

MDPO: OVERCOMING THE TRAINING-INFERERENCE DIVIDE OF MASKED DIFFUSION LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Diffusion language models, as a promising alternative to traditional autoregressive (AR) models, enable faster generation and richer conditioning on bidirectional context. However, they suffer from a key discrepancy between training and inference: during inference, MDLMs progressively reveal the structure of the generated sequence by producing fewer and fewer masked tokens, whereas this structure is ignored in training as tokens are masked at random. Although this discrepancy between training and inference can lead to suboptimal performance, it has been largely overlooked by previous works, leaving closing this gap between the two stages an open problem. To address this, we frame the problem of learning effective denoising trajectories as a sequential decision-making problem and use the resulting framework to apply reinforcement learning. We propose a novel Masked Diffusion Policy Optimization (MDPO) to exploit the Markov property diffusion possesses and explicitly train the model under the same progressive refining schedule used at inference. MDPO matches the performance of the previous state-of-the-art (SOTA) method with 60× fewer gradient updates, while achieving average improvements of 9.6% on MATH500 and 54.2% on Countdown over SOTA when trained within the same number of weight updates. Additionally, we improve the remasking strategy of MDLMs as a plug-in inference replacement to overcome the limitation that the model cannot refine tokens flexibly. This training-free method, termed Running Confidence Remasking (RCR), consistently enhances performance and provides further improvements when used with MDPO. Our findings establish great potential for investigating the discrepancy between pre-training and inference of MDLMs. To support reproducibility, we will release code upon acceptance.

1 INTRODUCTION

Remarkable advances in language modeling have largely been driven by the autoregressive (AR) paradigm (Hochreiter & Schmidhuber, 1997; Vaswani et al., 2017), where tokens are generated causally. Diffusion Language Models (DLMs), as an alternative line of work, generate by iterative denoising to enable parallel prediction, bidirectional conditioning, and faster generation speed. Recently, DLMs such as Mercury Coder (Khanna et al., 2025), Google Gemini Diffusion (Google DeepMind, 2025), and Seed Diffusion (Song et al., 2025) have emerged as promising alternatives that excel at external math and coding benchmarks. Built upon theoretical insights and practical optimizations in the discrete diffusion framework (Lou et al., 2024; Sahoo et al., 2024; Ou et al., 2025), Masked Diffusion Language Models (MDLMs) lead the current development of DLMs. In the forward process of MDLMs, tokens are randomly chosen to either remain stationary or transition to a predefined absorbing state, such as a special masking token, motivated by the success of BERT (Devlin et al., 2019). Then, in the backward process, the model learns to invert the noise by predicting *all* masked tokens conditioned on partially observed (unmasked) context. For generating a sequence of tokens, current MDLMs alternate between predicting masked tokens and selectively remasking a fraction of the predictions based on heuristics and uncalibrated token-wise scores. Representative works such as LLaDA-8B (Nie et al., 2025) and Dream-7B (Ye et al., 2025) demonstrate competitive performance to their similarly sized AR counterparts.

Despite promising empirical results, we identify two fundamental, largely overlooked problems of MDLMs that hinder effective denoising trajectories. (1) The *training–inference divide*: at inference

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107

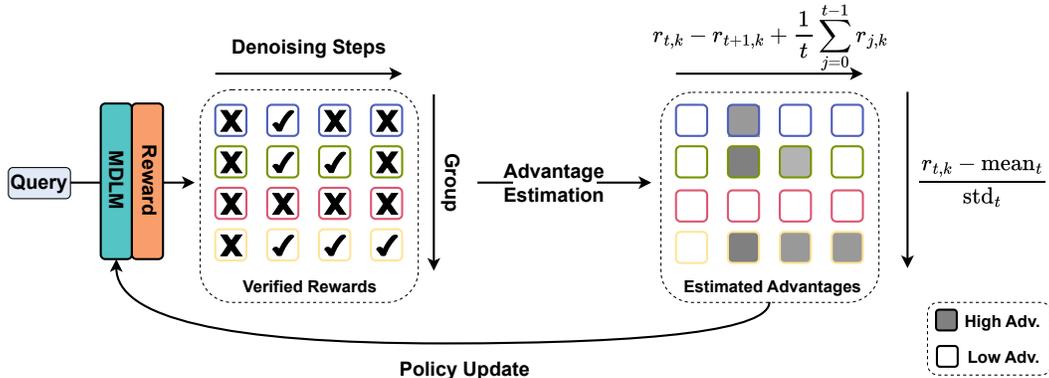


Figure 1: **MDPO** generates a group of answers given a query for RL rollouts. Then all completions at intermediate and final steps are verified with a reward model. Based on verified rewards, MDPO estimates the advantage of step t by considering rewards of the other steps in the current rollout and step t from other rollouts in the group. These estimated advantages are used for policy optimization.

time, MDLMs follow a model-dependent, confidence-guided remasking schedule that progressively reveals the structure of the generated sequence. By contrast, tokens are randomly masked at sampled diffusion steps in training, ignoring the progressive schedule. This mismatch prevents the model from learning effective denoising trajectories and often causes over-denoising: correct intermediate solutions are ‘refined’ into wrong final solutions (see Fig. 3 for statistical analysis). (2) The common remasking approaches freeze tokens once predicted, making it impossible to revise early low-confidence predictions at later denoising steps. We address these problems with complementary methods in this work. Masked Diffusion Policy Optimization (MDPO) casts denoising as a multi-step decision-making process, explicitly training the model under the same progressive, model-dependent remasking schedule used at inference. By using reinforcement learning with intermediate rewards to optimize these trajectories, MDPO aligns training with actual inference dynamics, closing the training–inference divide. Running Confidence Remasking (RCR) is a training-free decoding strategy that tracks per-position confidence over time, enabling flexible remasking and revision of low-confidence tokens at any step. In summary, our contributions are threefold:

- **Over-denoising:** We observe a novel and previously unknown phenomenon in MDLMs: models occasionally produce correct answers at intermediate steps but ‘refine’ them into incorrect results. We refer to this phenomenon as over-denoising, which inspires us to supervise models not only on the final results but also intermediate denoising steps.
- **Masked Diffusion Policy Optimization (MDPO):** We propose MDPO to learn effective denoising trajectories without direct supervision. By exploiting the fact that MDLMs yield complete generations at every inference step, MDPO optimizes the model with policy gradients via intermediate-step rewards. Unlike prior RL fine-tuning on MDLMs, MDPO explicitly targets addressing the *training-inference divide* overlooked by previous works.
- **Running Confidence Remasking (RCR):** Instead of freezing tokens based on their single-step confidence, RCR allows flexible remasking by continuously tracking the running lowest confidence over denoising trajectories. Our experiments demonstrate consistently improved performance of using RCR on both LLaDA pre-trained as well as MDPO fine-tuned models.

