

Empirical Analysis of Decoding Biases in Masked Diffusion Models

Pengcheng Huang¹, Tianming Liu², Zhenghao Liu^{1*},
Yukun Yan³, Shuo Wang³, Tong Xiao¹, Zulong Chen⁴, Maosong Sun³

¹School of Computer Science and Engineering, Northeastern University, Shenyang, China

²Software College, Northeastern University, Shenyang, China

³Department of Computer Science and Technology, Tsinghua University, Beijing, China

⁴Alibaba Group, Hangzhou, China

Abstract

Masked Diffusion Models (MDMs) have recently emerged as a promising non-autoregressive paradigm for sequence generation. However, their performance is highly sensitive to the choice of decoding strategy. In this work, we reveal that prevalent uncertainty-based decoding strategies induce two decoding biases in MDMs: rigid boundary bias and trivial token bias. These biases limit the model’s reasoning ability and ultimately degrade generation quality. To address these challenges, we propose UNmasking Calibration for Decoding DEbiasing (UNCODE), a decoding calibration framework that regularizes uncertainty-based decoding by incorporating two complementary priors to shape global decoding trajectories and promote content informativeness. Extensive experiments on three advanced MDMs across seven reasoning- and planning-intensive benchmarks demonstrate that UNCODE consistently outperforms existing decoding strategies by more than 7%, while achieving performance comparable to autoregressive models of similar parameter scales. The code is publicly available at <https://github.com/NEUIR/Uncode>.

1 Introduction

Diffusion Large Language Models (dLLMs) have recently emerged as a competitive alternative to conventional autoregressive models (ARMs) (Li et al., 2025b; Zhu et al., 2025), offering non-autoregressive decoding that updates multiple tokens per step and excels at global reasoning tasks (Li et al., 2023). Among dLLMs, Masked Diffusion Models (MDMs) are particularly appealing due to their simple masked-denoising formulation and strong empirical performance (Khanna et al., 2025; Nie et al., 2025b; Ye et al., 2025b). Unlike autoregressive models that generate text

* indicates corresponding author.

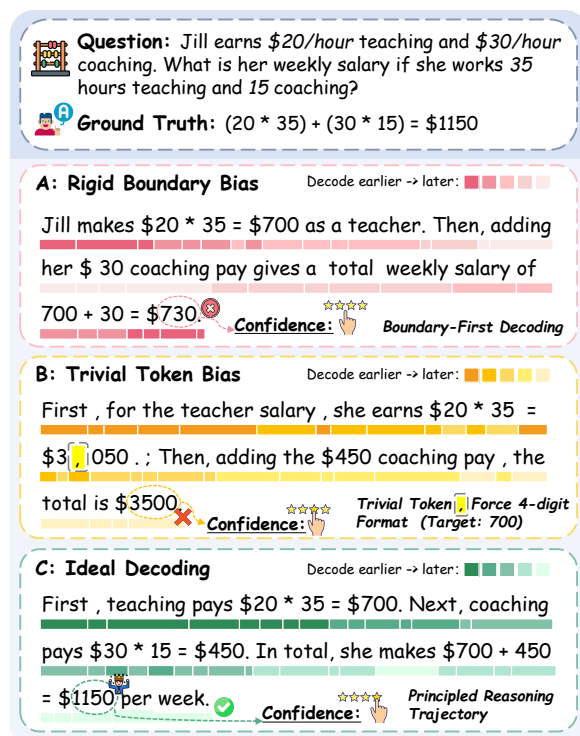


Figure 1: Illustration of decoding biases during reasoning. Uncertainty-based decoding prioritizes (A) rigid boundaries and (B) trivial tokens, causing systematic deviations from optimal problem-solving paths. In contrast, panel (C) presents the ideal unmasking pattern for producing coherent logical chains to achieve reliable inference.

strictly left-to-right, MDMs operate through iterative unmasking (Wu et al., 2025a; Lou et al., 2024). This design removes sequential constraints, permitting tokens to be generated in a flexible order (Kim et al., 2025). In practice, however, most existing MDMs employ uncertainty-based decoding strategies (Chang et al., 2022; Nie et al., 2025b), which greedily unmask tokens at positions with the lowest predictive uncertainty during each step.

Recent studies indicate that modern language models exhibit intrinsic “reasoning trajectories,” where the generation follows a logical reason-then-answer sequence and model confidence in the cor-

rect answer progressively increases as intermediate steps are generated (Wang and Zhou, 2024). In contrast, as illustrated in Figure 1, we identify two structural biases in uncertainty-based decoding that hinder MDMs from following such trajectories and thereby prevent the formation of effective reasoning chains that are essential for complex problem-solving. Specifically, we observe: (1) *Rigid Boundary Bias*, where the decoding trajectory follows a rigid pattern: boundary tokens are consistently decoded first, with updates progressively moving inward toward the center. This behavior prevents the model from adapting its decoding order to task-specific structural requirements and can cause answers to be resolved prematurely without establishing sufficient supporting rationale. (2) *Trivial Token Bias*, where high-frequency but semantically trivial tokens (e.g., punctuation or structural markers) are disproportionately prioritized, hindering the formation of critical reasoning steps and ultimately leading to degraded performance.

To mitigate these biases, we propose UNmasking Calibration for DecODing DEbiasing (UNCODE), a decoding calibration framework that calibrates uncertainty-based decoding for MDMs. Specifically, UNCODE adjusts the raw uncertainty scores using two complementary priors. The first prior, the Positional Trajectory Prior, introduces a position-dependent reweighting to discourage a rigid boundary-first decoding pattern and allow the model to flexibly tailor its trajectory to diverse task structures. The second prior, the Semantic Informativeness Prior, modulates uncertainty scores based on corpus-level frequency statistics, penalizing high-frequency, semantically uninformative tokens to redirect the focus of refinement toward tokens that provide more informative content.

To validate the effectiveness of UNCODE, we conduct extensive experiments on three representative MDMs across seven diverse benchmarks covering a broad range of reasoning and planning tasks. The results demonstrate that UNCODE consistently improves overall generation performance, achieving an average gain of over 7% compared to state-of-the-art decoding baselines, while achieving performance comparable to autoregressive models at a similar scale. By mitigating decoding biases, UNCODE overcomes the limitations imposed by rigid boundary-first decoding in existing MDMs, enabling more effective construction of reasoning trajectories and reducing uncertainty in answer prediction, thereby leading to pronounced improvements.

Furthermore, our analysis shows that the advantages of UNCODE can be kept when integrated with different efficient sampling frameworks, achieving substantial speedups while simultaneously improving generation quality.

2 Related Work

Masked Diffusion Models (MDMs) have emerged as a promising non-autoregressive alternative for text generation (Song et al., 2025; Nie et al., 2025b). Unlike autoregressive models that generate sequences in a left-to-right order, MDMs utilize an iterative unmasking process, allowing tokens to be revealed in a non-sequential and arbitrary manner. This flexibility makes the decoding strategy pivotal to MDM performance (Garg et al., 2025; Kim et al., 2025; Lee et al., 2025), as it directly controls the generation order and the effective context available at each step. As a result, numerous studies have explored how to leverage this flexibility to improve MDM performance (Li et al., 2025b).

Early implementations of MDMs adopted a simple uniform sampling approach (Austin et al., 2021a), where tokens are randomly selected for unmasking at each step. While this strategy aligns with the order-agnostic training distribution (Garg et al., 2025), its unguided nature often leads to suboptimal generation quality (Peng et al., 2025). To address this limitation, research has shifted towards uncertainty-driven heuristics that optimize the unmasking process by utilizing the model’s intrinsic prediction signals (Kim et al., 2025; Ye et al., 2025b; Ben-Hamu et al., 2025). Central to these strategies are three key metrics: confidence (Chang et al., 2022), which prioritizes tokens with the maximum predicted probability; margin (Kim et al., 2025), which measures the probability gap between the top two candidates; and entropy (Ye et al., 2025b), which quantifies overall predictive uncertainty. These approaches have been shown to improve stability and performance across a range of tasks (Ben-Hamu et al., 2025).

Building upon these foundations, a variety of heuristics (Wu et al., 2025a; Li et al., 2025a; Yu et al., 2025; Wei et al., 2025; Hong et al., 2025; Horvitz et al., 2025; Israel et al., 2025) have further refined these frameworks to enable adaptive unmasking, achieving notable efficiency gains with minimal performance degradation. In parallel, semi-autoregressive approaches (Arriola et al., 2025; Nie et al., 2025b) incorporate local

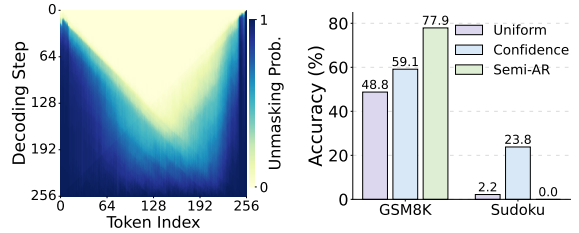
causal dependencies into the unmasking process and often improve performance on logical reasoning tasks. Despite these advances, both lines of work still rely on uncertainty-based scoring to prioritize which positions to unmask. While prior research mainly optimizes *how* uncertainty is used during decoding, we study systematic biases induced by this uncertainty-driven prioritization in MDMs, and identify two decoding biases inherent to uncertainty-based samplers.

3 Motivating Analysis: Decoding Biases in Masked Diffusion Models

In this section, we examine two decoding biases that arise when uncertainty-based samplers are applied to MDMs: *Rigid Boundary Bias* and *Trivial Token Bias*. Our empirical analyses reveal that both biases systematically shape the decoding behavior of MDMs, suppressing the emergence of critical content and ultimately undermining the performance of MDMs across different tasks.

Rigid Boundary Bias. Recent studies indicate that achieving strong performance with MDMs requires the model to use differentiated decoding orders tailored to each task (Kim et al., 2025). Nevertheless, we observe that uncertainty-based samplers consistently induce a boundary-first decoding pattern, which we refer to as *Rigid Boundary Bias*.

To illustrate this phenomenon, we conduct a series of controlled experiments, as shown in Figure 2. We first visualize LLaDA-8B-Instruct’s unmasking dynamics on the GSM8K reasoning benchmark using the confidence-based sampler in Figure 2(a). Specifically, for each token position and decoding step, we track whether the token is unmasked (1 if unmasked, 0 otherwise) and report the average unmasking probability across all instances. The results reveal a pronounced pattern: tokens near the sequence boundaries are systematically prioritized early, yielding a characteristic “U-shaped” decoding trajectory. Notably, we observe this pattern consistently across diverse datasets and uncertainty-based decoding strategies (Appendix A.3). As further analyzed in Appendix A.2, this bias stems from the premature unmasking of boundary tokens (BOS and EOS), where the attention mechanism’s local positional bias leads to elevated confidence for tokens near the sequence boundaries. This effect subsequently causes the unmasking process to move from both ends toward the center, resulting in a rigid boundary-first decoding pattern.



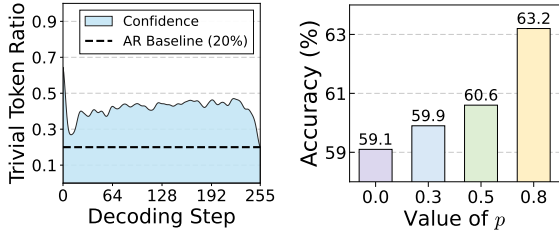
(a) Heatmap of unmasking prob. (b) Accuracy on reasoning vs. abilities.

Figure 2: Visualization of the Rigid Boundary Bias and its impact on downstream performance. (a) Unmasking probability for each token position across decoding steps on GSM8K, with both sequence length and decoding steps set to 256, where darker blue intensities denote higher unmasking probabilities. (b) Accuracy comparison of different decoding strategies on GSM8K (reasoning) and Sudoku (planning).

To better understand the impact of rigid boundary bias, we evaluate several decoding strategies on two representative datasets: GSM8K, which requires strict step-by-step sequential reasoning, and 4x4 Sudoku, a logic puzzle that benefits from globally coordinated planning. As shown in Figure 2(b), we compare three distinct strategies: the uniform sampler, the greedy confidence-based sampler, and the semi-autoregressive (Semi-AR) sampler. The results reveal markedly different performance across both tasks. On GSM8K, the uncertainty-based sampler struggles, achieving only 59.1% accuracy, substantially lower than the 77.9% attained by Semi-AR decoding. Conversely, on Sudoku, the uncertainty-based sampler obtains the highest accuracy, outperforming both Semi-AR and uniform decoding. This task-dependent divergence suggests that while Semi-AR decoding facilitates sequential reasoning by introducing coarse-grained step-wise structure, uncertainty-based decoding may be advantageous for globally coordinated planning. These findings highlight the rigidity inherent in uncertainty-based decoding as a key challenge, emphasizing the need for decoding flexibility to fully unleash the reasoning capability of MDMs.

