Specifically, given the Fano's inequality:

774

$$H \leq H_b(p_e) + p_e \log(m_\varepsilon - 1), \quad (15)$$

775

776 it follows that  $H \rightarrow 0 \Rightarrow p_e \rightarrow 0$ . Conversely,  $p_e \rightarrow 1$  implies  $H \rightarrow H_{\max}$ . In this sense, the  
 777 condition is satisfied:

778

$$H_P(\tilde{S}_k^{(\varepsilon)} \mid \tilde{S}_1^{(\varepsilon)} = \phi_\varepsilon(s'_1)) \rightarrow \log(m_\varepsilon - 1), \text{ if } p_{\max}^S(\varepsilon; s'_1) \rightarrow 0, \quad (16)$$

779

780 which reflects the extreme dispersion dictated by sensitivity. Therefore, it satisfies:

781

$$\mathbb{E}_{s'_1 \sim N'(s_1)}[H_S(\tilde{S}_k^{(\varepsilon)} \mid \tilde{S}_1^{(\varepsilon)} = \phi_\varepsilon(s_1))] - \mathbb{E}_{s'_1 \sim N'(s_1)}[H_P(\tilde{S}_k^{(\varepsilon)} \mid \tilde{S}_1^{(\varepsilon)} = \phi_\varepsilon(s_1))] > 0. \quad (17)$$

782

783 Q.E.D. □

784

## 785 B.5 CONSEQUENCES FOR OPTIMAL STRATEGY

786

787 **Formal chain rule of entropy.** To rigorously establish why sequential generation is superior for  
 788 serial tasks, we invoke the chain rule of entropy:

789

$$H(S_k \mid S_1) = \sum_{t=1}^{k-1} H(S_{t+1} \mid S_{\leq t}). \quad (18)$$

790

791 While our proof shows that the long-range conditional entropy  $H(S_k \mid S_1)$  is high due to cascading  
 792 sensitivity, the **local continuity** assumption (revised Definition 1) ensures that each step-wise conditional  
 793 entropy  $H(S_{t+1} \mid S_{\leq t})$  remains low and learnable. The autoregressive strategy minimizes the  
 794 decomposed loss:

795

$$L_{\text{AR}} = \sum_{t=1}^{k-1} \mathbb{E}[-\log p(S_{t+1} \mid S_{\leq t})], \quad (19)$$

796

797 where each term has low entropy and is thus efficiently learnable. In contrast, direct parallel prediction  
 798 of  $p(S_k \mid S_1)$  must contend with the high entropy  $H(S_k \mid S_1)$ , requiring exponentially more samples to  
 799 achieve the same precision.

800 **Optimal strategy selection.** Since the minimum expected NLL equals the conditional entropy, the  
 801 high conditional entropy of skip-step predictions in serial tasks implies a high irreducible loss for any  
 802  $p_\theta(S_k \mid S_1)$  objective. A loss-minimizing learner thus prefers to factorize predictions into a sequence  
 803 of low-entropy one-step conditionals, i.e., an autoregressive strategy consistent with the chain-rule  
 804 factorization and standard maximum-likelihood training.

805 By contrast, in parallel tasks, the existence of a stable high-probability outcome for some downstream  
 806 state  $S_k$  produces a low-entropy, high-confidence conditional, so optimizing  $p_\theta(S_k \mid S_1)$  can be  
 807 preferable and can support non-autoregressive or partially parallel generation plans.

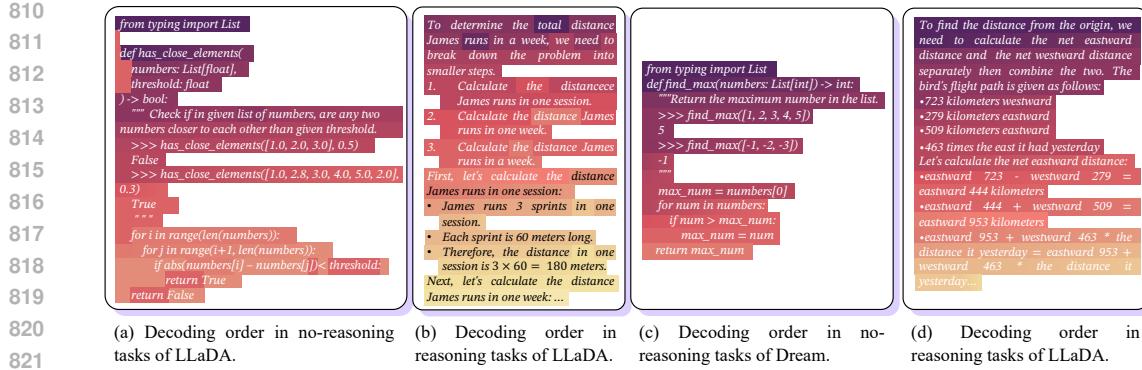


Figure 10: Decoding order of Dream and LLaDA on BigGSM (Chen et al., 2024).

**Takeaway.** Task-intrinsic sensitivity versus invariance dictates the shape of the optimal conditional distributions; via the cross-entropy/KL equivalence, this in turn selects the generation procedure that globally minimizes expected loss, with serial tasks forcing autoregression and parallel tasks permitting advantageous skip-step or parallel predictions.

## C EXPERIMENT SETTING

Due to their strong fundamental performance and stable generation, this study adopts Dream-7B-Instruct (Ye et al., 2025) and LLaDA-8B-Instruct (Nie et al., 2025b) as representative diffusion language models for most experimental evaluations.

In Section 3.1, to systematically examine how PSC affects DLLM reasoning, we select D1 (Zhao et al., 2025) and TraDo (Wang et al., 2025b) for diffusion-sequence analysis: D1 refines reasoning trajectories via reinforcement learning-based reward reshaping, whereas TraDo uses dynamic programming to coordinate denoising in latent space, thereby establishing a theoretical link between noise scheduling and reasoning stability.

In Section 3.3, when assessing the effects of different prompt strategies on DLLMs, we employ LLaDA-v1.5 (Zhu et al., 2025) and LLaDOU-Math (Huang et al., 2025). LLaDA-v1.5 improves mathematical and coding abilities through dynamic mask scheduling and hierarchical denoising, while LLaDOU-Math applies a reinforcement learning-based noise scheduling scheme and attains state-of-the-art performance on the MATH benchmark.