Empirical results on challenging mathematical and reasoning benchmarks show that both MDPO and the training-free RCR substantially improve generation quality and sample efficiency. Our findings underscore the value of addressing the training-inference divide for masked diffusion language models and open new directions for research in this emerging paradigm.

2 METHODOLOGY

2.1 PRELIMINARY: MASKED DIFFUSION LANGUAGE MODELS

We introduce our methods based on the formulation of LLaDA (Nie et al., 2025), a pre-trained MDLM that models clean data x_0 via a discrete diffusion process (Ou et al., 2025; Austin et al.,

2021), where tokens are progressively masked in the forward process and then recovered by iterative prediction in the reverse process. While specific methodological variations may exist, the methods proposed in this paper are broadly applicable to the majority of existing MDLM frameworks.

Training: The training objective optimizes a mask predictor $p_\theta(x_0 | \bar{x}_\tau)$ that takes partially masked data \bar{x}_τ ¹ as input and predicts all masked tokens simultaneously using a cross-entropy loss:

$$\mathcal{L}(\theta) \triangleq -\mathbb{E}_{x_0, \tau} \left[\frac{1}{\tau} \sum_{i=1}^L \mathbb{1}[\bar{x}_\tau^i = \mathcal{M}] \log p_\theta(x_0^i | \bar{x}_\tau) \right] \quad \bar{x}_\tau \triangleq \tilde{\gamma}(x_0, \tau) \quad (1)$$

Here, the expectation is computed over the unmasked ground truth sequences $\{x_0\}$ in the dataset and a randomly drawn masking ratio $\tau \sim U(0, 1)$. $\tilde{\gamma}(x_0, \tau)$ is a masking function that sets each token of x_0 to the mask token \mathcal{M} independently with probability τ . L is the sequence length and superscript i for both x and \bar{x} denotes the i 'th element in the sequence. The cross-entropy loss is only computed on the masked tokens as indicated by the function $\mathbb{1}[\cdot]$. In practice, $p_\theta(x_0 | \bar{x}_\tau)$ often depends on an additional input prompt which we drop here for notational clarity.

Inference: Inference of LLaDA involves iterative refinement from a fully masked input, progressively reducing the number of masked tokens at each step t . This is necessary, as the initial noise contains no structure and the mapping from noise to the generated sequence is highly non-linear and multimodal (Xiao et al., 2022; Song et al., 2021). Hence, multiple iterations (each conditioning on the structure generated so far) are required to achieve coherent, contextually and syntactically consistent outputs (Nie et al., 2025; Feng et al., 2025). Given T denoising steps, the inference starts from a fully masked sequence $\bar{x}_T = (\mathcal{M}, \dots, \mathcal{M})$ and alternates the following two steps:

1. Given a masked sequence \bar{x}_t , draw all masked tokens by sampling $x_{t-1} \sim p_\theta(\cdot | \bar{x}_t)$ to produce a *fully unmasked* sequence x_{t-1} .
2. Remask x_{t-1} using a confidence-based remasking function γ yielding a *partially masked* sequence \bar{x}_{t-1} . More specifically, with decreasing t , LLaDA remasks a decreasing number of tokens based on the current model's confidence score $p_\theta(x_{t-1}^i | \bar{x}_t)$

LLaDA continues this process until x_0 is sampled. The confidence-based remasking strategy in the second step above reveals more and more structure for $p_\theta(\cdot | \bar{x}_t)$ in the first step to condition on over time. T is a hyperparameter that trades off computation and quality (Nie et al., 2025) in practice. We refer readers to Section 2.4 for details about the implementation of the remasking function γ .

2.2 THE TRAINING-INFERENCE DIVIDE

The above formulation of LLaDA reveals an important discrepancy between training and inference: During training, the model is optimized to predict all masked tokens at randomly sampled masking ratios τ . In contrast, inference involves multiple iterative denoising steps informed by the model's confidence score, forming a trajectory with fewer and fewer masked tokens to be predicted, and more and more structure being revealed. However, this structure is ignored during training where tokens are masked at random at each iteration. This *training-inference divide* limits the performance of the model as the model is unable to learn effective denoising trajectories.

We observe symptoms of this divide by inspecting denoising trajectories of pre-trained MDLMs. Table 1 shows the accuracy of LLaDA on two reasoning tasks and the additional accuracy if we count intermediate-step correct answers that are 'refined' into wrong final answers. We term this phenomenon as over-denoising and conduct detailed studies in Section 3.3.

A simple solution to address this discrepancy would be to fine-tune the model with ground-truth trajectories. However, such trajectories are inherently unavailable, as human-generated data does not capture iterative denoising paths. In this paper, we hence investigate the following question:

Task	Acc.
MATH-500	18.2
w/ inter. ans	+9.8
Countdown	38.7
w/ inter. ans	+6.2

Table 1: Accuracy of LLaDA on two verifiable tasks and the ratio of at least one out of 256 intermediate steps is correct (w/ inter.ans).

¹We use the bar notation \bar{x} to denote (partially) masked and x for unmasked sequences throughout the paper.

How can we learn effective discrete diffusion trajectories without direct supervision?

A key advantage of masked diffusion language models over traditional auto-regressive models is that they yield *complete* text generations at *every* inference step. This allows for evaluating the quality of intermediate generations using an appropriately chosen reward model. In this paper, we propose to exploit this property to effectively fine-tune diffusion policies using deep reinforcement learning. To focus our analysis, we investigate reasoning tasks with verifiable answers, such as math problems. However, we remark that other problems without verifiable answers can also be addressed within our framework using modern validation concepts such as LLM-as-a-judge (Li et al., 2024).