Trivial Token Bias. We further observe that uncertainty-based samplers disproportionately prioritize semantically trivial, high-frequency tokens (e.g., \n, <space>, the, ., !) throughout the decoding process. We refer to this phenomenon as *Trivial Token Bias* (Figure 3).

We first quantify this bias by tracking the proportion of trivial tokens (definition provided in Appendix A.4) unmasked by the standard uncertainty-



(a) Over selection of trivial tokens. (b) Impact of trivial token suppression.

Figure 3: Verification of trivial token bias and trivial-token suppression efficacy. (a) The proportion of trivial tokens unmasked at each decoding step under the uncertainty-based sampler consistently exceeds the AR baseline (Qwen-2.5-7B-Instruct, dashed line). (b) GSM8K accuracy improves monotonically with suppression probability p .

based (Confidence) sampler as decoding proceeds. As illustrated in Figure 3(a), trivial tokens consistently constitute nearly 40% of the selections in MDMs. This significantly exceeds the $\sim 20\%$ baseline observed in ARMs. This striking discrepancy raises a key question: *Is the over-selection of trivial tokens in MDMs necessary for effective problem-solving, or is it instead detrimental?*

To investigate the impact of this bias, we conduct an intervention experiment on GSM8K by probabilistically suppressing the unmasking of trivial tokens. Specifically, at each decoding step, if the sampler selects a trivial token to unmask, we reject this choice with probability p and instead re-sample from the subset of candidate positions whose predictions are non-trivial. As shown in Figure 3(b), the performance of MDMs improves monotonically as trivial-token selections are suppressed more frequently, indicating that over-selecting trivial tokens is not always necessary and can sometimes hinder reasoning. We hypothesize that trivial tokens are prioritized because greedy uncertainty-based decoding selects, at each step, the positions with the lowest uncertainty globally. Since trivial tokens (e.g., whitespace and punctuation) are typically easier to predict than content-bearing tokens, they are assigned lower uncertainty scores and are consequently unmasked earlier. Although such tokens carry little semantic information, revealing them early can impose surface-level structural commitments (e.g., punctuation and line breaks), which may reshape the uncertainty landscape and subsequent position-selection dynamics. This reduces flexibility in forming informative intermediate states and can lead to suboptimal reasoning

trajectories. We provide a case study in Figure 16 (Appendix A.16) to illustrate this phenomenon.

4 Mitigating Decoding Bias with Unmasking Calibration

In this section, we briefly review the essential background of Masked Diffusion Models (MDMs) (§4.1), and then introduce UNCODE, which addresses both the rigid boundary bias and the trivial-token bias in MDMs by calibrating the unmasking schedule during decoding (§4.2).

4.1 Background of Masked Diffusion Models

This section establishes the formal training objective of MDMs and describes the decoding paradigm that our method aims to optimize.

Unlike autoregressive models, which are trained via next-token prediction, MDMs learn to reconstruct original tokens from partially masked sequences. During training, a sequence x_0 is corrupted by independently masking tokens with a masking ratio $\tau \sim \mathcal{U}(0, 1]$, yielding a partially masked sequence x_τ . The model p_θ is trained to reconstruct the masked tokens by minimizing cross-entropy loss on the masked positions:

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

where x^i denotes the token at the i -th position, \mathbf{M} denotes the mask token, and $\mathbf{1}[\cdot]$ is the indicator function. By sampling τ over a wide range, the model learns to handle various partial-context conditions, enabling it to predict any masked token from an arbitrary set of revealed tokens. This order-agnostic training naturally supports iterative, any-order decoding strategies at inference time.

During inference, the decoding process is formulated as iterative unmasking that gradually recovers a sequence from a fully masked state x_T to a complete sequence x_0 using a predefined unmasking sampler. Specifically, let \mathcal{M}_t denote the set of positions that remain masked at step t . The sampler first assigns an unmasking priority s_t^i to all masked tokens. A subset $\mathcal{I}_t \subseteq \mathcal{M}_t$ is then selected for unmasking by identifying the K positions within \mathcal{M}_t that possess the highest priority scores:

$$\mathcal{I}_t = \left\{ i \in \mathcal{M}_t \mid \text{rank}(s_t^i) \leq K \right\}. \quad (2)$$

For each selected position $i \in \mathcal{I}_t$, the token x_{t-1}^i is sampled according to the corresponding condi-

tional probability distribution:

$$x_{t-1}^i \sim p_\theta(\cdot | x_t, i), \quad (3)$$

and used to update the sequence, yielding a less-masked version, while the remaining positions stay masked. This procedure iterates until every position has been unmasked.

4.2 Calibrating Decoding Biases in MDMs for Effective Inference

Given the current set of remaining masked positions \mathcal{M}_t , existing MDMs typically use an uncertainty-based sampler by assigning a priority score to each masked position and unmasking the positions with the highest priority scores. Formally, at each step t , the sampler assigns a priority score s_t^i to every masked position $i \in \mathcal{M}_t$ using a scoring function $\mathcal{F}(\cdot)$ (e.g., confidence or negative entropy) derived from the model’s predictive distribution (Chang et al., 2022; Ye et al., 2025b):

$$s_t^i = \mathcal{F}(p_\theta(\cdot | x_t, i)), \quad \forall i \in \mathcal{M}_t. \quad (4)$$

However, this uncertainty-driven heuristic induces two systematic biases in MDM decoding: rigid boundary bias and trivial token bias (Section 3). To mitigate these biases, we propose UNCODE, a lightweight calibration framework that adjusts the unmasking priorities by incorporating both a positional trajectory prior and a semantic informativeness prior. The calibrated unmasking priority is defined as:

$$\tilde{s}_t^i = \mathcal{P}^i \cdot \mathcal{S}_t^i \cdot \underbrace{\mathcal{F}(p_\theta(\cdot | x_t, i))}_{\text{Raw Uncertainty Score}}, \quad (5)$$

where \mathcal{P}^i and \mathcal{S}_t^i are positive scalars that rescale the raw uncertainty score. Specifically, \mathcal{P}^i is a *positional trajectory prior* designed to mitigate rigid boundary-first decoding, while \mathcal{S}_t^i is a *semantic informativeness prior* aiming to reduce the dominance of trivial tokens during inference.

Positional Trajectory Prior (\mathcal{P}^i). To mitigate the rigid boundary-first bias observed in uncertainty-based decoding, we introduce a positional trajectory prior \mathcal{P}^i to modulate the spatial distribution of unmasking priorities, enabling consistent improvements across diverse tasks. This prior adjusts the unmasking priority of each position based on its absolute position, effectively reshaping the unmasking order. We instantiate \mathcal{P}^i as a simple position-dependent decay function:

$$\mathcal{P}^i = e^{-\lambda \cdot i}, \quad (6)$$

where i denotes the token index and λ controls the strength of positional regularization. When $\lambda \rightarrow 0$, the prior becomes uniform, recovering uncertainty-driven unmasking and permitting arbitrary generation trajectories. Conversely, larger values of λ impose a stronger sequential bias, encouraging more left-to-right decoding. In this way, UNCODE provides explicit control over the trade-off between non-autoregressive flexibility and causal dependency across diverse tasks.

Semantic Informativeness Prior (\mathcal{S}_t^i). The uncertainty-based decoding strategy usually prioritises high-frequency and semantically trivial tokens, thereby squandering valuable reasoning steps. To alleviate this effect, we introduce a semantic informativeness prior that reweights unmasking priorities based on token-level self-information (Cover, 1999). Specifically, for each masked position i at decoding step t , we consider the model’s greedy predicted $\hat{x}_t^i = \arg \max_{v \in \mathcal{V}} p_\theta(v | x_t, i)$, where v represents a candidate token from the vocabulary \mathcal{V} . We then assign an informativeness weight based on the token’s prior frequency, estimated from a large-scale general corpus \mathcal{D}' :

$$\mathcal{S}_t^i = -\log p_{\mathcal{D}'}(\hat{x}_t^i). \quad (7)$$

This prior downweights frequent, low-information tokens and correspondingly increases the relative priority of semantically informative content, encouraging such tokens to be resolved earlier during decoding. To prevent excessively large weights for rare tokens, we follow Zhang et al. (2024) and clip the informativeness scores at a threshold α , i.e., $\mathcal{S}_t^i \leftarrow \min(\mathcal{S}_t^i, \alpha)$.

5 Experimental Methodology

This section describes datasets, baselines, and implementation details used in our experiments.

Datasets. We evaluate UNCODE on seven benchmarks across four domains, selected to represent diverse structural dependencies: (1) Mathematical reasoning. To assess step-by-step logical deduction, we use GSM8K (Cobbe et al., 2021) and MATH500 (Lightman et al., 2024). (2) Code generation. Focusing on strict syntactic and causal dependencies, we evaluate on HumanEval (Chen et al., 2021) and MBPP (Austin et al., 2021b). (3) Scientific reasoning. To test knowledge-intensive reasoning that combines parametric knowledge with logical deduction, we use GPQA (Rein et al., 2023). (4) Planning. For

tasks prioritizing global constraint satisfaction, we include 4×4 Sudoku (Ye et al., 2025a) and Countdown (Nolte et al., 2024). We report pass@1 for code generation and accuracy for all other tasks. Detailed dataset descriptions and statistics are provided in Appendix A.5.

Baselines. We compare UNCODE against strong baselines from two families: (i) autoregressive models of comparable parameter scale, and (ii) representative decoding strategies for MDMs. To ensure a fair comparison, all MDM baselines are evaluated on the same backbone model under a matched decoding budget. Additional details about the baselines are provided in Appendix A.7.

Autoregressive baselines. We include three instruction-tuned ARMs of similar size: LLaMA-3.1-8B-Instruct (Dubey et al., 2024), Mistral-7B-Instruct (Jiang et al., 2023), and Qwen-2.5-7B-Instruct (Yang et al., 2024).

MDM decoding baselines. We further compare with several representative MDM decoding strategies: (1) Uniform (Austin et al., 2021a), which selects masked positions uniformly at random. (2) Uncertainty-based methods, which greedily rank masked positions by predicted uncertainty. We evaluate three representative proxies: maximum confidence (Chang et al., 2022), negative entropy (Ye et al., 2025b), and margin (Kim et al., 2025). (3) Semi-autoregressive (Semi-AR) decoding (Nie et al., 2025b), which generates the sequence block by block. It enforces a global autoregressive order across blocks, while determining the unmasking positions within each block according to the MDM’s prediction confidence. (4) Adaptive multi-token samplers, which unmask multiple tokens per step using uncertainty thresholds to improve decoding efficiency without significantly compromising generation quality. We evaluate two representative methods: Entropy-Bounded (EB) Sampler (Ben-Hamu et al., 2025) and FastdLLM (Wu et al., 2025a).

Implementation Details. We conduct experiments on three state-of-the-art masked diffusion language models: LLaDA-8B-Instruct (Nie et al., 2025b), LLaDA-1.5-8B (Zhu et al., 2025), and Dream-7B (Ye et al., 2025b). During inference, we adopt the same generation length settings as in Nie et al. (2025b) and set the number of denoising steps to be equal to the sequence length. We also verify the generalization to longer sequences in Appendix A.11. We use prediction confidence as the base uncertainty proxy for cal-

ibration. For the semantic-informativeness prior, we estimate token-frequency statistics by fitting $p_{\mathcal{D}'}(\cdot)$ on a 16GB subset of the C4 corpus (Rafael et al., 2020), following the protocol of Zhang et al. (2024) (see Appendix A.14 for an analysis of the calibration-corpus choice). We fix the clipping threshold to $\alpha = 10$ across all experiments. Regarding the positional-regularization coefficient λ , we set $\lambda = 0$ for Sudoku to favor globally coordinated planning, $\lambda = 0.25$ for the remaining tasks to encourage more progressive decoding, and $\lambda = 0.5$ for Countdown. More implementation details are provided in Appendix A.6.

6 Experiment Results

In this section, we first present the overall performance of UNCODE, followed by ablation results. Subsequently, we analyze the decoding dynamics of UNCODE and assess its compatibility with efficient decoding strategies.