In Section 4.1, when probing the deep reasoning capabilities of DLLMs, we include the autoregressive models LLaMA3-8B-Instruct and Qwen-7B-Instruct (Bai et al., 2023) as baselines. These two models serve as mature autoregressive references, enabling a comprehensive comparison with DLLMs.

Comprehensive experimental configurations are documented in the appendices. For Section 3.2 (Diffusion-Step Dilemma), Section 3.3 (Prompting Strategies on DLLMs), Section 4.1 (Long CoT Capability of DLLMs), and Section 4.2 (Three-directional Inference-time Scaling on DLLMs), the corresponding details are provided in Appendix D, F, G, and H, respectively.

## D DIFFUSION-STEP EVALUATION DETAILS

In this section, we provide additional technical details on the methodology used to evaluate the impact of diffusion steps and sampling lengths in Diffusion-based Language Models (DLLMs). We specifically focus on how these parameters influence the efficiency and accuracy of the models when tackling complex reasoning tasks.

For our analysis, we utilize the BigGSM dataset (Chen et al., 2024), which includes a diverse range of complex reasoning tasks designed to test current models’ ability to perform long-form reasoning. In particular, we assess the performance of two representative DLLMs on these tasks and compare them against a standard ALLM. We systematically vary both the number of diffusion steps and the sampling lengths to evaluate their combined effects on the reasoning efficiency of DLLMs. The

864 number of diffusion steps tested ranges from 1 to 1024, while the maximum token lengths vary from  
 865 1 to 1024 tokens, with low-confidence remasking. In each experiment, the number of diffusion steps  
 866 is set equal to the maximum token length. This range allows us to assess the model’s performance  
 867 under different levels of token generation and diffusion refinement. For ALLMs, we adjust the  
 868 maximum token length between 1 and 1024 and the temperature between 0.2 and 0.7, aiming to  
 869 achieve comparable performance to that of the DLLMs.

870 For each setting, we track the following metrics:

871

- 872 • **Accuracy:** The percentage of correct answers generated by the model.
- 873 • **Model Output Length:** The number of tokens generated by the model before reaching the  
 874 stopping token (calculated using the GPT-4O tokenizer).

875 When the maximum token length is less than or equal to 512, the model output length typically  
 876 constitutes 50% to 80% of the maximum token length. Specifically, when generating a maximum  
 877 token length of 512, achieving optimal performance requires 512 diffusion steps combined with low-  
 878 confidence remasking strategies. We utilize this remasking approach to ensure the best performance.

879 Our evaluation demonstrates that the efficiency of DLLMs in reasoning tasks is strongly influenced  
 880 by the balance between diffusion steps and sampling lengths. While a higher number of diffusion  
 881 steps generally improves reasoning accuracy, it increases computational requirements. Thus, while  
 882 sufficient diffusion steps are essential for effective reasoning, an excessive number can significantly  
 883 reduce processing efficiency.

## 884 E THE IMPLEMENTATION OF EARLY-STOP STRATEGY

885 The early stopping mechanism is based on the dynamic stability of tokens, which monitors the  
 886 variation of the newly updated tokens during the diffusion process to judge whether the generation has  
 887 converged. We calculate the overlap ratio between the current step’s selected tokens (current\_tokens)  
 888 and the previous step’s tokens (prev\_tokens) in each diffusion step. When the overlap ratio of the  
 889 selected tokens remains stable over three consecutive steps and exceeds a threshold  $\theta = 0.99$ , early  
 890 stopping is triggered. The overlap ratio is calculated as:

$$894 \quad \text{overlap\_ratio} = \frac{1}{N} \sum_{j=1}^N \mathbb{I}(\text{current\_tokens}_j = \text{prev\_tokens}_j)$$

895 Where  $N$  is the number of tokens updated in the current step, and  $\mathbb{I}$  is the indicator function.  
 896 This mechanism is controlled by the parameter early\_stop\_threshold = 0.99, which controls the  
 897 sensitivity. The higher the threshold, the more stable the token sequence needs to be before triggering  
 898 early stopping.

902 The parameter settings use a block-based diffusion strategy: the total generation length of 512 tokens  
 903 is divided into blocks of length block\_length = 32. Temperature = 0.7 helps balance exploration  
 904 and exploitation. We choose low\_confidence strategy, which updates tokens with low confidence.  
 905 This combination ensures the quality of the generated text while improving efficiency by using fewer  
 906 diffusion steps, which typically converge to a value smaller than the maximum 512 steps.

## 907 F PROMPTING EXPERIMENT DETAILS

### 908 F.1 EXPERIMENTAL SETUP

912 In this study, we employ the following models: Dream-7B-Instruct (Ye et al., 2025), LLaDA-8B-  
 913 Instruct (Nie et al., 2025b), LLaDA-v1.5 (Zhu et al., 2025), and LLaDOU-Math (Huang et al., 2025).  
 914 To optimize performance, we experiment with a temperature range of [0, 1], choose top-p=0.95  
 915 and block-length=32, and select the maximum token length from the set {128, 256, 512}, as well  
 916 as the diffusion step from {128, 256, 512}. For each model, we use the default decoding settings.  
 917 Additionally, we apply low-confidence remasking to explore the scaling behavior. All experiments  
 918 conduct on a single A100 or A800 80G GPU.