2.3 MASKED DIFFUSION POLICY OPTIMIZATION (MDPO)

While reinforcement learning (RL) has originally been applied to address sequential decision-making tasks, it emerges as a promising alternative for optimizing deep neural networks to maximize non-differentiable objectives (Mnih et al., 2013; François-Lavet et al., 2018; Ouyang et al., 2022). Recognizing inference in MDLMs as a sequential decision-making problem, we propose to explicitly train the mask prediction network as an RL policy (Jaeger & Geiger, 2024). Policy gradient methods train a policy network $\pi(a | s)$ that predicts a probability distribution over actions using non-differentiable rewards by maximizing the expected return:

$$\mathcal{J}(\pi) = \mathbb{E}_{(s,a)} [\pi(a | s) r(s, a)] \quad (2)$$

where the expectation is computed over state-action pairs (s, a) collected on-policy and $r(s, a)$ is a reward function. Linking MDLMs to policy gradient methods, each denoising step corresponds to an action that predicts all masked tokens in a (partially) masked sequence \bar{x}_t , with reward $r(x_{t-1})$ via an evaluation model $r(\cdot)$.

Based on this, we derive the expected return of our Masked Diffusion Policy Optimization:

$$\mathcal{J}(\theta) \triangleq \mathbb{E}_{\{x_T, \dots, x_0\} \sim p_\theta} \left[\sum_{t=1}^T \sum_{i=1}^L \mathbb{1}[\bar{x}_t^i = \mathcal{M}] \log p_\theta(x_{t-1}^i | \bar{x}_t) r(x_{t-1}) \right] \quad \bar{x}_t \triangleq \gamma(x_t, t) \quad (3)$$

where the expectation is computed over denoising trajectories $\{x_T, \dots, x_0\}$ generated by the current denoising policy p_θ . We use the same masking function $\gamma(x_t, t)$ as used during inference, which takes an unmasked sequence x_t as input and converts it into a (partially) masked sequence \bar{x}_t conditioned on the current time step t . The smaller the timestep, the fewer tokens are masked (see Section 2.4 for details). The indicator function aggregates the return only at locations where the input is masked. The denoising policy models the probability of the denoised token x_{t-1}^i at location i given the masked input at the previous time step \bar{x}_t . Again, we drop the prompt here for notational clarity.

As the gradients must be computed from data generated by the current policy p_θ , the policy gradient estimator defined in Eq. (3) only allows a single optimization step for each round of data collection. To enable multiple steps of optimization, we use importance sampling (Kakade & Langford, 2002; Black et al., 2024) and the clipped surrogate objective from Proximal Policy Optimization (PPO) (Schulman et al., 2017). Moreover, inspired by the effectiveness of group-relative advantage estimation (Shao et al., 2024), we sample a group of G trajectories $\{x_{t,1}, \dots, x_{t,G}\}$ at each step t from the old policy $p_{\theta_{\text{old}}}$ to update the current policy p_θ , see Fig. 1 for demonstration. Our final objective is given by:

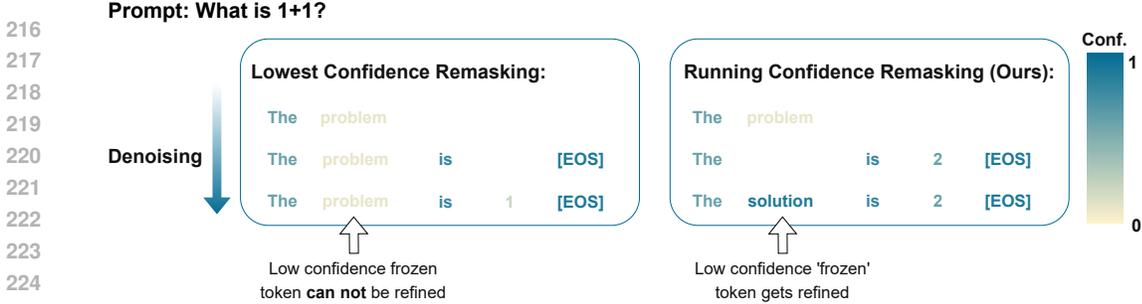
$$\mathcal{J}(\theta) \triangleq \mathbb{E}_{\{x_{T,k}, \dots, x_{0,k}\}_{k=1}^G \sim p_{\theta_{\text{old}}}} \left[\frac{1}{G} \sum_{k=1}^G \sum_{t=1}^T \sum_{i=1}^L \mathbb{1}[\bar{x}_{t,k}^i = \mathcal{M}] (Z_{t,k}^i - \beta \mathbb{D}_{KL}[p_\theta || p_{\text{ref}}]) \right] \quad (4)$$

Here, p_{ref} is the reference policy which is usually the initial model before training. $\mathbb{D}_{KL}[p_\theta || p_{\text{ref}}]$ is the KL divergence between the trained policy p_θ and the reference policy to avoid large deviation. We define the clipped and weighted advantage (Schulman et al., 2017) as

$$Z_{t,k}^i = \min \left[\frac{p_\theta(x_{t-1,k}^i | \bar{x}_{t,k})}{p_{\theta_{\text{old}}}(x_{t-1,k}^i | \bar{x}_{t,k})} A_{t-1,k}, \text{clip} \left(\frac{p_\theta(x_{t-1,k}^i | \bar{x}_{t,k})}{p_{\theta_{\text{old}}}(x_{t-1,k}^i | \bar{x}_{t,k})}, 1 - \epsilon, 1 + \epsilon \right) A_{t-1,k} \right] \quad (5)$$

where ϵ and β are hyper-parameters. $A_{t,k}$ is the advantage of $x_{t,k}$, the k 'th rollout in the group, calculated based on relative rewards inside the group:

$$A_{t,k} = \frac{s_{t,k} - \text{mean}(\{s_{t,1}, s_{t,2}, \dots, s_{t,G}\})}{\text{std}(\{s_{t,1}, s_{t,2}, \dots, s_{t,G}\})} \quad s_{t,k} = r(x_{t,k}) - r(x_{t+1,k}) + \frac{1}{t} \sum_{j=0}^{t-1} r(x_{j,k}) \quad (6)$$