6.1 Overall Performance

We present the main results of UNCODE across all benchmarks in Table 1.

Overall, UNCODE consistently outperforms all baseline decoding strategies, achieving over a 7% average performance improvement. Notably, this advantage generalizes well across different backbone models and task categories, with UNCODE achieving near-best performance among all baselines across the seven benchmarks. In contrast, existing decoding methods exhibit pronounced task specialization. For example, uncertainty-based samplers perform well on planning tasks such as Sudoku, outperforming the vanilla uniform decoding strategy by over 20%, but tend to underperform on reasoning-intensive tasks, including code generation and mathematical reasoning. Semi-AR methods, which introduce a degree of left-to-right constraint on top of confidence-based sampling, substantially improve reasoning performance; for instance, achieving over 10% average gains on HumanEval and GSM8K compared to the confidence sampler. However, these gains do not transfer to planning tasks, where the imposed sequential bias becomes restrictive.

By explicitly controlling the generation order and suppressing trivial token selections, UNCODE enables MDMs to achieve strong and stable performance across both reasoning and planning tasks. On LLaDA-1.5-8B, it even achieves compara-

Methods & LLMs	HumanEval	MBPP	GSM8K	MATH500	GPQA	Countdown	Sudoku	Avg.
<i>Autoregressive LLMs</i>								
LLaMA-3.1-8B	53.1	56.7	83.9	23.8	31.0	27.0	0.0	39.4
Mistral-7B	43.9	37.0	49.4	7.2	28.1	<u>22.7</u>	0.0	26.9
Qwen-2.5-7B	78.1	62.8	<u>71.9</u>	64.2	32.8	7.7	0.0	45.3
<i>LLaDA-8B-Instruct</i>								
Uniform (2021a)	15.2	24.6	48.8	15.0	29.0	14.4	2.2	21.3
Confidence (2022)	27.4	42.4	59.1	20.8	27.9	34.0	23.8	33.6
Entropy (2025b)	28.1	42.2	60.9	11.2	28.4	33.8	1.6	29.4
Margin (2025)	32.3	42.4	58.3	19.8	28.4	33.9	<u>26.6</u>	34.5
EB-Sampler (2025)	26.8	43.3	61.2	11.6	29.5	<u>34.1</u>	24.2	33.0
Semi-AR [†] (2025b)	39.0	<u>45.2</u>	77.9	27.6	27.7	32.6	0.0	<u>35.7</u>
Fast-dLLM [†] (2025a)	35.4	44.7	<u>78.2</u>	<u>28.4</u>	28.6	23.6	0.0	34.1
UNCODE	42.1	47.8	79.2	34.8	<u>29.2</u>	36.3	29.8	42.7
<i>LLaDA-1.5-8B</i>								
Uniform (2021a)	17.7	23.0	52.7	20.0	28.1	15.8	3.4	23.0
Confidence (2022)	28.1	43.3	60.7	22.8	<u>28.7</u>	33.8	24.8	34.6
Entropy (2025b)	32.9	44.0	60.3	11.2	26.6	<u>34.7</u>	0.2	30.0
Margin (2025)	25.0	43.3	57.5	23.2	28.4	31.8	33.6	34.7
EB-Sampler (2025)	32.9	43.6	61.1	13.4	26.6	34.6	0.2	30.3
Semi-AR [†] (2025b)	39.6	<u>46.8</u>	80.7	<u>34.2</u>	26.1	32.4	0.0	<u>37.1</u>
Fast-dLLM [†] (2025a)	37.2	46.1	80.8	31.2	27.9	23.6	0.0	36.7
UNCODE	46.3	49.9	82.2	37.4	28.8	35.0	<u>33.4</u>	44.7

Table 1: Experimental results on seven different benchmarks. The best score is highlighted in **bold**, and the second-best is underlined. Following prior practices (Nie et al., 2025b), we use 4-shot settings for GSM8K and MATH500, 5-shot for GPQA and Sudoku, 0-shot for HumanEval and MBPP, and 3-shot for Countdown. Methods marked with [†] denote samplers using the Semi-AR strategy. Results on Dream-7B are provided in Table 10 (Appendix A.15).

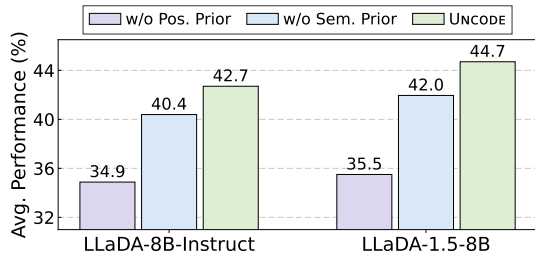


Figure 4: Ablation results of individual modules on LLaDA-8B-Instruct and LLaDA-1.5-8B, reporting the average performance across all evaluation benchmarks.

ble overall performance to Qwen-2.5-7B-Instruct. These results suggest that decoding-side improvements alone can better unlock the potential of MDMs, enabling performance that is comparable to that of autoregressive models.

6.2 Ablation Studies

We conduct ablation studies to verify the effectiveness of UNCODE; the results are shown in Figure 4.

To understand the significance of the positional trajectory prior (Pos. Prior), we compare UNCODE with models that are solely calibrated with the Semantic Informativeness Prior. As shown in Figure 4, omitting the positional trajectory prior leads to a substantial decrease in performance (42.7% \rightarrow 34.9% for LLaDA and 44.7% \rightarrow 35.5% for

LLaDA-1.5, respectively). This demonstrates the critical role of flexibly regulating the decoding order in breaking rigid boundary-first trajectory and fully harnessing the potential of MDMs in various practical scenarios. To further verify that UNCODE genuinely leverages this non-autoregressive flexibility rather than degenerating into autoregressive-like generation, we provide additional comparisons against strict fixed-order baselines in Appendix A.10. Complementing the positional adjustment, we evaluate the importance of the Semantic Informativeness Prior (Sem. Prior) by comparing UNCODE with a model calibrated exclusively by the positional prior. As shown in Figure 4, introducing the Sem. Prior further improves average performance by 2.3% and 2.7% in LLaDA and LLaDA-1.5, respectively. This demonstrates that the semantic informativeness prior, by mitigating the over-selection of trivial tokens, more effectively enhances MDM performance.

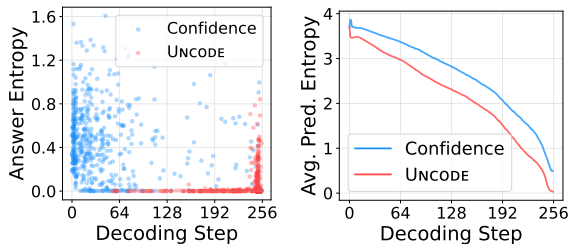
6.3 Further Analysis: Decoding Dynamics and Efficiency of UNCODE

In this section, we first delve into the internal mechanisms driving the superior performance of UNCODE, and subsequently demonstrate its compatibility with efficient decoding strategies.

Decoding Dynamics of UNCODE. As illustrated

Methods	HumanEval	MBPP	GSM8K	MATH500	GPQA	Countdown	Sudoku	Avg.
τ -leaping (2023)	17.7	38.9	55.2	16.2	28.9	32.1	3.2	27.5 (2.14x)
+ UNCODE	22.6	40.5	75.4	30.8	29.5	28.4	25.4	36.1 (1.99x)
EB-sampler (2025)	26.8	43.3	61.2	11.6	29.5	34.1	24.2	33.0 (2.32x)
+ UNCODE	41.5	46.6	79.3	35.2	28.4	36.2	25.6	41.8 (2.28x)
Fast-dLLM (2025a)	35.0	45.9	77.4	25.2	27.8	24.4	0.0	33.9 (4.21x)
+ UNCODE	36.0	48.2	77.8	29.2	28.4	36.3	0.6	36.7 (3.99x)

Table 2: Performance and efficiency analysis of UNCODE with integrated efficient decoding strategies. The last column reports the average performance, with the speedup factor relative to vanilla decoding shown in parentheses.



(a) Distribution of answer to- (b) Evolution of predictive un-
ken entropy. ken certainty.

Figure 5: Analysis of uncertainty dynamics during decoding. (a) Answer Entropy vs. Decoding Step. While the baseline (Confidence, blue) often prematurely un-masks answer tokens with high uncertainty, UNCODE (red) delays these answer tokens to later steps (bottom-right cluster), ensuring they are generated with high confidence. (b) UNCODE achieves faster global uncertainty reduction than the baseline.

in Figure 5, we conduct experiments to investigate the effectiveness of UNCODE in regulating the generation dynamics of MDMs.

To assess whether the model adheres to a logical “reasoning-then-answer” generation order, we examine answer token dynamics by tracking their unmasking steps and predictive entropies. As shown in Figure 5(a), the confidence sampler exhibits a distinct pattern of premature unmasking, where answer tokens appear at early steps despite having high entropy. This suggests the model is attempting to “guess” the answer without sufficient context. This phenomenon is driven by a rigid boundary bias, which compels the model to follow a “U-shaped” trajectory that forces answer prediction before the reasoning chain is fully constructed. Conversely, UNCODE leverages the positional trajectory prior to break this constraint, effectively deferring answer revelation until the global context is established, yielding high-confidence predictions grounded in sufficient reasoning.

We further evaluate the efficiency of global context construction by tracking the average predictive entropy of all remaining masked tokens. As shown

in Figure 5(b), the confidence sampler maintains a consistently higher level of entropy throughout the process. This suggests that its tendency to over-select trivial tokens causes the model to remain uncertain about the global context throughout the decoding process. In contrast, UNCODE consistently achieves lower predictive entropy, validating the effectiveness of our Semantic Informativeness Prior. By redirecting focus toward semantically rich tokens, UNCODE allows the model to rapidly establish a robust context, guiding effective reasoning for the subsequent generation.

Integrating UNCODE with Efficient Decoding Strategies. Furthermore, we evaluate the potential of UNCODE to serve as a plug-and-play unmasking policy within existing efficient decoding frameworks, as inference efficiency is a critical prerequisite for the practical deployment of MDMs (Israel et al., 2025). Our evaluation encompasses three representative paradigms: the deterministic τ -leaping (Chen et al., 2023), the adaptive EB-Sampler (Ben-Hamu et al., 2025), and Fast-dLLM (Wu et al., 2025a) (with KV cache optimization). Detailed implementations are provided in Appendix A.8. As demonstrated in Table 2, UNCODE seamlessly integrates with these efficient decoding frameworks, consistently outperforming their native policies across most benchmarks, yielding an average performance improvement of over 3%. Crucially, the integrated variants preserve the efficiency benefits, achieving inference speeds exceeding $2\times$ that of vanilla decoding.

7 Conclusion

In this work, we identify two decoding biases that arise when uncertainty-based samplers are applied to Masked Diffusion Models (MDMs): rigid boundary bias and trivial token bias, both of which limit the performance potential of MDMs. To address these challenges, we propose UNCODE, a decoding framework that incorporates a Positional Trajec-

tory Prior and a Semantic Informativeness Prior to mitigate these biases. Extensive experiments on three advanced MDMs across diverse benchmarks demonstrate that UNCODE consistently outperforms existing methods, achieving significant improvements in generation quality and comparable performance to autoregressive models of similar size. Additionally, we demonstrate that UNCODE is compatible with current efficient decoding frameworks, enabling high-quality, accelerated inference without the need for additional training.

Limitations

While UNCODE achieves promising results, it introduces a weight λ for the Positional Trajectory Prior. Although a lightweight search over a small set of candidates (e.g., four values) proves effective, it requires ground-truth labels for validation, which limits its applicability in fully unsupervised settings. To address this, we introduce an automated, adaptive strategy in Appendix A.13, eliminating the dependency on annotated data. Future work may further replace this heuristic with a learned controller. Next, although UNCODE integrates seamlessly with efficient decoding strategies, computing the Positional Trajectory and Semantic Informativeness priors introduces a small per-step overhead. However, as detailed in Appendix A.12, our empirical results indicate that this cost is negligible relative to the model’s forward pass, allowing the overall inference efficiency to remain comparable to standard decoding baselines.

Acknowledgments

This work is partly supported by the National Natural Science Foundation of China (No. 62576082) and the Alibaba Innovative Research Program. This work is also supported by the AI9Stars community.