918 F.2 SEQUENTIAL REASONING PROMPTING  
919920 These methods were originally designed primarily for Autoregressive Large Language Models  
921 (ALLMs) and have played a key role in optimizing their reasoning capabilities:  
922923  
924 **Zero-CoT (Kojima et al., 2022)** This strategy uses a simple natural-language instruction (e.g.,  
925 “Let’s think step by step”) to elicit the inherent sequential reasoning capability of autoregressive  
926 models, enabling them to generate coherent reasoning chains without in-context examples.  
927**Zero-CoT**928 Question: {question}  
929930 Let’s think step by step:  
931932  
933 **Plan-and-Solve (Wang et al., 2023)** By explicitly separating the reasoning process into a planning  
934 phase and an execution phase, this strategy prompts autoregressive models to first outline a solution  
935 framework and then complete the details. This improves structural integrity and global consistency,  
936 and is particularly effective for tasks requiring multi-step reasoning and long-range dependency.  
937**Plan-and-Solve**938 Let’s first understand the problem, extract relevant variables and their corresponding numerals,  
939 and make a complete plan. Then, let’s carry out the plan, calculate intermediate variables (pay  
940 attention to correct numerical calculation and commonsense), solve the problem step by step,  
941 and show the answer.  
942943 Question: {question}  
944945  
946 **Least-to-Most (Zhou et al., 2022)** In autoregressive models, this method decomposes complex  
947 problems into a sequence of simpler subproblems, guiding the model to solve them incrementally  
948 rather than tackling the full problem in a single step.  
949**Least-to-Most**950 Q: Elsa has 5 apples. Anna has 2 more apples than Elsa. How many apples do they have  
951 together?  
952953 A: Let’s break down this problem:  
954955 1. How many apples does Anna have?  
956 2. How many apples do Elsa and Anna have together?  
957 1. Anna has 2 more apples than Elsa. So Anna has  $2 + 5 = 7$  apples.  
958 2. Elsa and Anna have  $5 + 7 = 12$  apples together.  
959960 Q: question  
961962 A: Let’s break down this problem:  
963964 Together, these strategies enhance the sequential reasoning behavior of autoregressive models, using  
965 prompt design to induce systematic and logically sound reasoning trajectories. However, we find that  
966 these strategies are not well suited to DLLMs.  
967968 F.3 CONSTRAINT-GUIDED REASONING PROMPTING  
969970 **Complex-CoT:** The original version of Complex-CoT (Fu et al., 2022) leverages a few-shot  
971 reasoning technique to prompt LLMs into performing more sophisticated reasoning processes. This  
approach enhances the model’s ability to handle tasks that require a series of logical inferences or

972 multi-step reasoning, thereby improving the overall performance on complex questions. Specifically,  
 973 by providing a few-shot example that demonstrates how to perform intricate reasoning, the model  
 974 learns to apply similar patterns to new, unseen problems.

975 In contrast, the Constrained-Guided Version of Complex-CoT introduces a crucial modification to  
 976 meet specific requirements. Rather than using few-shot examples, we reframe the prompting method  
 977 as instruction-based, zero-shot constraints created by human experts. These constraints guide the  
 978 model’s reasoning process without the need for training on a set of example problems. To implement  
 979 this approach, the following prompting structure is used to ensure that the model approaches each  
 980 question with the necessary depth and detail:

### 982 Complex-CoT (Constrained-Guided Version)

984 You should think about the following question as thoroughly and in as much detail as possible.

985 Question: {question}

988 **MARP:** The original MARP (Chen et al., 2024) employs an instruction-based, in-context-learning  
 989 approach to guide LLMs in structuring and constraining each step of the reasoning process. This  
 990 method decomposes complex problems into manageable components by promoting multi-step rea-  
 991 soning, while ensuring each step is focused and achievable. By constraining reasoning at each stage,  
 992 MARP prevents overgeneralization and ensures logical, organized outputs.

993 To meet the requirements of the Constrained-Guided Version, we modify MARP in two ways: first, by  
 994 organizing reasoning into discrete steps, and second, by enabling parallel processing within each step.  
 995 This approach allows the model to perform multiple operations simultaneously without compromising  
 996 clarity or precision. The key concept is to balance step-by-step reasoning with parallel processing,  
 997 enhancing task efficiency. Each reasoning step involves multiple basic operations, ensuring clarity  
 998 and minimizing computational overhead.

999 The following prompt structure is used to guide the model’s reasoning process:

### 1001 MARP (Constrained-Guided Version)

1003 Reason step by step, but process operations in parallel.

- 1004 • At each step, you may perform multiple simple operations (up to 5).
- 1005 • Each operation must remain basic and not involve excessive complexity.
- 1006 • If you choose to perform more operations in a single step, then each operation must be  
 1007 correspondingly smaller in scope.

1009 Question: {question}

## 1012 F.4 PARALLEL-ENCOURAGING PROMPTING

1014 This parallel-encouraging strategy is essential for reducing PSC in reasoning. To adapt MARP for  
 1015 DLM, we enable the model to process multiple operations concurrently, avoiding the bottleneck of  
 1016 sequential processing, where each step depends on the completion of the previous one. This parallel  
 1017 processing speeds up reasoning and enhances scalability. At the same time, operational complexity  
 1018 constraints ensure the reasoning process remains clear and manageable. This method balances parallel  
 1019 execution with simplicity, allowing for effective multi-step reasoning without overwhelming the  
 1020 model with overly complex tasks. Moreover, the model adjusts operation complexity dynamically.  
 1021 When tasked with more operations in a given step, each operation must be simpler, preventing  
 1022 cognitive overload and helping the model stay focused on individual tasks.

1023 Ultimately, this approach enables the model to execute parallel reasoning efficiently while maintaining  
 1024 clarity and precision. The detailed prompting for implementation is as follows:

1025 **The key to Diff-MARP’s performance improvement lies in the phrase “Reasoning in parallel,” which  
 explicitly elicits the model’s parallel reasoning capability. Diff-MARP formalizes and systematizes**

1026

1027

1028

**Diff-MARP**

1029

Reasoning in parallel.

1030

- At each step, you may carry out several small operations at the same time.
- Keep each operation simple.

1031

Question: {question}

1032

1033

1034

1035

1036

1037

1038

this instruction, constraining the number of parallel operations to avoid overload and thereby improving robustness relative to using the instruction alone. Here, we present a case using Diff-MARP.

1039

1040

**Diff-MARP Case Study: Using Diff-MARP Strategy**

1041

Step 1: Calculate the total sleep time from 10 p.m. to 6 a.m.

1042

Kim's usual sleep time is from 10 p.m. to 6 a.m., which is 8 hours. Since there are 60 minutes in an hour, her usual sleep time is  $8 * 60 = 480$  minutes.