226 Figure 2: Comparison between Low-Confidence Remasking (LCR) and our proposed Running
 227 Confidence Remasking (RCR) during iterative denoising. For each step, we show tokens that are *not*
 228 remasked. LCR freezes low-confidence tokens once unmasked, preventing further refinement, which
 229 potentially accumulates early-stage noise. For example, the token ‘problem’ predicted and frozen
 230 by LCR at step 1 is wrong but maintained until the end of the denoising, which leads to the final
 231 wrong answer. Whereas RCR tracks the running maximum confidence for each position, allowing
 232 persistently low-confidence tokens to be refined in later steps, leading to higher-quality completions.
 233

234 The original group-relative advantage estimation optimizes the memory usage of training LLMs
 235 with RL by dropping the jointly trained value model that is used to predict the advantages. Beyond
 236 this, another intuitive reason to use group-relative advantage estimation in the diffusion setting is to
 237 incentivize better generation with fewer denoising steps. Specifically, by comparing to a group of
 238 trajectories, steps reaching the final reward more rapidly are assigned higher advantages. In addition,
 239 Eq. (6) encourages each denoising step to (i) make an immediate improvement over the previous,
 240 noisier prediction and (ii) steer the trajectory toward cleaner states that score well overall.
 241

242 **2.4 REMASKING**

243 As introduced in Section 2.1, inference in LLaDA alternates between predicting all masked tokens
 244 in \bar{x}_t and selectively remasking a fraction of the prediction x_{t-1} via $\bar{x}_{t-1} = \gamma(x_{t-1}, t)$ for iterative
 245 refinement. We first discuss random and confidence-based remasking as introduced in LLaDA. Next,
 246 we propose a simple yet effective improvement to address a key limitation of confidence-based
 247 remasking, which consistently leads to better performance for both LLaDA pre-trained as well as
 248 MDPO fine-tuned models. Formally, the remasking function γ assigns the mask symbol \mathcal{M} to the n_t
 249 tokens with the highest masking scores:

250
$$\gamma(x_t, t) = \begin{cases} \mathcal{M} & \text{if } m_t^i \geq (n_t\text{'th highest value of } m_t) \\ x_t^i & \text{otherwise} \end{cases} \quad (7)$$

253 where $n_t = \lfloor \phi(t, T)L \rfloor$ determines the number of tokens to be masked depending on the time step t
 254 and a scheduling function ϕ . n_t decreases as t increases (more tokens are masked for larger t). m_t^i
 255 is a generic definition of the masking score of token i at step t and remasking strategies vary in how
 256 they compute m_t^i . LLaDA applies a linear function for ϕ so that an equal amount of $\lfloor \frac{L}{T} \rfloor$ tokens are
 257 *not* remasked every step, i.e., $n_t = \lfloor \frac{T-t}{T}L \rfloor$. We investigate various schedules in Appendix B.2.

258 **Random Remasking (RR):** A simple baseline to construct masking scores is random sampling:

259
$$m_{t-1}^i \sim \mathcal{U}(0, 1) \quad \forall_{t,i} : \bar{x}_t^i = \mathcal{M} \quad (8)$$

261 While random sampling is consistent with the LLaDA training objective (Section 2.1), which masks
 262 tokens randomly, this strategy performs poorly in practice (Nie et al., 2025) as it doesn’t exploit the
 263 structure of natural diffusion trajectories.

264 **Low-Confidence Remasking (LCR):** Inspired by the annealing tricks of sampling in AR LLMs
 265 (Holtzman et al., 2020), LLaDA proposes to utilize the complement of the predicted confidence of
 266 the model as the masking scores:

267
$$m_{t-1}^i = 1 - p_\theta(x_{t-1}^i | \bar{x}_t) \quad \forall_{t,i} : \bar{x}_t^i = \mathcal{M} \quad (9)$$

268 Empirical evaluations demonstrate that this model-derived low-confidence remasking significantly
 269 outperforms random masking in downstream tasks.

Running Confidence Remasking (RCR): An important observation is that both remasking strategies above (used in LLaDA) assign masking scores only to predicted tokens in x_t that were masked before. The unmasked tokens remain fixed until the end of the denoising process. In other words, if x_t^i is not remasked by $\gamma(x_t, t)$, it will be unmasked (frozen) in the following steps $x_{<t}^i$. We consider this a crucial limitation, as the predicted tokens, particularly in the early steps, tend to be highly noisy due to the limited structure revealed at that stage. Freezing these noisy tokens until the end of the denoising process makes it more difficult to produce high-quality generation in practice.

To address this, we propose Running Confidence Remasking (RCR). Instead of deciding based solely on the confidence at the current step, we track for each position i the highest confidence it has achieved so far for predicting masked tokens during the denoising process. At each step, we identify the n_t positions whose running maximum confidence is the lowest and remask tokens at these positions. Formally, for token i at step t , let t' index any earlier step in the denoising process ($t \leq t' \leq T$), the masking score is defined as:

$$m_{t-1}^i = 1 - \max_{t' \geq t} (p_\theta(x_{t'-1}^i | \bar{x}_{t'})) \quad \forall_{t,i} : \bar{x}_t^i = \mathcal{M} \quad (10)$$

Tokens at positions with higher running maximum confidence get lower masking scores (less likely to be remasked), while those that have never reached high confidence remain candidates for remasking.

Empirically, with more structure being revealed along with denoising, earlier steps often yield low-confidence predictions, whereas later steps tend to converge to higher confidence. Under the LCR strategy, early tokens are retained despite their relatively low confidence if they happen to fall within the top- n_t set at that step and can not be refined in future steps. In contrast, as Fig. 2 shows, RCR allows such tokens to be remasked in later steps if tokens at other positions surpass them in running confidence, enabling the model to revise uncertain predictions before producing the final output.

3 EXPERIMENTS

Our experiments are designed to address three core questions:

1. How do MDPO and RCR improve generation performance and sampling efficiency, and are their effects complementary?
2. Recognizing over-denoising as a manifestation of the training–inference divide, how can we leverage it to close this gap?
3. What is the impact of different sampling settings during MDPO training?