References

Marianne Arriola, Aaron Gokaslan, Justin T. Chiu, Zhihan Yang, Zhixuan Qi, Jiaqi Han, Subham Sekhar Sahoo, and Volodymyr Kuleshov. 2025. [Block diffusion: Interpolating between autoregressive and diffusion language models](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*.

Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. 2021a. [Structured denoising diffusion models in discrete state-spaces](#).

In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 17981–17993.

- Jacob Austin, Augustus Odena, Maxwell I. Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie J. Cai, Michael Terry, Quoc V. Le, and Charles Sutton. 2021b. [Program synthesis with large language models](#). *ArXiv preprint*, abs/2108.07732.
- Heli Ben-Hamu, Itai Gat, Daniel Severo, Niklas Nolte, and Brian Karrer. 2025. [Accelerated sampling from masked diffusion models via entropy bounded unmasking](#). *ArXiv preprint*, abs/2505.24857.
- Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T. Freeman. 2022. [Maskgit: Masked generative image transformer](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 11305–11315. IEEE.
- Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. 2023. [Accelerating large language model decoding with speculative sampling](#). *CoRR*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, and Harri Edwards. 2021. [Evaluating large language models trained on code](#). *ArXiv preprint*, abs/2107.03374.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *ArXiv preprint*, abs/2110.14168.
- Thomas M Cover. 1999. *Elements of information theory*. John Wiley & Sons.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, and Akhil Mathur. 2024. [The llama 3 herd of models](#). *CoRR*.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, and 5 others. 2024. [The language model evaluation harness](#).
- Prateek Garg, Bhavya Kohli, and Sunita Sarawagi. 2025. [Masked diffusion models are secretly learned-order autoregressive models](#). *ArXiv preprint*, abs/2511.19152.

- Feng Hong, Geng Yu, Yushi Ye, Haicheng Huang, Huangjie Zheng, Ya Zhang, Yanfeng Wang, and Jiangchao Yao. 2025. [Wide-in, narrow-out: Revokable decoding for efficient and effective llms](#). *ArXiv preprint*, abs/2507.18578.
- Zachary Horvitz, Raghav Singhal, Hao Zou, Carles Domingo-Enrich, Zhou Yu, Rajesh Ranganath, and Kathleen McKeown. 2025. [No compute left behind: Rethinking reasoning and sampling with masked diffusion models](#). *ArXiv preprint*, abs/2510.19990.
- Daniel Israel, Guy Van den Broeck, and Aditya Grover. 2025. [Accelerating diffusion llms via adaptive parallel decoding](#). *CoRR*.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L elio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth ee Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *CoRR*.
- Yuchu Jiang, Yue Cai, Xiangzhong Luo, Jiale Fu, Jiarui Wang, Chonghan Liu, and Xu Yang. 2025. [d2 cache: Accelerating diffusion-based llms via dual adaptive caching](#). *ArXiv preprint*, abs/2509.23094.
- Samar Khanna, Siddhant Kharbanda, Shufan Li, Harshit Varma, Eric Wang, Sawyer Birnbaum, Ziyang Luo, Yanis Miraoui, Akash Palrecha, Stefano Ermon, and 1 others. 2025. [Mercury: Ultra-fast language models based on diffusion](#). *arXiv preprint arXiv:2506.17298*, 1.
- Jaeyeon Kim, Kulin Shah, Vasilis Kontonis, Sham M. Kakade, and Sitan Chen. 2025. [Train for the worst, plan for the best: Understanding token ordering in masked diffusions](#). *CoRR*.
- Sanghyun Lee, Seungryong Kim, Jongho Park, and Dongmin Park. 2025. [Lookahead unmasking elicits accurate decoding in diffusion language models](#). *ArXiv preprint*, abs/2511.05563.
- Jinsong Li, Xiaoyi Dong, Yuhang Zang, Yuhang Cao, Jiaqi Wang, and Dahua Lin. 2025a. [Beyond fixed: Training-free variable-length denoising for diffusion large language models](#). *ArXiv preprint*, abs/2508.00819.
- Tianyi Li, Mingda Chen, Bowei Guo, and Zhiqiang Shen. 2025b. [A survey on diffusion language models](#). *ArXiv preprint*, abs/2508.10875.
- Yifan Li, Kun Zhou, Wayne Xin Zhao, and Ji-Rong Wen. 2023. [Diffusion models for non-autoregressive text generation: A survey](#). In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*, pages 6692–6701. ijcai.org.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2024. [Let’s verify step by step](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Aaron Lou, Chenlin Meng, and Stefano Ermon. 2024. [Discrete diffusion modeling by estimating the ratios of the data distribution](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Richard Diehl Martinez, Z ebulon Goriely, Andrew Caines, Paula Buttery, and Lisa Beinborn. 2024. [Mitigating frequency bias and anisotropy in language model pre-training with syntactic smoothing](#). *ArXiv preprint*, abs/2410.11462.
- Shen Nie, Fengqi Zhu, Chao Du, Tianyu Pang, Qian Liu, Guangtao Zeng, Min Lin, and Chongxuan Li. 2025a. [Scaling up masked diffusion models on text](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*.
- Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. 2025b. [Large language diffusion models](#).
- Niklas Nolte, Ouail Kitouni, Adina Williams, Mike Rabbat, and Mark Ibrahim. 2024. [Transformers can navigate mazes with multi-step prediction](#). *CoRR*.
- Zhangzhi Peng, Zachary Bezemek, Sawan Patel, Jarrid Rector-Brooks, Sherwood Yao, Alexander Tong, and Pranam Chatterjee. 2025. [Path planning for masked diffusion model sampling](#). *CoRR*.
- Open R1 Project. 2025. [Openr1-math-220k: A large annotated dataset for math reasoning](#). Accessed: 2025-11-24.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2023. [GPQA: A graduate-level google-proof q&a benchmark](#). *CoRR*.
- Daria Soboleva, Faisal Al-Khateeb, Robert Myers, Jacob R Steeves, Joel Hestness, and Nolan Dey. 2023. [SlimPajama: A 627B token cleaned and deduplicated version of RedPajama](#).
- Yuxuan Song, Zheng Zhang, Cheng Luo, Pengyang Gao, Fan Xia, Hao Luo, Zheng Li, Yuehang Yang, Hongli Yu, Xingwei Qu, and 1 others. 2025. [Seed diffusion: A large-scale diffusion language model with high-speed inference](#). *ArXiv preprint*, abs/2508.02193.

- Chenglong Wang, Yang Gan, Hang Zhou, Chi Hu, Yongyu Mu, Kai Song, Murun Yang, Bei Li, Chunliang Zhang, Tongran Liu, and 1 others. 2025. [Mro: Enhancing reasoning in diffusion language models via multi-reward optimization](#). *ArXiv preprint*, abs/2510.21473.
- Xuezhi Wang and Denny Zhou. 2024. [Chain-of-thought reasoning without prompting](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Qingyan Wei, Yaojie Zhang, Zhiyuan Liu, Dongrui Liu, and Linfeng Zhang. 2025. [Accelerating diffusion large language models with slowfast: The three golden principles](#). *ArXiv preprint*, abs/2506.10848.
- Chengyue Wu, Hao Zhang, Shuchen Xue, Zhijian Liu, Shizhe Diao, Ligeng Zhu, Ping Luo, Song Han, and Enze Xie. 2025a. [Fast-dllm: Training-free acceleration of diffusion llm by enabling kv cache and parallel decoding](#). *ArXiv preprint*, abs/2505.22618.
- Xinyi Wu, Yifei Wang, Stefanie Jegelka, and Ali Jababai. 2025b. [On the emergence of position bias in transformers](#). *ArXiv preprint*, abs/2502.01951.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, and Chengyuan Li. 2024. [Qwen2 technical report](#). *ArXiv preprint*, abs/2407.10671.
- Jiacheng Ye, Jiahui Gao, Shansan Gong, Lin Zheng, Xin Jiang, Zhenguo Li, and Lingpeng Kong. 2025a. [Beyond autoregression: Discrete diffusion for complex reasoning and planning](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*.
- Jiacheng Ye, Zihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng Kong. 2025b. [Dream 7b: Diffusion large language models](#). *arXiv preprint arXiv:2508.15487*.
- Runpeng Yu, Xinyin Ma, and Xinchao Wang. 2025. [Dimple: Discrete diffusion multimodal large language model with parallel decoding](#). *ArXiv preprint*, abs/2505.16990.
- Weichao Zhang, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Yixing Fan, and Xueqi Cheng. 2024. [Pre-training data detection for large language models: A divergence-based calibration method](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 5263–5274.
- Siyao Zhao, Devansh Gupta, Qinqing Zheng, and Aditya Grover. 2025. [dl: Scaling reasoning in diffusion large language models via reinforcement learning](#). *CoRR*.
- Fengqi Zhu, Rongzhen Wang, Shen Nie, Xiaolu Zhang, Chunwei Wu, Jun Hu, Jun Zhou, Jianfei Chen, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. 2025. [Llada 1.5: Variance-reduced preference optimization for large language diffusion models](#). *CoRR*.
- Yukun Zhu, Ryan Kiros, Richard S. Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 19–27. IEEE Computer Society.

A Appendix

A.1 License

The licenses for the datasets used in this study are as follows: HumanEval, MBPP, GSM8K, and MATH500 are released under the MIT License; GPQA is released under the CC BY 4.0 License; and Countdown and Sudoku are released under the Apache License 2.0.

A.2 In-Depth Examination of the Rigid Boundary Bias Phenomenon

In Section 3, we observe that uncertainty-based samplers induce a rigid boundary bias in the decoding behavior of MDMs. Building on this empirical observation, we now examine why uncertainty-based decoding, when applied to MDMs, so consistently gives rise to a boundary-first trajectory. Our analysis identifies two interacting factors that jointly contribute to this phenomenon: the high predictability of boundary tokens and inherent local positional biases in the model.

The first factor arises from the use of fixed training templates in MDMs. Training sequences are typically augmented with explicit boundary tokens such as `<bos>` and `<eos>`, which appear at highly regular positions and with high frequency throughout the training corpus. This strong positional regularity makes boundary tokens structurally predictable, leading the model to assign them high confidence during inference. Under uncertainty-based decoding, these high-confidence boundary positions are therefore selected for unmasking at the earliest decoding steps, causing both ends of the sequence to be fixed first and resulting in a boundary-first decoding pattern.

Second, once boundary tokens are fixed, inherent local positional biases in the model’s attention mechanism (Wu et al., 2025b; Jiang et al., 2025) make it easier for the model to assign higher confidence to positions near already revealed tokens than on those farther away. As a result, the decoding process tends to propagate inward from the confirmed boundaries toward the center of the sequence, producing a characteristic “U-shaped” decoding trajectory. To test this hypothesis, we conduct a controlled intervention experiment on 100 GSM8K examples using confidence-based decoding. Before decoding begins, we manually unmask a single token at a specified relative position (25% or 75% of the sequence length), initializing it as a trivial token (e.g., a space). We then run the standard decoding

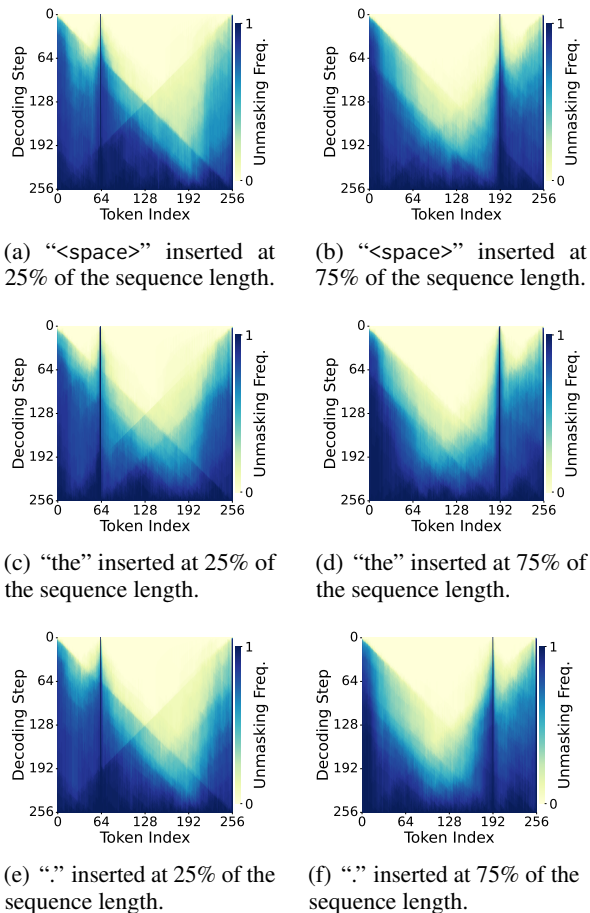
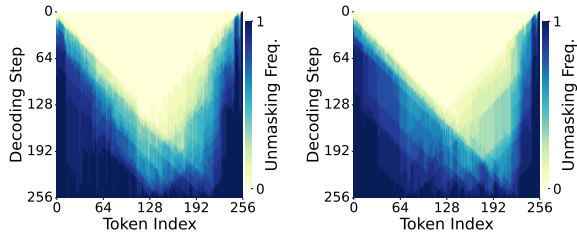


Figure 6: Unmasking dynamics under anchor insertion. A trivial high-frequency token (e.g., space, “the”, or “.”) is manually unmasked before decoding at an intermediate position (25% or 75% of the sequence length). The resulting unmasking patterns reveal a strong local positional bias, with decoding trajectories propagating outward from the inserted anchor.