1043

Step 2: Calculate the time she spent sleepwalking.

1044

Kim was sleepwalking from 2:15 to 2:38 am, which is 23 minutes.

1045

Step 3: Calculate the time she woke up earlier than usual.

1046

Kim woke up 5 minutes earlier than usual to go to the bathroom.

1047

Step 4: Calculate the total time she spent on her bed that day.

1048

The total time she spent on her bed that day is the usual sleep time minus the time she spent sleepwalking and the time she woke up earlier than usual. So,  $480$  minutes -  $23$  minutes -  $5$  minutes =  $452$  minutes.

1049

1050

1051

1052

1053

1054

1055

1056

**Diff-MARP Case Study: Without Using Diff-MARP Strategy**

1057

1. Kim usually goes to sleep at 10 p.m. and wakes up at 6 a.m. This is a total of 8 hours.

1058

2. On the day in question, she woke up 5 minutes earlier than usual to go to the bathroom.

1059

3. She was sleepwalking from 2:15 to 2:38 am, which is a total of 23 minutes.

1060

4. To find out how many minutes she slept on her bed that day, we need to subtract the time she was sleepwalking from the total time she was supposed to sleep.

1061

5. Total sleep time =  $8$  hours \*  $60$  minutes/hour =  $480$  minutes

1062

6. Sleepwalking time =  $23$  minutes

1063

7. Sleep on bed time = Total sleep time - Sleepwalking time =  $480$  minutes -  $23$  minutes =  $457$  minutes

1064

8. However, she woke up 5 minutes earlier than usual, so we need to subtract those 5 minutes from the sleep on bed time.

1065

9. Final sleep on bed time =  $457$  minutes -  $5$  minutes =  $452$  minutes.

1066

1067

1068

1069

1070

1071

1072

1073

- **Using the Diff-MARP Strategy:** We observe that each calculation step has clear labeling and explanations. This decomposition helps ensure transparency in the reasoning process, and the logical steps from start to finish are easy to follow.
- **Without Using the Diff-MARP Strategy:** When Diff-MARP is not used, although the result is the same, the steps are more condensed, and the structure of the reasoning process is weaker. The key difference lies in how Diff-MARP systematically guides the reasoning process, providing clearer explanations for each sub-step.

1080  
1081 F.5 VERIFICATION OF PSC-SPECIFIC ENHANCEMENT1082 To verify that the observed performance differences arise from the PSC rather than a general limitation  
1083 of DLLMs, we conducted additional experiments on the autoregressive model Qwen2.5-7B on the  
1084 GSM8K dataset using the same prompts as in the main text.

Prompting Strategy	Accuracy
Baseline	56.8%
Complex-CoT	41.5%
MARP	58.8%
Diff-MARP	54.2%

1092 Table 2: Performance of Qwen2.5-7B on GSM8K under different prompting strategies.  
10931094 The results are presented in Table 2. If the performance difference stemmed only from a general  
1095 capability gap, the autoregressive model should also show substantial gains under these prompts.  
1096 However, its improvement is minimal or even absent. This supports our hypothesis that these strategies  
1097 enhance DLLM performance mainly by mitigating PSC, rather than acting as general-purpose prompt  
1098 or capability improvements.  
10991100 F.6 DLLM PERFORMANCE UNDER HIGH-COMPUTATIONAL BUDGETS  
11011102 To exclude the alternative explanation that limited iteration steps hinder long-text generation, we  
1103 evaluated DLLMs under high-computation budgets (max\_length=1024, diffusion\_steps=1024) using  
1104 sequential prompts (e.g., Zero-CoT, Least-to-Most, Plan-and-Solve). As summarized in Table ??, per-  
1105 formance remains constrained even with substantially increased computation. This result strengthens  
1106 the causal attribution of DLLM performance limits in sequential reasoning tasks primarily to PSC.  
1107

Prompting Strategy	DREAM	LLaDA
Baseline	80.6%	75.6%
Zero-CoT	75.8%	72.3%
Least-to-Most	69.3%	73.5%
Plan-and-Solve	72.4%	71.6%

1114 Table 3: DLLM performance on sequential prompts under high-computation budgets.  
11151116 F.7 ELIMINATING THE INFLUENCE OF DECODING METHODS  
11171119 We conducted prompt experiments on the BigGSM dataset (Chen et al., 2024) using two decoding  
1120 methods, entropy and topk\_margin. The results shown in table 4 align with those obtained using  
1121 confidence-based decoding, supporting the conclusion that PSC is an inherent property of the DLLM  
1122 architecture rather than a byproduct of the default decoding method.  
1123

Prompting Strategy	Entropy	Topk_margin
Baseline	41.15	41.97
Zero-CoT	41.15	39.84
Least-to-Most	34.75	32.62
Plan-and-Solve	35.25	41.48
Complex-CoT	41.80	43.44
MARP	42.95	46.07
DIFF-MARP	47.21	48.19

1132 Table 4: Performance comparison of different decoding methods.  
1133

1134 F.8 ENVISIONING IDEAL DLLM-FRIENDLY TRAINING DATA  
1135

1136 The essence of DLLM parallel generation is global optimization and the joint generation of all tokens.  
1137 Sequential reasoning requires localized and conditional generation. PSC occurs precisely between  
1138 these two paradigms. Therefore, we argue that an ideal, DLLM-friendly training dataset should not  
1139 follow the traditional linear chain structure but should possess the characteristic of "global conditional  
1140 independence." Our envisioned design for an ideal synthetic dataset is as follows:

1141  
1142 **Pre-planning Topology Structure**

- 1143 • **Data Format:** [Question] → [Global Plan] → [Detailed Steps] →  
1144 [Answer].
- 1145 • **Design Principle:** The model is required to first output a high-level, low-dimensional "solution  
1146 blueprint" or "global plan" in parallel before generating specific content. Once this plan is  
1147 determined, it becomes the common condition for generating all subsequent detailed steps. Under  
1148 this condition, the interdependency between individual steps is significantly reduced, making  
1149 them suitable for parallel filling.