3.1 EXPERIMENTAL SETUP

Inference: The performance of MDLMs depends on several inference-time hyperparameter choices. A key distinction lies in whether denoising is performed over the full sequence at once or in blocks. Previous works apply a semi-autoregressive strategy where the sequence is divided into several blocks and generated from left to right. Within each block, tokens are denoised in parallel. We refer to this setting as *semi-AR*, and contrast it with the setting of denoising all tokens in the entire sequence simultaneously (*pure-Diff*). In semi-AR, the same remasking strategies used in pure-Diff apply, but each block denoises fewer tokens with fewer steps. To generate a sequence of length L with block size B in T steps, each block is allocated approximately $\lfloor \frac{T \cdot B}{L} \rfloor$ steps. As shown in Fig. 6, performance varies significantly depending on the choice of B and T . For controlled comparisons, we adopt two standard configurations throughout our experiments: pure-Diff ($B = L$) and semi-AR ($B = 128$). Extensive analysis on choices of block size is in Appendix B.3. Though empirically the pure-Diff setting usually underperforms semi-AR, we keep it as an interesting setting in our main evaluation because of the potential speed advantage that pure-Diff possesses at generating long responses.

Data: We investigate reasoning tasks with answers that can be verified by static reward functions: (1) MATHEMATICAL REASONING where the model generates solutions to mathematical problems and we verify if the ground truth answers are in the solutions with a robust mathematical expression evaluation system (Kydliček & Gandenberger, 2025). Note that different math verifiers can lead to changes in evaluation results and we control the experiments by consistently use the same verifier. (2) PLANNING with Countdown (Pan et al., 2025), an easy-to-verify combinatorial arithmetic game in

Table 2: **Model performance on Mathematics and Countdown: Best** and second-best methods in each setting are shaded. For each task we report results on semi-AR (Block Size = 128) and pure-Diff (Block Size = 512) given the generation length of 512. We also compare the performance across multiple choices of denoising steps.

Block Size	MATH500						Countdown					
	128			512			128			512		
	64	128	256	64	128	256	64	128	256	64	128	256
LLaDA-8B-Instruct	20.0	36.2	39.4	20.0	22.6	18.2	5.5	18.8	20.7	14.8	33.6	38.7
+ SFT (Nie et al., 2025)	18.0	33.8	41.8	22.2	24.6	25.4	12.3	20.7	28.6	35.5	43.1	49.6
+ diffu-GRPO (Zhao et al., 2025a)	23.8	34.4	40.2	21.2	26.2	20.6	18.2	35.9	40.1	49.4	55.0	58.9
+ RCR (training-free)	25.4	37.6	40.8	22.6	22.6	19.2	9.4	15.2	26.6	32.8	43.8	45.7
+ MDPO	26.2	38.6	42.8	23.4	25.2	26.2	64.8	70.7	70.7	68.4	68.8	71.1
+ MDPO + RCR	30.4	38.4	44.2	24.8	24.0	26.4	69.9	70.7	73.4	67.2	69.5	57.4

which the model tries to reach target numbers using basic arithmetic operations on a given set (3 or 4) of numbers. Evaluation is on MATH-500 (Lightman et al., 2024), a curated dataset of 500 problems from the MATH dataset (Hendrycks et al., 2021); and the Countdown test set by Zhao et al. (2025a), specifically. For training on the mathematical reasoning task, we utilize (part of) the math data from Huggingface OpenR1 (Hugging Face, 2025) that consists of 93.7k competition-level math problems from NuminaMath 1.5 (LI et al., 2024). For Countdown, we directly use the respective training split released by Pan et al. (2025).

Model and Training: We compare our methods to two baseline methods: diffu-GRPO (Zhao et al., 2025a), which is the first integration of policy gradient methods trained on MDLMs, and LLaDA with supervised finetuning (SFT). We (re)train both the baseline methods and MDPO initialized with LLaDA-8B-Instruct (Nie et al., 2025) under the same settings. Training of all methods is conducted under a fixed compute budget using 8 NVIDIA H100 GPUs with 100 weight update steps. We use a batch size of 128 and apply gradient accumulation when memory constraints occur. For both diffu-GRPO and our proposed MDPO, the group size for group-relative advantage estimation is 8, and the number of denoising steps for sampling rollouts is 128. One special setting that will be discussed in detail in Section 3.3 is that we train MDPO with *only* over-denoising samples, which shows better sample efficiency.

3.2 IMPROVEMENTS IN GENERATION PERFORMANCE AND SAMPLE EFFICIENCY

We first show the results (see Table 2) of our proposed Masked Diffusion Policy Optimization (MDPO) and training-free Running Confidence Remasking (RCR) compared to the baselines on MATH-500 and Countdown. Across all configurations, both MDPO and RCR individually improve substantially upon the LLaDA initialization, with RCR often achieving performance comparable to MDPO despite requiring no additional training. We remark that RCR, as a training-free method, even outperforms the training baselines in most of the settings for the MATH task. Notably, combining MDPO with RCR consistently yields further gains over either method alone, achieving the best or second-best performance in nearly all settings, which demonstrates that MDPO and RCR are complementary. We further observe that the relative performance gains are more pronounced in settings with fewer inference steps, indicating improved *sampling efficiency*.

Findings from Countdown: Furthermore, in Countdown, we observe a surprising behavioral shift: the model trained with MDPO transitions from generating explicit step-by-step reasoning to producing direct answers in a few (<10) steps. This echoes the advantages of diffusion models for latent reasoning discussed by previous work (Zhu et al., 2025). Another interesting observation is that pure-Diff (of LLaDA) substantially outperforms semi-AR on the Countdown task, which is rare on our math task and other tasks evaluated by the LLaDA paper. We hypothesize that this is due to the short expected reasoning traces for Countdown, where semi-AR uses fewer *effective* blocks as well as fewer *effective* denoising steps, limiting its ability to refine predictions. After MDPO training, as the model tends to directly produce answers without explicit reasoning, the performance