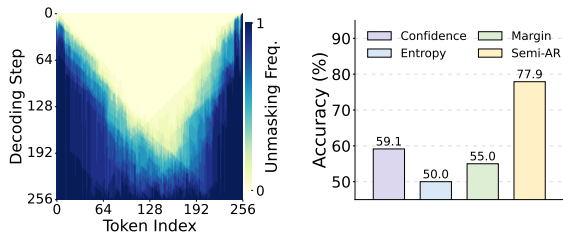
procedure and visualize the resulting unmasking dynamics. As shown in Figure 6, the artificially inserted token consistently acts as a new anchor in the decoding trajectory. During the earliest decoding steps, unmasking probabilities peak at anchor positions, including the sequence boundaries and the inserted token, and subsequently spread to neighboring positions. The same qualitative behavior is observed when other frequent tokens, such as “the” or “.”, are used as anchors.

A.3 Unmasking Dynamics of Uncertainty-Based Samplers Across Datasets

Following our observation of the rigid boundary bias on GSM8K using the confidence-based sampler in Section 3, we now systematically assess its pervasiveness across uncertainty-based decoding in MDMs. Specifically, we extend our vi-



(a) Confidence-based Sampling (b) Entropy-based Sampling



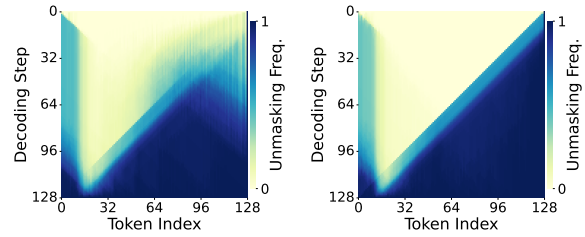
(c) Margin-based Sampling (d) Performance Comparison

Figure 7: Consistent boundary-first unmasking patterns across various uncertainty metrics. (a-c) Visualization of unmasking probability for Confidence, Entropy, and Margin samplers on GSM8K. (d) Corresponding performance comparison.

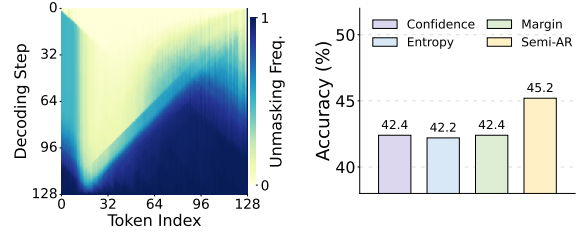
ualization of unmasking dynamics to all benchmarks employed in our main experiments, examining three common uncertainty-based sampling criteria: confidence-based, entropy-based, and margin-based samplers. The resulting heatmaps for GSM8K (Figure 7), MBPP (Figure 8), HumanEval (Figure 9), GPQA (Figure 10), Countdown (Figure 11), and Sudoku (Figure 12) consistently exhibit a boundary-first decoding pattern, regardless of the specific uncertainty metric or task domain.

Furthermore, this uniformity suggests that the bias is not an artifact of a specific calculation method (e.g., confidence vs. entropy) but rather stems from the underlying probability landscape of the model itself. Since MDMs are typically trained to predict tokens in any order, they tend to exhibit high confidence (low entropy/ large margin) for syntactically deterministic tokens, such as punctuation or structural boundaries, while remaining uncertain about the semantically dense tokens required for reasoning. Consequently, all uncertainty-based metrics inevitably prioritize these “easy” tokens, leading to the observed premature unmasking of boundary and other easy tokens.

Beyond the decoding dynamics themselves, this bias also has clear performance implications. Across reasoning-intensive tasks such as code generation and mathematical problem solving, uncertainty-based samplers underperform semi-

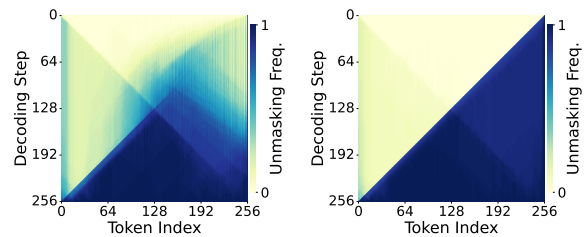


(a) Confidence-based Sampling (b) Entropy-based Sampling

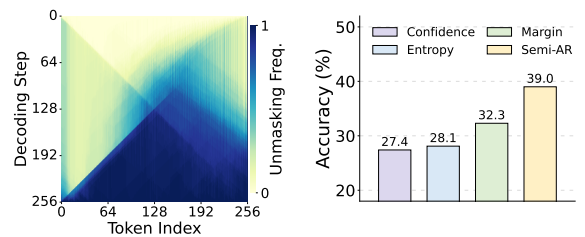


(c) Margin-based Sampling (d) Performance Comparison

Figure 8: Consistent boundary-first unmasking patterns across various uncertainty metrics. (a-c) Visualization of unmasking probability for Confidence, Entropy, and Margin samplers on MBPP. (d) Corresponding performance comparison.



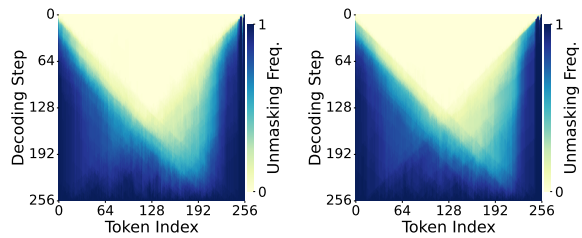
(a) Confidence-based Sampling (b) Entropy-based Sampling



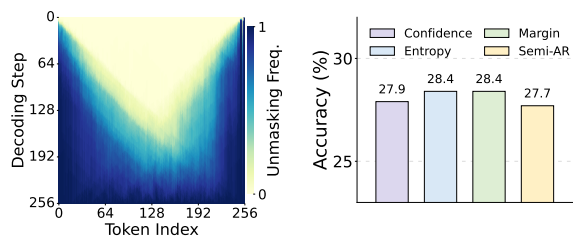
(c) Margin-based Sampling (d) Performance Comparison

Figure 9: Consistent boundary-first unmasking patterns across various uncertainty metrics. (a-c) Visualization of unmasking probability for Confidence, Entropy, and Margin samplers on HumanEval. (d) Corresponding performance comparison.

autoregressive methods by over 5 points on average. In contrast, on globally constrained planning tasks such as Countdown and Sudoku, uncertainty-based samplers outperform semi-autoregressive decoding, where strict left-to-right constraints prove limiting. These contrasting behaviors further highlight the central role of decoding order in determining MDM

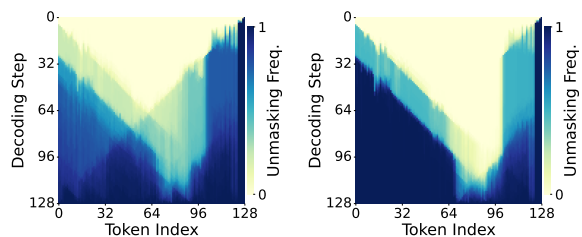


(a) Confidence-based Sampling (b) Entropy-based Sampling

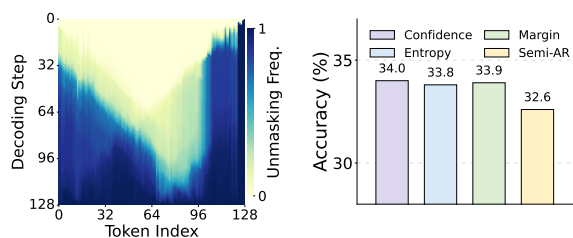


(c) Margin-based Sampling (d) Performance Comparison

Figure 10: Consistent boundary-first unmasking patterns across various uncertainty metrics. (a-c) Visualization of unmasking probability for Confidence, Entropy, and Margin samplers on GPQA. (d) Corresponding performance comparison.



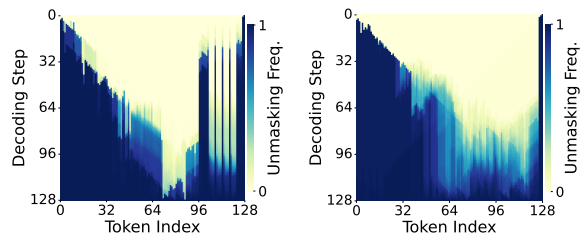
(a) Confidence-based Sampling (b) Entropy-based Sampling



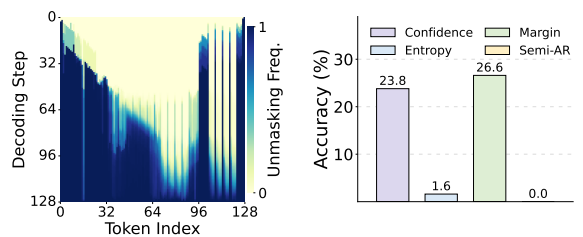
(c) Margin-based Sampling (d) Performance Comparison

Figure 11: Consistent boundary-first unmasking patterns across various uncertainty metrics. (a-c) Visualization of unmasking probability for Confidence, Entropy, and Margin samplers on Countdown. (d) Corresponding performance comparison.

performance. At the same time, they reveal a key limitation of existing approaches: both uncertainty-based samplers and semi-autoregressive methods impose a single, fixed decoding trajectory, which restricts the ability of MDMs to adapt their generation order to the structural demands of different tasks.



(a) Confidence-based Sampling (b) Entropy-based Sampling



(c) Margin-based Sampling (d) Performance Comparison

Figure 12: Consistent boundary-first unmasking patterns across various uncertainty metrics. (a-c) Visualization of unmasking probability for Confidence, Entropy, and Margin samplers on Sudoku. (d) Corresponding performance comparison.

A.4 Definition of Trivial Tokens

For our analysis, we define trivial tokens as semantically uninformative, high-frequency tokens that appear frequently in the corpus but contribute little meaningful content, following the protocol established by [Martinez et al. \(2024\)](#). These tokens primarily consist of common structural symbols, punctuation marks, and a small set of function or filler words. Notably, although `<endof text>` appears frequently in model outputs, we exclude it from this definition, as MDMs typically utilize it as a functional padding token during instruction tuning ([Nie et al., 2025b](#)). We intentionally adopt a conservative definition, restricting our scope to tokens that are unambiguously structural or uninformative, to avoid conflating triviality with task-relevant function words. The complete list of trivial tokens used in our study is provided below. Note that this definition of “trivial tokens” is solely used for the *empirical analysis* in Section 3 to demonstrate the existence of the trivial-token bias in MDMs.

List of Trivial Tokens

`<SPACE>` `\n` `\t` `.` `,` `?` `!` `:` `;` `\`
`"` `'` `$` `is` `the` `of` `a` `an`

To verify that the observed trivial-token bias is a robust phenomenon and is insensitive to the spe-

cific definition of trivial tokens, we analyze the generation behavior under two extreme definitions: (1) expanding the list to include all common NLTK stopwords, and (2) restricting the list to only whitespace tokens. As shown in Table 3, while the absolute percentages vary depending on the definition, the core conclusion remains robust: MDMs consistently generate a significantly higher proportion of low-information tokens compared to AR models across all definitions.

Definition	MDM	AR
Conservative definition (Ours)	41.7%	20.9%
Expanded w/ NLTK stopwords	49.0%	33.7%
Whitespace-only	8.9%	7.2%

Table 3: Robustness of the trivial-token bias observation across different definitions. Values denote the average ratio of trivial tokens generated by MDM and AR.