1150  
1151 **Modular Parallel Structure**

- 1152 • **Data Format:** Explicitly label parallel blocks, e.g., [Subtask A] || [Subtask B] ||  
1153 [Subtask C] → [Result Aggregation].
- 1154 • **Design Principle:** Decompose complex problems into independently solvable submodules and  
1155 explicitly declare in the data that these submodules can be processed in parallel. This forcibly  
1156 guides the model to learn the independence between tasks rather than an inherent sequential  
1157 order.

1158 By systematically constructing synthetic data incorporating the above explicit independence assumptions  
1159 and using it for fine-tuning the model, we can fundamentally reshape the model's attention  
1160 mechanism from the root. This enables the model to better adapt to the parallel reasoning paradigm  
1161 of DLLMs, thereby fundamentally mitigating the PSC problem.

1162  
1163 **G DLLM'S LIMITED CAPABILITIES OF LONG CoT REASONING**1164  
1165 **G.1 LONG CHAIN-OF-THOUGHT CAPABILITIES**

1166 Following Chen et al. (2025), the Long Chain-of-Thought (Long CoT) reasoning capabilities comprise  
1167 three linked components: deep reasoning, exploration, and reflection.

1168 **Deep Reasoning.** Given  $s_i$  as the  $i$ -th reasoning step, Deep reasoning models the conditional probability  
1169  $p_\theta(s_0, s_1, \dots, s_K | s_0)$ , facilitating multi-step logical inference through iterative refinement.  
1170 The associated reverse process can be characterized by a factorization:

$$1171 p_\theta(s_0, s_1, \dots, s_K | s_0) = \prod_{i=0}^K p_\theta(s_{i+1} | s_i). \quad (20)$$

1172 **Exploration.** Exploration stems from the probabilistic nature of the reverse process. At each  
1173 exploration step  $s_j$ , multiple samples  $s_j^k$  can be drawn from the conditional distribution  $p_\theta(s_j | s_i, i < j)$ ,  
1174 enabling the model to explore diverse plausible continuations or solutions. This is formalized as:

$$1175 s_j^k \sim p_\theta(s_j | s_i, i < j), \quad k = 1, \dots, K, \quad (21)$$

1176 where  $K$  controls the breadth of exploration. This sampling diversity enhances robustness by covering  
1177 multiple reasoning paths and mitigating premature convergence to suboptimal outputs.

1178 **Reflection.** We view reflection as a self-correction mechanism arising from iterative conditioning  
1179 on latent states. At each reverse step, the model revises its belief about the target sequence using the  
1180 immediately previous state and, via the accumulated latent trajectory, all prior estimates. Formally,  
1181 this corresponds to implicit message passing:

$$1182 \hat{s}_j \sim p_\theta(s_j | s_i, i \geq j), \quad (22)$$

1188 where  $\hat{s}_j$  denotes the corrected state at step  $j$ , enabling iterative error correction and refinement.  
 1189

1190 Together, these components yield a procedure that combines structured logics with stochastic explo-  
 1191 ration and continual self-correction, enabling effective reasoning on complex multi-step tasks.  
 1192

## 1193 G.2 STATEGIES FOR SELF-REFLECTION AND SELF-EXPLORATION EXPERIMENTS ON DLLMs

1194 To investigate whether DLLMs truly possess the fundamental capabilities for Long CoT Reasoning,  
 1195 we designed two sets of experiments: self-reflection and self-exploration, using two distinct prompting  
 1196 strategies to examine the basic abilities of DLLMs.  
 1197

1198 **Self-Reflection:** (1) Prompting Reflection, structured reflection prompts are embedded within initial  
 1199 instructions, requiring the model to perform logical self-checking during generation. (2) Autoregressive  
 1200 Forcing Reflection, correction prompts (e.g., "Wait...there might be something wrong") are replaced  
 1201 with the end-of-sequence (EOS) token as a post-generation intervention strategy.  
 1202

1202 **Self-Exploration:** (1) Prompting Exploration, which embeds exploration prompts in initial instructions  
 1203 to activate multi-path reasoning. (2) Autoregressive Forcing Exploration, which replaces EOS token  
 1204 to "Let's think in another way..." to induce exploratory reasoning.  
 1205

## 1206 G.3 EVALUATION OF SELF-REFLECTION AND SELF-EXPLORATION CAPABILITIES

1207 In evaluating the self-reflection capabilities of the LLaDA-8B-Instruct (Nie et al., 2025a) and Dream-  
 1208 7B-Instruct (Ye et al., 2025) models, the BigGSM dataset was utilized. During the generation process,  
 1209 we employed a temperature of 0.7 for self-reflection and 0.2 for self-exploration, coupled with top-p  
 1210 sampling set to 0.95. Additionally, diffusion steps were configured to 512, and the generation length  
 1211 was fixed at 512. For the investigation into the models' self-exploration capabilities, the experimental  
 1212 settings were identical, with the sole distinction being the substitution of the reflection strategy with  
 1213 an exploration strategy. Based on the setting of Qin et al. (2023), we utilize the following reasoning  
 1214 metrics for deeper analysis:  
 1215

1215 **Semantic Alignment:** The semantic alignment metrics (Golovneva et al., 2022) lies in the *reasoning*  
 1216 *alignment vector*, which spans from the  $N$ -step hypothesis  $\mathbf{h}$  to the source  $\mathbf{s}$  of length  $T$ :  
 1217

$$r\text{-}align(\mathbf{h} \rightarrow \mathbf{s}) = \{\alpha_1, \alpha_2, \dots, \alpha_N\}, \quad (23)$$

1219 where each alignment value can be calculated as:  
 1220

$$\alpha_i = r\text{-}align(h_i \rightarrow \mathbf{s}) = \frac{[1 + \max_{j=1}^T \cos(h_i, s_j)]}{2} \in [0, 1]. \quad (24)$$

1223 Here, such an alignment value is the normalized cosine similarity between the reference step and the  
 1224 most similar sentence in a context, and explicitly measures the *grounding* of the step-wise reasoning  
 1225 with respect to the source text. The alignment vector  $r\text{-}align(h \rightarrow \mathbf{s})$  is estimated by matching the  
 1226 source text and the reasoning chain on the embeddings of the tokens and individual reasoning steps.  
 1227 A similar confidence alignment score is introduced in CTC to measure whether the information of  
 1228 the  $i$ -th source document token  $s_j$  is supported by the hypothesis token  $h_i$ , assessing whether the  
 1229 reasoning step  $h_i$  supports the source context  $s$ .  
 1230