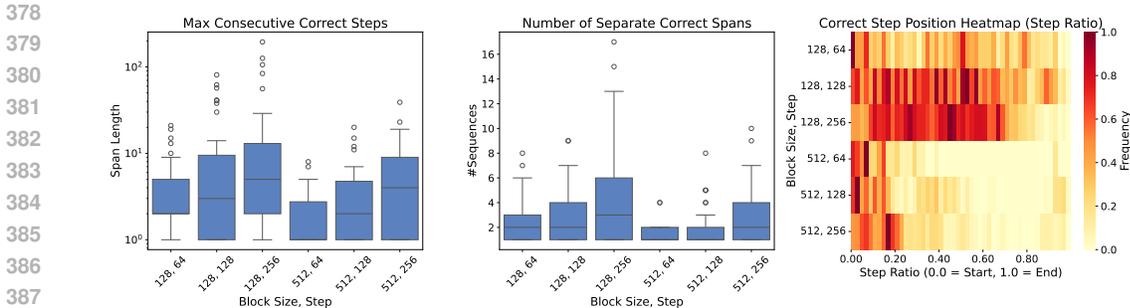


Figure 3: **Analysis on the occurrence of over-denoising samples across multiple inference settings (Block size, Step).** (Left) Maximum length of consecutive correct steps (span) in the denoising trajectories. (Middle) Number of separate correct spans over trajectories. (Right) Heatmap of the relative position of correct steps across the denoising process.

between pure-Diff and semi-AR gets closer. This observation highlights the flexibility of pure-Diff, particularly for tasks that do not require extended reasoning chains.

3.3 OVER-DENOISING AS AN EFFECTIVE DATA FILTER

We analyse the over-denoising phenomenon in more detail in Fig. 3 which reveals two key insights. First, the left and middle subfigures show that correct spans are typically short and fragmented, with trajectories often containing multiple separate correct spans instead of steadily accumulating correct steps. Such fragmentation increases the likelihood of losing correct tokens before reaching the final step. Second, the heatmap shows that correct answers often appear surprisingly very early but tend to decay over subsequent steps instead of steadily accumulating.

These findings not only underline the necessity to address the training-inference divide for MDLMs, but also reveal that over-denoising provides highly informative signals of intermediate steps for improving MDLMs. Inspired by this, we train MDPO with only over-denoising samples instead of all data so that MDPO can learn from these intermediate signals more efficiently. Specifically, we first use the initialized model to run inference on the whole dataset to identify over-denoising samples and then train only on this subset, which constitutes roughly 10% of the original data. The comparison between MDPO-all-data and MDPO in Fig. 4 shows that the model trained on only over-denoising excels in most settings, highlighting the effectiveness of over-denoising as a data filter for MDPO.

3.4 IMPACT OF SAMPLING SETTINGS ON MDPO

The heatmap in Fig. 3 reveals that semi-AR tends to distribute correctness more evenly but with greater fragmentation, while pure-Diff produces an early burst of correctness that decays sharply and rarely recovers. This raises a question of how different sampling settings during the rollout collection of RL affect the final performance. To investigate this, we compare MDPO variants trained on (i) pure-Diff rollouts only, (ii) semi-AR rollouts only, and (iii) an even mixture of both. Fig. 4 shows that training on a single mode yields the largest improvement in the evaluation setting of the respective mode, but often at the cost of performance in the other modes. The mixture strategy used in our main MDPO setting achieves a more balanced performance, matching or exceeding the best single-mode results in several configurations. This suggests that mixed sampling allows the policy to learn denoising behaviors that generalize across all inference strategies.

4 RELATED WORK

Diffusion probabilistic models (Sohl-Dickstein et al., 2015; Ho et al., 2020) have become the de facto standard for generative modeling of continuous signals, including images, videos, 3D shapes, and robotic trajectories. Building on early efforts in discrete diffusion (Austin et al., 2021; Hoogeboom et al., 2021), several works have adapted diffusion models for text generation. Diffusion-LM (Li et al., 2022) and Self-Conditioned Embedding Diffusion (Strudel et al., 2022) have pioneered embedding-based diffusion approaches. Liu et al. (2025) introduces a novel framework that separates the

432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

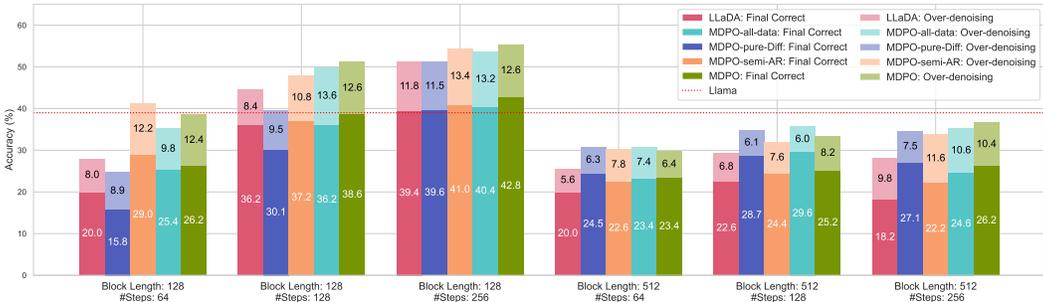


Figure 4: **Comparison of MDPO variants.** We report the final accuracy and proportion of over-denoising cases for all models. MDPO-all-data is trained on all data samples, whereas MDPO is trained on *only* over-denoising data, both with rollouts sampled from a mixture of semi-AR and pure-Diff. MDPO-pure-Diff and MDPO-semi-AR represent models that are trained on rollouts sampled from *only* pure-Diff, and semi-AR, respectively.

generation process into two models: a planner and a denoiser. Meanwhile, masked diffusion (Austin et al., 2021; Sahoo et al., 2024; Ou et al., 2025) has been established as a special case of discrete diffusion, with recent research significantly scaling their capabilities. Notably, DiffuLLaMA (Gong et al., 2025) initializes masked diffusion language models using pre-trained LLaMA weights. LLaDA (Nie et al., 2025) and Dream (Ye et al., 2025) scale MDLMs by pre-training from scratch to the billion-parameter regime, achieving performance comparable to similarly sized AR LLMs.