A.5 Dataset Descriptions and Statistics

We evaluate UNCODE on a diverse set of benchmarks covering code generation, mathematical reasoning, scientific question answering, and planning tasks. Dataset statistics are summarized in Table 4.

Code Generation. HumanEval (Chen et al., 2021) is a widely used benchmark for Python code generation, where models are required to complete function implementations from natural language descriptions and are evaluated using pass@1 based on unit tests. MBPP (Mostly Basic Programming Problems) (Austin et al., 2021b) evaluates code generation on a diverse collection of entry-level programming tasks, also assessed using pass@1.

Mathematical Reasoning. GSM8K (Cobbe et al., 2021) consists of grade-school-level math word problems that require multi-step numerical reasoning, and performance is measured by exact-match accuracy on the final answer. MATH500 (Lightman et al., 2024) is a curated subset of the MATH dataset containing challenging competition-level problems, evaluating advanced mathematical reasoning via accuracy.

Scientific Question Answering. GPQA (Rein et al., 2023) is a scientific QA benchmark covering biology, physics, and chemistry, designed to test domain knowledge and multi-step reasoning, with accuracy as the evaluation metric.

Planning Tasks. Countdown (Nolte et al., 2024) is a planning-oriented task where models must construct a sequence of arithmetic operations to reach a target number from a given set of numbers. Su-

Dataset	Task Type	#Test	Metric
HumanEval	Code Generation	164	Pass@1
MBPP	Code Generation	427	Pass@1
GSM8K	Math Reasoning	1,319	Acc.
MATH500	Math Reasoning	500	Acc.
GPQA	Science QA	448	Acc.
Countdown	Planning	1,000	Acc.
Sudoku	Planning	501	Acc.

Table 4: Statistics of evaluation datasets.

doku (Ye et al., 2025a) evaluates global constraint satisfaction by requiring models to complete partially filled Sudoku grids. Performance on both planning tasks is measured by accuracy.

A.6 Details of Implementation

In this subsection, we provide a detailed description of our experimental setup (§5) and the implementation specifics of our proposed UNCODE.

All experiments were conducted on four NVIDIA A100 GPUs with 80GB of memory. To ensure reproducibility, we use a fixed random seed of 42. During inference, the temperature is set to 0.0 to ensure deterministic generation. Following Nie et al. (2025b), we set the maximum response length to 256 tokens for HumanEval, GSM8K, and GPQA; 128 tokens for MBPP, Countdown, and Sudoku; and 1024 tokens for MATH500. For the positional regularization coefficient λ , we consider a small set of fixed values $\{0, 0.25, 0.5, 1.0\}$ to control the degree of sequential bias in decoding. When $\lambda = 0$, no positional trajectory prior is applied, resulting in fully order-agnostic decoding. Empirically, setting $\lambda = 1.0$ yields behavior that closely resembles strictly left-to-right decoding. Based on these observations, we set $\lambda = 0$ for Sudoku to favor globally coordinated planning, $\lambda = 0.25$ for the remaining tasks to encourage more progressive decoding, and $\lambda = 0.5$ for Countdown. All λ values are fixed across backbone models.

Evaluation Framework. To ensure a fair and rigorous comparison, all baseline methods were evaluated using the same evaluation procedures. For GSM8K and GPQA, we employed the Im-evaluation-harness framework (Gao et al., 2024). For other tasks requiring specialized metrics or formats, we adopted publicly available or officially recommended evaluation scripts, following (Nie et al., 2025b; Zhao et al., 2025; Wang et al., 2025).

Prompting Strategies. For GSM8K, GPQA, and MATH500, we follow Nie et al. (2025b)

and use the standard prompts provided in the Im-evaluation-harness framework. For Sudoku, we follow the evaluation protocol of [Zhao et al. \(2025\)](#) and use their corresponding prompt. For HumanEval and MBPP, since [Nie et al. \(2025b\)](#) does not release the corresponding prompt templates, we adopt the prompt format from [Wang et al. \(2025\)](#). The prompt is shown below:

Prompt for HumanEval and MBPP

Role: You are a professional Python coding assistant

Task: Complete the following function implementation strictly and clearly without any additional comments or explanations.

For the Countdown task, we follow [Wang et al. \(2025\)](#) and adopt a prompt that explicitly instructs the model to construct a sequence of arithmetic operations leading to the target number, as shown below.

Prompt for Countdown Task

For the given numbers, find a sequence of arithmetic operations that results in the target number. Show your reasoning and conclude with "The answer is: "

A.7 Implementation Details of Baselines

This subsection describes the implementation details of the baseline decoding strategies evaluated in our experiments.

Uniform Sampler. This strategy serves as the most basic decoding baseline for MDMs. At each decoding step, a fixed number of masked positions are selected uniformly at random and unmasked.

Confidence-Based Sampler. Confidence-based decoding is one of the most widely used sampling strategies for MDMs and is adopted by models such as LLaDA ([Nie et al., 2025b](#)). At each decoding step t , it assigns an unmasking priority to every masked position based on the model’s predictive confidence, and then selects the positions with the highest confidence to unmask. Formally, let x_t denote the partially revealed sequence at step t . For a masked position i , the confidence score is defined as the maximum probability of the most likely token under the model’s predictive distribution:

$$s_{\text{conf},t}^i = \max_{v \in \mathcal{V}} p_{\theta}(v | x_t, i), \quad (8)$$

where \mathcal{V} denotes the vocabulary. The sampler then unmask the position with the largest score. Unless otherwise specified, we unmask one position per decoding step for this sampler.

Entropy-Based Sampler. This strategy evaluates the uncertainty of the entire predictive distribution by using entropy as a proxy ([Ben-Hamu et al., 2025](#)) at each position. A lower entropy indicates a more peaked and confident distribution. The score is given by the negative entropy of the predictive distribution after applying the softmax function:

$$s_{\text{ent},t}^i = \sum_{v \in \mathcal{V}} p_{\theta}(v | x_t, i) \log p_{\theta}(v | x_t, i). \quad (9)$$

Margin-Based Sampler. This alternative measures the model’s uncertainty by computing the probability margin between the two most confident candidate tokens at each position ([Kim et al., 2025](#)). A larger margin indicates a more decisive prediction. The score is defined as:

$$s_{\text{margin},t}^i = p_{\theta}(v_1 | x_t, i) - p_{\theta}(v_2 | x_t, i), \quad (10)$$

where v_1 and v_2 denote the two most likely tokens for position i according to the model’s predictive distribution $p_{\theta}(\cdot | x_t, i)$.

EB-Sampler. The Entropy-Bounded Sampler (EB-Sampler) ([Ben-Hamu et al., 2025](#)) accelerates generation by unmasking a variable number of tokens at each step. The number of tokens unmasked is controlled by an error-tolerance hyperparameter γ , which constrains the total uncertainty to preserve generation quality. At each iteration, masked positions in \mathcal{M}_t are ranked by their entropy, and the largest subset is selected such that the cumulative entropy, used as a proxy for joint uncertainty, remains bounded as follows:

$$\sum_{i \in \mathcal{M}_t} H(p_{\theta}(\cdot | x_t, i)) - \max H(p_{\theta}(\cdot | x_t, i)) \leq \gamma, \quad (11)$$

where $H(\cdot)$ denotes the token entropy. This strategy enables more aggressive parallel decoding when predictions are confident. Following [Ben-Hamu et al. \(2025\)](#), we set $\gamma = 0.01$ for all datasets in our experiments.

Semi-Autoregressive Sampler. The semi-autoregressive sampler extends the confidence-based sampler by introducing a partial left-to-right constraint. Specifically, the response is divided into multiple contiguous blocks of equal size. Within each block, unmasking follows the confidence-based sampling strategy, allowing tokens to be revealed in an order-agnostic manner. Across blocks,

decoding proceeds sequentially from left to right, so that each block is generated conditioned on all previously completed blocks. In our experiments, we set the number of blocks to 8 for all tasks.

Fast-dLLM. Fast-dLLM (Wu et al., 2025a) builds upon the semi-autoregressive decoding paradigm and further introduces a confidence-aware parallel decoding scheme together with a KV-cache design tailored for MDMs to accelerate inference. Unlike samplers that unmask a fixed number of tokens at each step, Fast-dLLM dynamically unmask all positions whose predictive confidence exceeds a predefined threshold ϵ . As a result, the number of tokens revealed at each decoding step is adaptive and varies with the model’s confidence distribution. Formally, at decoding step t , all masked positions i satisfying the following condition are selected for unmasking:

$$\max_{v \in \mathcal{V}} p_{\theta}(v | x_t, i) > \epsilon, \quad (12)$$

are unmasked in parallel. Following Wu et al. (2025a), we set the confidence threshold to $\epsilon = 0.9$ in all experiments.

A.8 Details on Integrating with Efficient Decoding Strategies

In this subsection, we provide implementation details for integrating UNCODE with three efficient decoding strategies, as discussed in Section 6.3. For τ -leaping (Chen et al., 2023), we follow the original protocol and unmask a fixed number of tokens (τ) at each decoding step. In our experiments, we set $\tau = 2$. The specific positions to unmask are selected according to the calibrated score used by UNCODE. For the EB-Sampler (Ben-Hamu et al., 2025), we replace the entropy-based priority with our calibrated score for ranking masked positions (Eq. 5), while retaining the original cumulative budget constraint defined by the error tolerance. We set the error tolerance parameter to $\gamma = 0.01$ in all experiments. For Fast-dLLM (Wu et al., 2025a), we apply the unmasking threshold to the calibrated score used in our main experiments, rather than to the raw predictive confidence, enabling positions with high calibrated scores to be unmasked in parallel. We set the threshold to $\epsilon = 0.6$. To further accelerate decoding, we adopt the prefix KV cache mechanism proposed in Fast-dLLM.

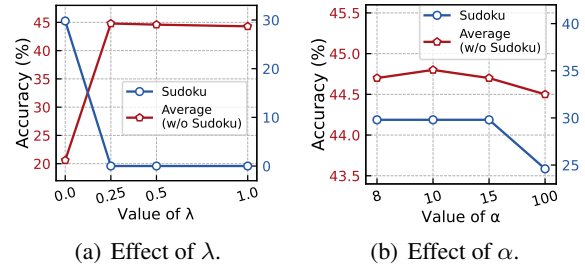


Figure 13: Analysis of decoding behavior under different λ and α settings. We find that λ substantially affects the induced decoding trajectory, whereas performance remains largely stable across a broad range of α , suggesting that UNCODE is not sensitive to the clipping threshold. To examine task-specific effects, we report average results across non-Sudoku datasets (blue) and separately for Sudoku (red).

A.9 Decoding Behavior under Different Positional and Clipping Settings

In this subsection, we analyze the effect of two key hyperparameters in UNCODE on decoding behavior, namely the positional coefficient λ and the clipping threshold α . Rather than identifying a single optimal setting, this analysis aims to characterize the decoding behaviors induced by different parameter choices. The results are summarized in Figure 13.

Figure 13(a) illustrates that the positional coefficient λ governs the degree of sequential inductive bias in decoding. For most tasks except Sudoku, moderate values (e.g., $\lambda = 0.25$) strike a balance between order-agnostic generation and progressive decoding, leading to strong overall performance. In contrast, larger values impose a stronger left-to-right constraint, which reduces decoding flexibility and can degrade performance on tasks that benefit from parallel or non-sequential reasoning. For globally constrained planning tasks such as Sudoku, smaller values of λ are preferable, as they avoid prematurely enforcing a sequential decoding order. These results suggest that different tasks naturally align with different decoding trajectories, underscoring the importance of decoding order in MDMs, consistent with observations in Kim et al. (2025).

Figure 13(b) examines the effect of the clipping threshold α . Performance remains stable across a broad range of values, except for extreme settings (e.g., $\alpha = 100$), indicating that UNCODE is robust to the choice of α . In our experiments, we fix α to 10 for simplicity and consistency across tasks.

Method	HumanEval	MBPP	GSM8K	Math500	GPQA	Countdown	Sudoku	Avg.
Confidence	27.4	42.4	59.1	20.8	27.9	34.0	23.8	33.6
L-2-R + Sem. Prior	40.2	45.2	78.2	32.8	27.9	35.7	0.0	37.1
Semi-AR + Sem. Prior	40.2	44.7	78.8	33.2	27.5	32.8	0.0	36.7
UNCODE	42.1	47.8	79.2	34.8	29.2	36.3	29.8	42.7

Table 5: Comparison of UNCODE against fixed-order baselines augmented with the semantic prior.