1230 **Repetition-word:** To identify repeated, or paraphrased steps, we look at the repetition word  
 1231 scores (Golovneva et al., 2022) between all steps in the hypothesis chain:  
 1232

$$1 - \max_{i=2 \dots N} \max_{j=1 \dots i-1} \left[ (1/M_i) \sum_{l=1}^{M_i} r_{\text{align}}^{\text{token}}(h_{i,l} \rightarrow h_j) \right].$$

1235 For each pair of sentences, we look at the mean token alignment and find those sentences that  
 1236 maximize this alignment score. In other words, Repetition-Token will punish chains where there are  
 1237 at least two steps with high overlap in token embeddings.  
 1238

1239 **Informativeness:** Measures how well information present in the source is used in the reasoning  
 1240 steps, we calculate informativeness (Golovneva et al., 2022):  
 1241

$$\frac{(1/T) \sum_{t=1}^T r_{\text{align}}(s_t \rightarrow h) + (1/N) \sum_{i=1}^N r_{\text{align}}(h_i \rightarrow s)}{2}.$$

1242 Info-step gives a higher score to reasoning steps that are well-grounded with respect to the source,  
 1243 and identifies the degree of information from source that is covered by the generated hypothesis. A  
 1244 lower Info-Step score corresponds to the reasoning steps that are not related to the source sentences  
 1245 or have missed information provided in the context.

1246 **Reasoning-Alignment** : The most straightforward way to evaluate the correctness of the hypothesis  
 1247 chain is to compare the degree of the overlap between the hypothesis and the reference. One way of  
 1248 doing that is to measure the reasoning alignment (Golovneva et al., 2022) between them:  
 1249

$$1250 \quad \frac{1}{N} \sum_{i=1}^N r_{\text{align}}(h_i \rightarrow r).$$

$$1251$$

$$1252$$

1253 **Token-Entropy** : To calculate the token entropy, we will utilize the pipeline as follows: First,  
 1254 calculate the probability of each token  $p(t_i)$ , which is the frequency of token  $t_i$  divided by the total  
 1255 number of tokens  $N$ :

$$1256 \quad p(t_i) = \frac{\text{count}(t_i)}{N}$$

$$1257$$

1258 Next, calculate the information content  $I(t_i)$  of each token, which reflects the uncertainty contribution  
 1259 of that token to the text:

$$1260 \quad I(t_i) = -\log(p(t_i))$$

$$1261$$

1262 Finally, token-entropy is the weighted average of the information content of all tokens, given by:

$$1263 \quad H = -\sum_{i=1}^N p(t_i) \log(p(t_i))$$

$$1264$$

$$1265$$

1266 where  $p(t_i)$  is the probability of token  $t_i$ , and  $\log(p(t_i))$  is the corresponding logarithmic information  
 1267 content. Token-entropy reflects the overall uncertainty of the text. A higher value indicates that  
 1268 the text is more random and diverse, while a lower value suggests that the text is more focused and  
 1269 repetitive.

1270 **Cosine similarity (Sim)** Cosine similarity measures the degree of similarity between two vectors  
 1271 encoded by BGE (Xiao et al., 2023) by calculating the cosine of the angle between them. For text  
 1272 embedding vectors, a value closer to 1 indicates greater semantic similarity. Let the two generated  
 1273 text vectors be  $\mathbf{A}$  and  $\mathbf{B}$ . The formula for calculating their cosine similarity is:

$$1274 \quad \text{cosine\_similarity}(\mathbf{A}, \mathbf{B}) = \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\| \cdot \|\mathbf{B}\|} = \frac{\sum_{i=1}^n A_i B_i}{\sqrt{\sum_{i=1}^n A_i^2} \cdot \sqrt{\sum_{i=1}^n B_i^2}} \quad (25)$$

$$1275$$

$$1276$$

1277 where  $\mathbf{A} \cdot \mathbf{B}$  is the dot product of vectors  $\mathbf{A}$  and  $\mathbf{B}$ .  $\|\mathbf{A}\|$  and  $\|\mathbf{B}\|$  are the Euclidean norms  
 1278 (magnitudes) of vectors  $\mathbf{A}$  and  $\mathbf{B}$ .  $A_i$  and  $B_i$  represent the components of vectors  $\mathbf{A}$  and  $\mathbf{B}$  along the  
 1279  $i$ -th dimension.

1280 **Perplexity (PPL) of the Model** Perplexity is a concept in information theory used to measure  
 1281 the uncertainty of a probabilistic model in predicting samples. In natural language processing, it is  
 1282 employed to evaluate how well a language model fits a set of test data.

1283 Given a sequence of  $N$  tokens  $W = w_1, w_2, \dots, w_N$ , where the language model predicts the  
 1284 probability of this sequence  $P(W)$ , the perplexity of the sequence is defined as:

$$1285 \quad \text{PPL}(W) = P(W)^{-\frac{1}{N}} = \exp\left(-\frac{1}{N} \log P(W)\right). \quad (26)$$

$$1286$$

$$1287$$

1288 Because of the sequence's independence assumption, we can compute  $P(W)$  as:

$$1289 \quad P(W) = \prod_{i=1}^N P(w_i | w_1, \dots, w_{i-1}). \quad (27)$$

$$1290$$

$$1291$$

$$1292$$

1293 Therefore, the commonly seen formula for perplexity is:

$$1294 \quad \text{PPL}(W) = \exp\left(-\frac{1}{N} \sum_{i=1}^N \log P(w_i | w_1, \dots, w_{i-1})\right), \quad (28)$$

$$1295$$

1296 where  $\log P(W)$  is the log probability of the entire sequence.  $\frac{1}{N} \sum_{i=1}^N \log P(w_i | w_1, \dots, w_{i-1})$  is  
 1297 the average log probability of the sequence under the model.  $\exp$  is the exponential function, used to  
 1298 transform the log probability back to its original scale.