Training diffusion models with RL has been primarily explored in image diffusion models (Black et al., 2024; Fan et al., 2023) on non-differentiable objectives such as human-perceived image quality. Represented by Group Relative Policy Optimization (GRPO) (Shao et al., 2024), RL has emerged as a promising tool to optimize sparse and task-specific reward signals for language models (Guo et al., 2025; Yu et al., 2025b; Luo et al., 2025b;a; Hu et al., 2025; Zeng et al., 2025), especially for post-training. These RL algorithms assume sequential token-by-token generation, a key property of AR models. However, MDLMs differ fundamentally: they denoise entire sequences from corrupted inputs, making AR-specific RL inapplicable. Recent attempts Zhao et al. (2025a;b) adapt GRPO to MDLMs with inserted partial ground-truth reasoning traces to improve sample efficiency, meanwhile concurrent work Wang et al. (2025b) also observes the over-denoising phenomenon and adds entropy-based rewards to optimize MDLMs with RL. In contrast, we introduce MDPO, a policy optimization method designed for MDLMs to learn effective denoising trajectories by optimizing the iterative denoising process as a multi-step decision-making problem. To our knowledge, MDPO is the first RL method explicitly designed to optimize MDLMs denoising trajectories as multi-step decision-making.

Kim et al. (2025); Wang et al. (2025a) show that though MDLMs face inherently harder training challenges due to solving computationally intractable infilling tasks, using adaptive decoding strategies effectively alleviates these. Following this, Block Diffusion (Arriola et al., 2025) models sequences in autoregressive blocks with intra-block diffusion, concurrent work Li et al. (2025) observes the same over-denoising phenomenon and applies a confidence-based early-commit strategy for faster denoising. Our proposed Running Confidence Remasking enables more flexible decoding to substantially enhance the denoising quality by enabling iterative revision of early-predicted ‘noisy’ tokens.

5 DISCUSSION & CONCLUSION

This paper tackles the training–inference divide in Masked Diffusion Language Models (MDLMs) by proposing Masked Diffusion Policy Optimization (MDPO) and Running Confidence Remasking (RCR). MDPO uses reinforcement learning to optimize denoising trajectories with intermediate rewards, while RCR flexibly revisits earlier noisy predictions to reduce error propagation. One limitation of our work is that only verifiable tasks with static reward functions are investigated. We leave the application of MDPO to general tasks using modern validation concepts such as LLM-as-a-judge for intermediate rewards as future work.

REFERENCES

- 486
487
488 Marianne Arriola, Subham Sekhar Sahoo, Aaron Gokaslan, Zhihan Yang, Zhixuan Qi, Jiaqi Han,
489 Justin T Chiu, and Volodymyr Kuleshov. Block diffusion: Interpolating between autoregressive
490 and diffusion language models. In *Proc. of the International Conf. on Learning Representations*
491 *(ICLR)*, 2025. 9
- 492 Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured
493 denoising diffusion models in discrete state-spaces. In *Advances in Neural Information Processing*
494 *Systems (NeurIPS)*, 2021. 2, 8, 9
- 495 Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models
496 with reinforcement learning. In *Proc. of the International Conf. on Learning Representations*
497 *(ICLR)*, 2024. 4, 9
- 498 Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T. Freeman. Maskgit: Masked generative
499 image transformer. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp.
500 11305–11315, 2022. 16
- 502 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of
503 deep bidirectional transformers for language understanding. In *North American Chapter of the*
504 *Association for Computational Linguistics (NAACL)*, 2019. 1
- 505 Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel,
506 Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. Reinforcement learning for fine-
507 tuning text-to-image diffusion models. In *Advances in Neural Information Processing Systems*
508 *(NeurIPS)*, 2023. 9
- 509 Guhao Feng, Yihan Geng, Jian Guan, Wei Wu, Liwei Wang, and Di He. Theoretical benefit and
510 limitation of diffusion language model. *arXiv.org*, 2025. 3
- 511 Vincent François-Lavet, Peter Henderson, Riashat Islam, Marc G. Bellemare, and Joelle Pineau. An
512 introduction to deep reinforcement learning. *Found. Trends Mach. Learn.*, 11(3-4):219–354, 2018.
513 4
- 514 Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An,
515 Peilin Zhao, Wei Bi, Jiawei Han, Hao Peng, and Lingpeng Kong. Scaling diffusion language
516 models via adaptation from autoregressive models. In *Proc. of the International Conf. on Learning*
517 *Representations (ICLR)*, 2025. 9
- 518 Google DeepMind. Gemini diffusion: our state-of-the-art experimental text diffusion model, 2025.
519 URL <https://deepmind.google/models/gemini-diffusion/>. 1
- 520 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
521 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
522 via reinforcement learning. *arXiv.org*, 2025. 9
- 523 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
524 and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In
525 *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 7
- 526 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in*
527 *Neural Information Processing Systems (NeurIPS)*, 2020. 8
- 528 Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, pp.
529 1735–1780, 1997. 1
- 530 Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text
531 degeneration. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2020. 5
- 532 Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows
533 and multinomial diffusion: Learning categorical distributions. In *Advances in Neural Information*
534 *Processing Systems (NeurIPS)*, 2021. 8