A.10 Validation of Non-Autoregressive Flexibility in UNCODE

To further isolate the performance contribution of the Positional Trajectory Prior and verify whether UNCODE leverages the flexibility of MDMs, we compare it against strict fixed-order baselines augmented with the semantic informativeness prior. Specifically, we evaluate a pure left-to-right strategy (L-2-R + Sem. Prior) and a block-wise left-to-right strategy (Semi-AR + Sem. Prior), reporting the results across all seven benchmarks on LLaDA-8B-Instruct.

As detailed in Table 5, UNCODE significantly outperforms both fixed-order baselines (L-2-R + Sem. and Semi-AR + Sem.) by 5.6% and 6.0%, respectively. This confirms that the dynamic decoding trajectory of the Positional Trajectory Prior enables UNCODE to leverage non-autoregressive flexibility of MDMs, which is essential for superior performance.

A.11 Long-Sequence Generalization

To evaluate whether the proposed method generalizes to longer sequences, we extend the maximum generation length to 512 while keeping all other settings unchanged. Table 6 summarizes the results on LLaDA-8B-Instruct.

As shown in Table 6, UNCODE exhibits robust performance when the generation length is extended, achieving an average score of 42.6%. This indicates that the proposed priors effectively guide the MDM to leverage the expanded generation capacity, preserving a substantial performance advantage over the baselines.

A.12 Inference Overhead Analysis

To assess the computational efficiency of UNCODE, we provide quantitative latency measurements comparing our approach against the standard confidence-based sampler. As shown in Table 7, introducing both the Positional Trajectory Prior and the Semantic Informativeness Prior yields only a negligible per-step overhead.

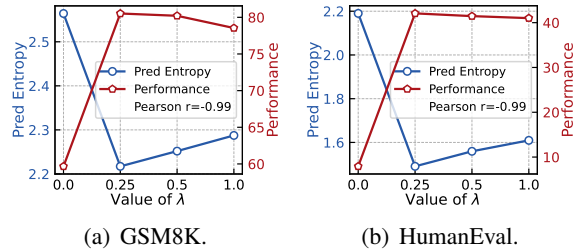


Figure 14: Correlation analysis between Mean Predictive Entropy and task performance across varying λ values. The blue curve denotes mean predictive entropy (left axis), while the red curve indicates performance (right axis) on (a) GSM8K and (b) HumanEval. The strong negative correlation ($r \approx -0.99$) empirically validates that lower mean predictive entropy serves as a reliable indicator of higher reasoning quality.

A.13 Adaptive λ Selection via Mean Predictive Entropy Minimization

In this subsection, we propose an adaptive strategy for selecting the positional coefficient λ . We first define the mean predictive entropy and show that it correlates well with generation quality across different λ settings. We then leverage entropy minimization to select λ automatically, without requiring manual tuning or ground-truth labels.

Definition of Mean Predictive Entropy. For an individual data instance, we define the *sequence-level predictive entropy* as the average entropy of the predictive distributions over all masked tokens at each decoding step. Specifically, for each decoding step t , we compute the entropy $\mathcal{H}_t(i)$ for each masked token $i \in \mathcal{M}_t$, where \mathcal{M}_t is the set of masked positions at step t , and $\mathcal{H}_t(i)$ denotes the entropy of the predictive distribution for token i at step t . The *sequence-level predictive entropy* for a given instance is then computed by averaging the entropies across all decoding steps t :

$$\mathcal{H}_{\text{seq}} = \frac{1}{T} \sum_{t=1}^T \frac{1}{|\mathcal{M}_t|} \sum_{i \in \mathcal{M}_t} \mathcal{H}_t(i), \quad (13)$$

where T is the total number of decoding steps, and $|\mathcal{M}_t|$ is the number of masked positions at step t .

Method	HumanEval	MBPP	GSM8K	Math500	GPQA	Countdown	Sudoku	Avg.
Confidence	18.9	40.5	58.2	24.0	27.2	22.8	0.0	27.4
Semi-AR	37.8	39.1	75.4	33.2	26.8	27.3	0.0	34.2
Fast-dLLM	38.4	40.3	75.7	34.2	26.8	33.8	0.0	35.6
UNCODE	43.3	48.2	79.9	34.4	30.4	35.7	26.0	42.6

Table 6: Performance under an extended maximum generation length of 512.

Method	MBPP	Math500
Confidence	1.03	48.88
UNCODE	1.04	49.08

Table 7: Latency comparison (seconds per iteration) on LLaDA-8B-Instruct with 4×A100 GPUs.

Next, to compute the *mean predictive entropy* at the dataset level, we extend this definition to the entire dataset. Let $\mathcal{D}_{\text{cal}} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ be the calibration subset of N instances. The *mean predictive entropy* is then defined as the average of the sequence-level predictive entropies across all instances in the calibration set:

$$\bar{\mathcal{H}} = \frac{1}{N} \sum_{j=1}^N \mathcal{H}_{\text{seq}}^{(j)}, \quad (14)$$

where $\mathcal{H}_{\text{seq}}^{(j)}$ is the sequence-level predictive entropy for the j -th instance. The resulting $\bar{\mathcal{H}}$ serves as a proxy for the overall generation quality, which is used for selecting the guidance weight λ in the subsequent strategy.

Empirical Validation of the Minimum Entropy Principle. To assess whether mean predictive entropy is a reliable proxy for output quality, we perform a correlation analysis on a calibration subset of 100 instances randomly sampled from GSM8K and HumanEval, as shown in Figure 14. For each candidate $\lambda \in \Lambda = \{\lambda_1, \dots, \lambda_K\}$, we decode the entire subset and compute: (i) the mean predictive entropy $\bar{\mathcal{H}}(\lambda)$, and (ii) the corresponding task metric (Accuracy for GSM8K, pass@1 for HumanEval). We then compute the Pearson correlation coefficient (r) between the mean predictive entropy and task performance across the K candidate λ values. As shown in Figure 14(a) for GSM8K and Figure 14(b) for HumanEval, we observe a strong negative correlation, indicating that lower mean entropy corresponds to better task performance. This finding motivates entropy minimization as a label-free criterion for selecting λ .

Adaptive Selection via Entropy Minimization. Based on the observation above, we select the

guidance weight λ by minimizing the mean predictive entropy on a small calibration set. Let $\mathcal{D}_{\text{cal}} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ denote a small subset of validation inputs and $\Lambda = \{\lambda_1, \dots, \lambda_K\}$ be the candidate set. For each $\lambda \in \Lambda$, we decode \mathcal{D}_{cal} to obtain outputs $\{\mathbf{y}_\lambda^{(j)}\}_{j=1}^N$. We then compute the mean predictive entropy:

$$\bar{\mathcal{H}}(\lambda) = \frac{1}{N} \sum_{j=1}^N \mathcal{H}_{\text{seq}}(\mathbf{y}_\lambda^{(j)}), \quad (15)$$

where $\mathcal{H}_{\text{seq}}(\mathbf{y})$ denotes the *average predictive entropy* of a decoded sequence. Finally, we select the entropy-minimizing weight:

$$\lambda^* = \arg \min_{\lambda \in \Lambda} \bar{\mathcal{H}}(\lambda). \quad (16)$$

By selecting λ via entropy minimization on a small calibration subset, this strategy avoids manual tuning and does not require ground-truth labels during selection.

Robustness Under Limited Calibration Data.

We further evaluate the robustness of the adaptive strategy under limited calibration data. Specifically, we randomly sample only 10 instances to form a calibration set and compute the mean predictive entropy for each candidate λ , alongside the model’s corresponding performance. As shown in Table 9, even when relying on merely 10 instances to estimate the mean predictive entropy, our adaptive strategy consistently selects the optimal λ . This demonstrates that entropy minimization remains effective even with highly limited unlabeled calibration data.

A.14 Impact of Calibration Corpus

In this subsection, we investigate the impact of the calibration corpus \mathcal{D}' on UNCODE through a series of controlled ablation studies. Specifically, we examine two key dimensions: (1) the source of the calibration corpus, and (2) the corpus size. The results are summarized in Table 8.

We first assess the robustness of UNCODE across varied calibration corpora. Specifically, we compare the default C4 (16GB) against two alternatives: (i) a domain-specific mixture of BookCorpusOpen and OpenR1-Math-220K (denoted as Book

Calibration Corpus	HumanEval	MBPP	GSM8K	MATH500	GPQA	Countdown	Sudoku	Avg.
<i>LLaDA-Instruct-8B</i>								
Book & Code (2015; 2025)	42.1	47.3	79.2	35.0	27.9	36.3	24.4	41.7
SlimPajama (2023)	40.9	47.1	79.1	35.2	28.4	36.3	30.8	42.5
C4 (16 GB) (2020)	42.1	47.8	79.2	34.8	29.2	36.3	29.8	42.7
C4 (50 GB) (2020)	42.2	47.8	79.2	36.0	29.0	36.3	30.0	42.9
C4 (200 GB) (2020)	42.9	47.8	79.2	36.2	29.0	36.3	30.0	43.1
C4 (400 GB) (2020)	42.9	47.8	79.2	36.2	29.0	36.3	30.0	43.1

Table 8: Evaluation of UNCODE across various calibration corpora. ‘Book & Code’ refers to a combination of BookCorpusOpen and OpenR1-Math-220K datasets. C4 (x GB) denotes a sampled subset of the C4 dataset, where x represents the size of the subset in gigabytes.

λ	HumanEval	MBPP	GSM8K	Math500	GPQA	Countdown	Sudoku	Avg.
0.00	7.93 (2.397)	33.5 (1.538)	59.7 (2.565)	4.0 (2.482)	29.0 (3.852)	34.0 (2.532)	29.8 (1.288)	28.28
0.25	42.1 (1.465)	47.8 (0.495)	79.2 (2.014)	34.8 (1.187)	29.2 (3.821)	34.7 (1.641)	0.2 (2.366)	38.29
0.50	41.5 (1.556)	47.8 (0.497)	78.9 (2.017)	34.0 (1.318)	28.8 (3.871)	36.3 (1.529)	0.0 (2.791)	38.19
1.00	40.2 (1.575)	45.4 (0.543)	77.7 (2.041)	34.4 (1.266)	28.4 (3.895)	35.7 (1.506)	0.0 (2.474)	37.93

Table 9: Adaptive λ selection with only 10 unlabeled calibration instances. Values in parentheses are mean predictive entropy measured on the 10-instance calibration subset.

Methods & LLMs	HumanEval	MBPP	GSM8K	MATH500	GPQA	Countdown	Sudoku	Avg.
<i>Autoregressive LLMs</i>								
LLaMA-3.1-8B-Instruct	<u>53.1</u>	<u>56.7</u>	83.9	<u>23.8</u>	<u>31.0</u>	27.0	0.0	39.4
Mistral-7B-Instruct	43.9	37.0	49.4	7.2	28.1	<u>22.7</u>	0.0	26.9
Qwen-2.5-7B-Instruct	78.1	62.8	<u>71.9</u>	64.2	32.8	7.7	0.0	45.3
<i>LLaDA-8B-Instruct</i>								
Uniform (2021a)	15.2	24.6	48.8	15.0	29.0	14.4	2.2	21.3
Confidence (2022)	27.4	42.4	59.1	20.8	27.9	34.0	23.8	33.6
Entropy (2025b)	28.1	42.2	60.9	11.2	28.4	33.8	1.6	29.4
Margin (2025)	32.3	42.4	58.3	19.8	28.4	33.9	<u>26.6</u>	34.5
EB-Sampler (2025)	26.8	43.3	61.2	11.6	29.5	<u>34.1</u>	24.2	33.0
Semi-AR [†] (2025b)	<u>39.0</u>	<u>45.2</u>	77.9	27.6	27.7	32.6	0.0	<u>35.7</u>
Fast-dLLM [†] (2025a)	35.4	44.7	<u>78.2</u>	<u>28.4</u>	28.6	23.6	0.0	34.1
UNCODE	42.1	47.8	79.2	34.8	<u>29.2</u>	36.3	29.8	42.7
<i>LLaDA-1.5-8B</i>								
Uniform (2021a)	17.7	23.0	52.7	20.0	28.1	15.8	3.4	23.0
Confidence (2022)	28.1	43.3	60.7	22.8	<u>28.7</u>	33.8	24.8	34.6
Entropy (2025b)	32.9	44.0	60.3	11.2	26.6	<u>34.7</u>	0.2	30.0
Margin (2025)	25.0	43.3	57.5	23.2	28.4	<u>31.8</u>	33.6	34.7
EB-Sampler (2025)	32.9	43.6	61.1	13.4	26.6	34.6	0.2	30.3
Semi-AR [†] (2025b)	<u>39.6</u>	<u>46.8</u>	80.7	<u>34.2</u>	26.1	32.4	0.0	<u>37.1</u>
Fast-dLLM [†] (2025a)	37.2	46.1	80.8	31.2	27.9	23.6	0.0	36.7
UNCODE	46.3	49.9	82.2	37.4	28.8	35.0	<u>33.4</u>	44.7
<i>Dream-v0-Instruct-7B</i>								
Uniform (2021a)	17.7	31.9	31.5	17.0	32.8	4.1	0.2	19.3
Confidence (2022)	27.4	41.5	45.4	20.8	<u>35.3</u>	19.8	0.0	27.2
Entropy (2025b)	26.2	42.4	36.8	17.0	33.5	<u>19.0</u>	0.0	25.0
Margin (2025)	<u>28.1</u>	41.7	48.3	<u>22.0</u>	35.7	<u>19.0</u>	0.0	27.8
EB-Sampler (2025)	26.8	<u>43.6</u>	37.5	17.4	33.3	18.6	0.0	25.3
Fast-dLLM (2025a)	12.8	<u>23.9</u>	46.1	19.2	34.4	11.6	0.0	21.1
UNCODE	57.9	56.4	76.4	37.8	33.9	18.4	0.0	40.1