1299 Observation of the results for both DREAM (Ye et al., 2025) and LLaDA (Nie et al., 2025a) models,  
 1300 under both self-reflection and self-exploration settings, the scores across all ROSCOE-SA evaluation  
 1301 metrics are highly similar. This indicates that current diffusion language models (DLLMs) have not  
 1302 yet genuinely acquired the deeper capabilities of self-reflection and self-exploration, as their outputs  
 1303 do not exhibit significant differences under varying strategic prompts.

#### 1305 G.4 EVALUATION OF DEEP-REASONING CAPABILITY

1307 Following Chen et al. (2024), we further investigate reasoning boundaries (RBs) in deep reasoning  
 1308 capabilities in mathematical reasoning. We prompt DLLMs to generate plans and assess their  
 1309 accuracy through manual evaluation. When the model meets the question with fewer than 1 reasoning  
 1310 steps, accuracy surpasses 80%. Conversely, when reasoning steps exceed 3-4, accuracy falls below  
 1311 10%. Moreover, we first randomly select 200 samples to generate examples and split steps from the  
 1312 DLLM-generated rationales based on ROSCOE (Golovneva et al., 2022). Further, we also manually  
 1313 identify the first model’s incorrect step position.

#### 1316 G.5 REFLECTION/EXPLORATION GAINS IN ALLMs AND DLLMs

1318 The core inference of this section is that the limited gains from reflection and exploration in diffusion-  
 1319 based decision-making large language models (DLLMs) arise mainly from their inherent PSC  
 1320 characteristics. To test this, we conducted a comparative analysis with an autoregressive baseline  
 1321 under matched settings. The results in Table 5 show that the autoregressive model achieves substan-  
 1322 tially larger performance improvements, supporting the view that the weak gains of reflection and  
 1323 exploration are driven by the PSC issue in diffusion models rather than by intrinsic shortcomings of  
 1324 these methods.

	Qwen2.5-7B	LLaDA-8B-Instruct	Dream-7B-Instruct
Baseline	56.8%	52.72%	50.7%
Prompting Reflection	63.3%	52.75%	47.0%
Prompting Exploration	63.5%	52.75%	48.75%
Autoregressive Forcing Reflection	73.9%	47.5%	37.25%
Autoregressive Forcing Exploration	72.1%	52.0%	44.5%

1332 Table 5: Reflection/Exploration Benefits: Autoregressive vs. Diffusion Models

## 1335 H THREE-DIRECTIONAL INFERENCE-TIME SCALING ON DLLMs

### 1338 H.1 PARALLEL SCALING EXPERIMENT DETAILS

1339 In the parallel scaling section, we utilize the dual-cache generation strategy from Fast-dLLM based  
 1340 on the diffusion language model LLaDA-8B-Instruct (Nie et al., 2025a), and perform batch process-  
 1341 ing on the BigGSM (Chen et al., 2024) dataset. Key configurations include: diffusion\_steps=256,  
 1342 gen\_length=256, block\_length=32, and a Dynamic Low-Confidence Remasking mechanism.

1343 We also employed the dual\_cache generation strategy from Fast-dLLM (Wu et al., 2025) on the  
 1344 Dream-7B-Instruct (Ye et al., 2025) model for testing on the BigGSM reasoning dataset. The core  
 1345 configuration includes: diffusion\_steps=256, gen\_length=256, and block\_length=32.

1347 The results are shown in Figure 11. It can be observed that the accuracy generally increases with  
 1348 higher k-values. At the initial attempts, the accuracy at Temperature 1.0 was relatively low. Although  
 1349 it showed significant improvement in the early stages, its later accuracy fell behind other temperatures.  
 At Temperature 0.1, the accuracy growth was more stable initially, but eventually plateaued at around

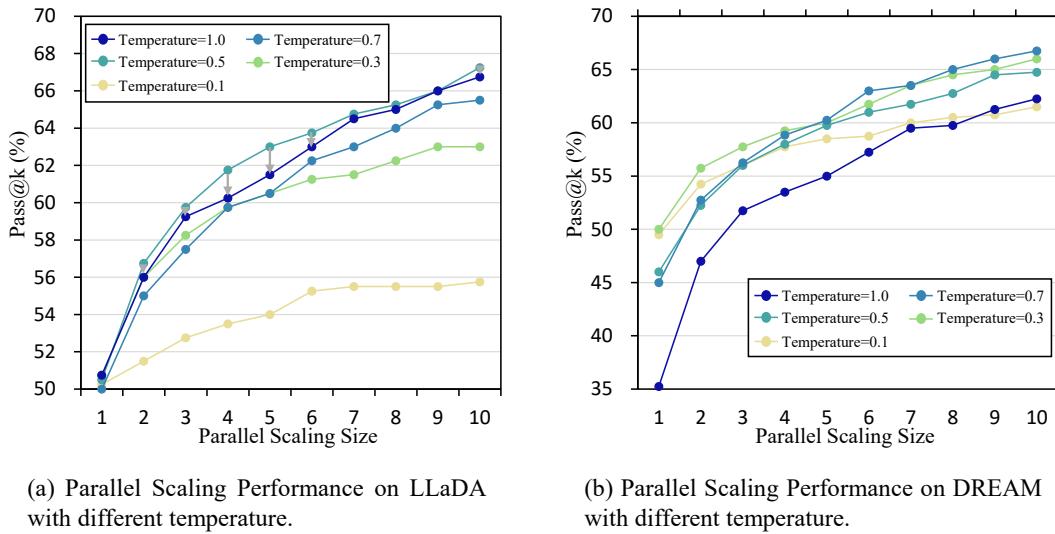


Figure 11: Parallel scaling performance of DLLMs under Different Temperature Settings

60%, similar to Temperature 1.0. Overall, intermediate temperatures demonstrated better pass@k accuracy performance, achieving higher accuracy with more consistent and stable growth.

## H.2 DIFFUSION SCALING EXPERIMENT DETAILS

We set the diffusion\_step to be between 1 and 4096 (with max-token-length equal to 512 for Figure 9 (a). The other settings are identical to those of parallel scaling. The results are shown in Figure 9 (a).