Table 10: Experimental results on seven different benchmarks. The best score is highlighted in **bold**, and the second-best is underlined. Following prior practices (Nie et al., 2025b), we use 4-shot settings for GSM8K and MATH500, 5-shot for GPQA and Sudoku, 0-shot for HumanEval and MBPP, and 3-shot for Countdown. Methods marked with [†] denote samplers using the Semi-AR strategy. All settings follow prior works (Nie et al., 2025b; Zhao et al., 2025; Wang et al., 2025) for fair comparison.

& Code), and (ii) SlimPajama (Soboleva et al., 2023). Across these settings, we observe comparable average performance. These results indicate that UNCODE is generally robust to the selection of calibration sources.

Next, we investigate the impact of corpus size on performance stability. We sample three larger subsets from C4, sized 50GB, 200GB, and 400GB. We observe that while increasing the corpus size initially improves accuracy, the performance gains saturate beyond the 200GB scale. We hypothesize that the estimated token distribution converges once the corpus reaches a sufficient size; consequently, further enlarging the dataset provides marginal benefit for the calibration signal. Overall, we observe that UNCODE benefits from improved calibration, but does not exhibit strong sensitivity to the specific calibration corpus used.

A.15 Results on Dream

To further evaluate the generalization capability of our approach, we apply UNCODE to Dream-v0-Instruct-7B (Ye et al., 2025b), a state-of-the-art MDM developed independently of the LLaDA series (Nie et al., 2025a). As shown in Table 10, UNCODE consistently achieves the best performance across all evaluated tasks, with an average score of 40.1%. This represents a substantial improvement over the strongest uncertainty-based baseline, margin-based sampling, which reaches 27.8%. In particular, UNCODE achieves 57.9% on HumanEval and 76.4% on GSM8K, demonstrating strong performance in both code generation and mathematical reasoning tasks.

These results indicate that UNCODE is not limited to a specific model family, but instead addresses a broader limitation inherent to masked diffusion models. Across different architectures, we consistently observe a tendency for uncertainty-based decoding to prioritize boundary tokens early in the decoding process and to favor semantically trivial tokens. This behavior suggests a systematic bias in how unmasking priorities are assigned, rather than an artifact of a particular model or implementation. By incorporating both positional awareness and semantic awareness into the sampling process, UNCODE mitigates these biases and enables the model to allocate its generation capacity more effectively. As a result, the decoding process is better aligned with the underlying reasoning or planning structure of the task, leading to more coherent and logically ordered outputs.

A.16 Case Study

In this subsection, we present three case studies, as shown in Figure 15, Figure 16, and Figure 17, drawn from the GSM8K (mathematical reasoning) and HumanEval (code generation) benchmarks. The first two cases are from GSM8K, and the third is from HumanEval. These case studies compare the baseline confidence-based sampling strategy with our proposed UNCODE approach.

As shown in Table 15, due to rigid boundary bias, the confidence sampler prematurely commits to a final answer of 2975 at an early stage, before computing the correct total number of classes and students. Once this premature decision is made, subsequent tokens are forced to justify the incorrect conclusion rather than revising it, resulting in a fluent yet numerically incorrect solution. In contrast, UNCODE incorporates the Positional Trajectory Prior, which delays emitting the final answer and first completes the intermediate computations (e.g., 25 weekday classes, 8 Saturday classes, 33 total classes, 495 students, and the resulting revenue). The decoding trajectory allocates more steps to these intermediate reasoning tokens, allowing for a coherent, step-by-step derivation that leads to the correct answer of 7425.

As illustrated in Table 16, due to the trivial token bias, the confidence sampler decodes trivial tokens at critical steps (e.g., the comma is decoded at step 121, at position 76, instead of the 7 at position 75), thereby locking the reasoning space. This leads to errors, such as the comma being used as a thousand separator, causing invalid representations like 0,700. Consequently, the model fails to predict 700, ultimately resulting in a calculation error. In contrast, UNCODE effectively suppresses the selection of trivial tokens, enabling the model to arrive at the correct solution.

For the HumanEval benchmark, as shown in Table 17, the baseline’s rigid boundary bias often leads to premature sequence termination and incomplete code, resulting in a significant waste of the generation budget. In contrast, UNCODE mitigates this issue by incorporating position-aware weighting and content-aware confidence calibration, which introduce an inductive bias toward sequential reasoning. As a result, UNCODE ensures that complete and correct code is generated before sequence termination.

These case studies demonstrate that explicit trajectory control and trivial token suppression are

Case1: Mathematical Reasoning (GSM8K)	
Question	Judy teaches 5 dance classes, every day, on the weekdays and 8 classes on Saturday. If each class has 15 students and she charges \$15.00 per student, how much money does she make in 1 week?
Ground Truth	7,425
Confidence (Baseline)	Judy teaches 5 dance classes every day on weekdays, so in one week, she teaches $5 * 5 = 25$ classes . Each class has 15 students, so in one week, she teaches $25 * 15 = 375$ students . She teaches 5 dance classes every day, so in one week, she teaches $5 * 15 = 75$ students . Additionally, she teaches 8 classes on Saturday. Each class has 15 students , so on Saturday, she teaches $8 * 15 = 120$ students . In total, in one week, she teaches $75 + 120 = 195$ students . Since she charges \$15.00 per student, she makes $195 * 15 = \$2975.00$ in 1 week. ##### 2,975 < eot_id > < endoftext >
UNCODE (Ours)	Judy teaches 5 dance classes every day on weekdays, for a total of $5 \text{ days} * 5 \text{ classes/day} = 25$ classes . On Saturday, she teaches 8 classes . So in one week, she teaches a total of $25 + 8 = 33$ classes . Each class has 15 students, so she teaches $33 * 15 = 495$ students . She charges \$15 per student, so she makes $495 * 15 = \$7425.00$ in one week. ##### 7,425 < eot_id > < endoftext >

Figure 15: A case study on the GSM8K dataset illustrating how the generation order impacts mathematical reasoning. The baseline’s premature commitment to an incorrect answer highlights a critical failure mode, whereas UNCODE’s coherent process leads to the correct solution. The generation timing is color-coded as follows: Early Stage , Middle Stage , and Late Stage .

critical for complex problem solving. By mitigating both the rigid boundary bias and the trivial token bias, UNCODE enables the model to produce outputs that are more coherent, accurate, and complete.

Case 2: Mathematical Reasoning (GSM8K)	
Question	Jill gets paid \$20 per hour to teach and \$30 to be a cheerleading coach. If she works 50 weeks a year, 35 hours a week as a teacher and 15 hours a week as a coach, what's her annual salary?
Ground Truth	57,500
Confidence (Baseline)	<p>1. Weekly earnings from teaching:</p> <ul style="list-style-type: none"> - Jill earns \$20 per hour and works 35 hours a week. - Weekly earnings from teaching = $\\$20/\text{hour} * 35 \text{ hours/week} = \\$3,050$ <p>2. Weekly earnings from coaching:</p> <ul style="list-style-type: none"> - Jill earns \$30 per hour and works 15 hours a week. - Weekly earnings from coaching = $\\$30/\text{hour} * 15 \text{ hours/week} = \\450 <p>3. Total weekly earnings:</p> <ul style="list-style-type: none"> - Total weekly earnings = Teaching earnings + Coaching earnings - Total weekly earnings = $\\$3,050 + \\$450 = \\$3,500$ <p>4. Annual earnings:</p> <ul style="list-style-type: none"> - There are 50 weeks in a year. - Annual earnings = $\\$3,500/\text{week} * 50 \text{ weeks/year} = \\$175,000$ <p>Therefore, Jill's annual salary is \$175,000. ❌</p>
UNCODE (Ours)	<p>To calculate Jill's annual salary, we need to break down her earnings into two parts: her earnings as a teacher and her earnings as a cheerleading coach.</p> <p>1. Earnings as a teacher:</p> <ul style="list-style-type: none"> - Hourly rate: \$20 - Hours per week: 35 - Weekly earnings as a teacher: $\\$20/\text{hour} * 35 \text{ hours/week} = \\$700/\text{week}$ <p>2. Earnings as a cheerleading coach:</p> <ul style="list-style-type: none"> - Hourly rate: \$30 - Hours per week: 15 - Weekly earnings as a cheerleading coach: $\\$30/\text{hour} * 15 \text{ hours/week} = \\$450/\text{week}$ <p>3. Total weekly earnings:</p> <ul style="list-style-type: none"> - Total weekly earnings: $\\$700/\text{week} + \\$450/\text{week} = \\$1,150/\text{week}$ <p>4. Total annual earnings:</p> <ul style="list-style-type: none"> - Number of weeks per year: 50 - Total annual earnings: $\\$1,150/\text{week} * 50 \text{ weeks/year} = \\$57,500/\text{year}$ <p>Therefore, Jill's annual salary is \$57,500. ✓</p>

	Step	Pos.	Token	Confidence
Cur token	121	76	,	0.79
Candidate	121	74	\$	0.63
Candidate	121	75	7	0.42

Figure 16: A case study on the GSM8K dataset illustrating how the trivial token bias impacts mathematical reasoning. The baseline decodes trivial tokens at critical steps (e.g., the comma is decoded at step 121, at position 76, instead of the 7 at position 75), locking the reasoning space (e.g., the comma, used as a thousand separator, causes failure as 0,700 is invalid), and leading to incorrect predictions. In contrast, UNCODE suppresses the selection of trivial tokens, resulting in the correct solution. The generation timing is color-coded as follows: **Early Stage** and **Middle Stage**.

Case 3: Code Generation (HumanEval)	
Task	Implement the body of the Python function 'has_close_elements'. The function should take a list of floats and a threshold, and return 'True' if any two distinct elements in the list are closer than the threshold, and 'False' otherwise.
Confidence (Baseline)	<pre> python def has_close_elements(numbers: List[float], threshold: float) -> bool: for num in numbers: for other in numbers: if abs(num - other) < threshold: return True return False </pre> <p><endoftext><endoftext> ... (many EOS tokens)</p> <p><i>Token Allocation: Code (~52 tokens) vs. Prematurely Generated EOS (~204 tokens)</i></p>
UNCODE (Ours)	<pre> python from typing import List def has_close_elements(numbers: List[float], threshold: float) -> bool: """ (A rigorous docstring with examples is generated here) """ for i in range(len(numbers)): for j in range(i + 1, len(numbers)): if abs(numbers[i] - numbers[j]) < threshold: return True return False </pre> <p><endoftext><endoftext></p> <p><i>Token Allocation: Code (~178 tokens) vs. EOS Generated After Reasoning (~78 tokens)</i></p>

Figure 17: Case study on the HumanEval dataset for code generation. The baseline method exhibits catastrophic failure due to the premature generation of trivial tokens, whereas UNCODE, through explicit trajectory control, produces robust and complete code. The generation timing is color-coded as follows: Early Stage, Middle Stage, and Late Stage.