## H.3 SEQUENTIAL SCALING EXPERIMENT DETAILS

We set the Max Token Length to be between 1 and 4096 (with diffusion-step equal to max-token-length for Figure 9 (d). The other settings are identical to those of parallel scaling. The results are shown in Figure 9 (d).

## I REINFORCEMENT LEARNING INFLUENCE ON PSC

To assess the impact of post-training methods on PSC, we compare the LLaDA-8B-Instruct (Nie et al., 2025b) model with its d1 variant fine-tuned by reinforcement learning. Experiments are conducted on the BigGSM dataset (Chen et al., 2024), from which 100 samples are randomly selected. Under identical inference hyperparameters, we examine the decoding order during generation. We quantify PSC using the following two metrics:

**Inversion Pair Ratio:** During generation, if a token at an earlier position is produced after a token at a later position (i.e., if  $i < j$ , but the generation step  $\text{step}[i] > \text{step}[j]$ ), it is counted as an inversion pair. A higher ratio indicates more severe sequential conflict in decoding.

**Spearman’s Rank Correlation Coefficient:** This metric measures the consistency between the model’s generation order and the natural left-to-right language order. A higher coefficient indicates that the generation process more closely follows habitual language sequencing.

Model	Inversion Pair Ratio	Spearman Coefficient
LLaDA-8B-Instruct	3.41%	98.62%
d1 (After RL Training)	<b>6.65% (+3.24)</b>	<b>97.27% (-1.35)</b>

Table 6: Comparison of PSC Phenomenon Before and After RL Training

1404 The results show that, after reinforcement learning training, the model’s inversion pair ratio increases  
 1405 from 3.41% to 6.65%, while the Spearman coefficient decreases from 98.62% to 97.27%. This  
 1406 suggests that reinforcement learning intensifies sequential conflicts during decoding and makes the  
 1407 PSC phenomenon more pronounced.

1408 Here, we explain why the inversion pair ratios remain relatively low and the Spearman coefficients  
 1409 remain close to 1. Diffusion language models are designed to generate coherent text sequences that  
 1410 follow linguistic conventions, rather than arbitrary permutations of words. To this end, the model  
 1411 architecture (for example, the Block mechanism) guides generation at a coarse level to follow the  
 1412 sequential structure of natural language. In particular, the generation order between different blocks  
 1413 is strictly constrained to proceed from left to right. Consequently, both the inversion ratio and the  
 1414 Spearman coefficient still reflect strong overall adherence to sequential generation.

## 1416 J THE RELATIONSHIP BETWEEN CONTEXT DEPENDENCY AND PSC

1418 Establishing how context dependency strength relates to DLLM performance is key to characterizing  
 1419 the PSC phenomenon. We quantify this relationship via an information-theoretic analysis, using  
 1420 conditional entropy to measure context dependency. Specifically, for each sample in the BigGSM  
 1421 dataset, we compute the conditional entropy at each generation step with a GPT-2 model and then  
 1422 average these values across steps to obtain the sample’s average conditional entropy.

1423 The conditional entropy for token  $x_i$  given context  $x_{<i}$  is:

$$1425 H(x_i|x_{<i}) = - \sum_{w \in V} P(w|x_{<i}) \cdot \log P(w|x_{<i}). \quad (29)$$

1427 The average conditional entropy over a text of length  $T$  tokens is:

$$1429 \text{Avg}_{\text{Entropy}} = \frac{1}{T-1} \sum_{i=1}^{T-1} H(x_i|x_{<i}), \quad (30)$$

1432 which quantifies how strongly the generated answer is constrained by the context: lower entropy  
 1433 indicates stronger contextual constraints and thus higher dependency.

1434 Using this metric, we partition the BigGSM dataset into three subsets and compute, for each subset,  
 1435 the average conditional entropy and the Dream accuracy, which reflects DLLM performance.

1438 Entropy Level	1438 Sample Size	1438 Avg. Entropy	1438 Dream Acc	1438 Avg. DI
1439 Low Entropy	203	2.5269	13.79%	27.41
1440 Medium Entropy	203	3.0766	35.29%	21.56
1441 High Entropy	203	3.4951	45.32%	19.06

1442 Table 7: Relationship between Context Dependency Strength and DLLM Performance

1445 The results are summarized in Table 7. As contextual dependency strengthens (i.e., conditional  
 1446 entropy decreases), DLLM performance drops sharply from 45.32% to 13.79%. In addition to  
 1447 conditional entropy, we introduce the Attention Distance Index (ADI) to further examine how  
 1448 dependency strength affects DLLM performance. The ADI measures whether the model’s attention  
 1449 at each layer preferentially targets long-distance tokens, thus characterizing its ability to model strong  
 1450 dependencies.

1451 We define the **Attention Distance Index (ADI)** to quantify the model’s tendency to attend to  
 1452 long-distance tokens. For each layer, the ADI is calculated as:

$$1453 \text{ADI}_{\text{layer}} = \frac{1}{H} \sum_{h=1}^H \left[ \sum_{i=1}^L \sum_{j=1}^L A_{h,i,j} \cdot \frac{|i-j|}{L-1} \right], \quad (31)$$

1455 where  $A_{h,i,j}$  is the attention weight from query token  $i$  to key token  $j$  in head  $h$ ;  $L$  is the sequence  
 1456 length;  $H$  is the number of attention heads; and  $\frac{|i-j|}{L-1}$  is the normalized distance between tokens.

1458     **Dependency Index (DI)** for the entire model is then computed as the average ADI across all layers:  
 1459

$$1460 \quad \text{DI} = \frac{1}{N} \sum_{l=1}^N \text{ADI}_l, \quad (32)$$

1461     where  $N$  is the total number of layers. A higher DI indicates stronger long-range dependency  
 1462     modeling requirements.

1463     Our analysis shows that stronger dependencies (lower conditional entropy or higher DI) markedly  
 1464     degrade DLLM performance, particularly when long-range dependencies are required. This behavior  
 1465     is consistent with intrinsic limitations of the parallel generation architecture and further supports our  
 1466     theoretical interpretation of the PSC phenomenon.

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511