- 540 Jingcheng Hu, Yinmin Zhang, Qi Han, Daxin Jiang, Xiangyu Zhang, and Heung-Yeung Shum.
541 Open-reasoner-zero: An open source approach to scaling up reinforcement learning on the base
542 model. *arXiv.org*, 2025. 9
- 543
544 Hugging Face. Open r1: A fully open reproduction of deepseek-r1, January 2025. URL <https://github.com/huggingface/open-r1>. 7
- 545
546 Bernhard Jaeger and Andreas Geiger. An invitation to deep reinforcement learning. *Found. Trends*
547 *Optim.*, 7(1):1–80, 2024. 4
- 548
549 Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In
550 *Proc. of the International Conf. on Machine learning (ICML)*, pp. 267–274, 2002. 4
- 551
552 Samar Khanna, Siddhant Kharbanda, Shufan Li, Harshit Varma, Eric Wang, Sawyer Birnbaum,
553 Ziyang Luo, Yanis Miraoui, Akash Palrecha, Stefano Ermon, Aditya Grover, and Volodymyr
554 Kuleshov. Mercury: Ultra-fast language models based on diffusion. *arXiv.org*, 2025. 1
- 555
556 Jaeyeon Kim, Kulin Shah, Vasilis Kontonis, Sham M. Kakade, and Sitan Chen. Train for the worst,
557 plan for the best: Understanding token ordering in masked diffusions. *arXiv.org*, 2025. 9
- 558
559 Hynek Kydlíček and Greg Ganderberger. Math-verify: A robust mathematical expression evaluation
560 system. <https://github.com/huggingface/Math-Verify>, 2025. Version 0.8.0, released July 2, 2025. 6
- 561
562 Haitao Li, Qian Dong, Junjie Chen, Huixue Su, Yujia Zhou, Qingyao Ai, Ziyi Ye, and Yiqun
563 Liu. Llm-as-judges: A comprehensive survey on llm-based evaluation methods. *arXiv.org*,
564 abs/2412.05579, 2024. 4
- 565
566 Jia LI, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Costa
567 Huang, Kashif Rasul, Longhui Yu, Albert Jiang, Ziju Shen, Zihan Qin, Bin Dong,
568 Li Zhou, Yann Fleureau, Guillaume Lample, and Stanislas Polu. NuminaMath.
569 [<https://huggingface.co/AI-MO/NuminaMath-1.5>](https://github.com/project-numina/aimo-progress-prize/blob/main/report/numina_dataset.pdf), 2024. 7
- 570
571 Pengxiang Li, Yefan Zhou, Dilxat Muhtar, Lu Yin, Shilin Yan, Li Shen, Yi Liang, Soroush Vosoughi,
572 and Shiwei Liu. Diffusion language models know the answer before decoding. *arXiv.org*,
573 abs/2508.19982, 2025. 9
- 574
575 Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B. Hashimoto. Diffusion-
576 lm improves controllable text generation. In *Advances in Neural Information Processing Systems*
577 (*NeurIPS*), 2022. 8
- 578
579 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
580 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *Proc. of the*
581 *International Conf. on Learning Representations (ICLR)*, 2024. 7
- 582
583 Sulin Liu, Juno Nam, Andrew Campbell, Hannes Stärk, Yilun Xu, Tommi S. Jaakkola, and Rafael
584 Gómez-Bombarelli. Think while you generate: Discrete diffusion with planned denoising. In *Proc.*
585 *of the International Conf. on Learning Representations (ICLR)*, 2025. 8
- 586
587 Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios
588 of the data distribution. In *Proc. of the International Conf. on Machine learning (ICML)*, 2024. 1
- 589
590 Michael Luo, Roy Huang Sijun Tan, Xiaoxiang Shi, Rachel Xin, Colin Cai, Ameen Patel, Alpay
591 Ariyak, Qingyang Wu, Ce Zhang, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepcoder: A
592 fully open-source 14b coder at o3-mini level, 2025a. Notion Blog. 9
- 593
594 Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai,
595 Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepscaler: Surpassing o1-preview
596 with a 1.5b model by scaling rl, 2025b. Notion Blog. 9
- 597
598 Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan
599 Wierstra, and Martin A. Riedmiller. Playing atari with deep reinforcement learning. *arXiv.org*,
600 2013. 4

- 594 Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin,
595 Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv.org*, 2502.09992, 2025.
596 1, 2, 3, 5, 7, 9, 15
597
- 598 Jingyang Ou, Shen Nie, Kaiwen Xue, Fengqi Zhu, Jiacheng Sun, Zhenguo Li, and Chongxuan Li.
599 Your absorbing discrete diffusion secretly models the conditional distributions of clean data. In
600 *Proc. of the International Conf. on Learning Representations (ICLR)*, 2025. 1, 2, 9
601
- 602 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
603 Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton,
604 Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and
605 Ryan Lowe. Training language models to follow instructions with human feedback. In *Advances
606 in Neural Information Processing Systems (NeurIPS)*, 2022. 4
607
- 608 Jiayi Pan, Junjie Zhang, Xingyao Wang, Lifan Yuan, Hao Peng, and Alane Suhr. Tinyzero.
609 <https://github.com/Jiayi-Pan/TinyZero>, 2025. Accessed: 2025-01-24. 6, 7
610
- 611 Subham S. Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T. Chiu,
612 Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language
613 models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 1, 9
614
- 615 John Schulman. Approximating kl divergence, 2020. URL <http://joschu.net/blog/kl-approx.html>. 14
616
- 617 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
618 optimization algorithms. *arXiv.org*, 2017. 4
619
- 620 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
621 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical
622 reasoning in open language models. *arXiv.org*, 2024. 4, 9
623
- 624 Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
625 learning using nonequilibrium thermodynamics. In *Proc. of the International Conf. on Machine
626 learning (ICML)*, 2015. 8
627
- 628 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Proc. of
629 the International Conf. on Learning Representations (ICLR)*, 2021. 3
630
- 631 Yuxuan Song, Zheng Zhang, Cheng Luo, Pengyang Gao, Fan Xia, Hao Luo, Zheng Li, Yuehang
632 Yang, Hongli Yu, Xingwei Qu, Yuwei Fu, Jing Su, Ge Zhang, Wenhao Huang, Mingxuan Wang,
633 Lin Yan, Xiaoying Jia, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Yonghui Wu, and Hao Zhou.
634 Seed diffusion: A large-scale diffusion language model with high-speed inference. *arXiv.org*, 2025.
635 1
636
- 637 Robin Strudel, Corentin Tallec, Florent Althé, Yilun Du, Yaroslav Ganin, Arthur Mensch, Will
638 Grathwohl, Nikolay Savinov, Sander Dieleman, Laurent Sifre, and Rémi Leblond. Self-conditioned
639 embedding diffusion for text generation. *arXiv.org*, abs/2211.04236, 2022. 8
640
- 641 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,
642 Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information
643 Processing Systems (NeurIPS)*, 2017. 1
644
- 645 Guanghan Wang, Yair Schiff, Subham Sekhar Sahoo, and Volodymyr Kuleshov. Remasking discrete
646 diffusion models with inference-time scaling. *arXiv.org*, abs/2503.00307, 2025a. 9
647
- 648 Wen Wang, Bozhen Fang, Chenchen Jing, Yongliang Shen, Yangyi Shen, Qiuyu Wang, Hao Ouyang,
649 Hao Chen, and Chunhua Shen. Time is a feature: Exploiting temporal dynamics in diffusion
650 language models. *arXiv.org*, abs/2508.09138, 2025b. 9
651
- 652 Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with
653 denoising diffusion gans. In *Proc. of the International Conf. on Learning Representations (ICLR)*,
654 2022. 3
655

- 648 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
649 Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
650 Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang,
651 Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia,
652 Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu
653 Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *arXiv.org*, abs/2412.15115, 2024.
654 17
- 655 Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng
656 Kong. Dream 7b: Diffusion large language models. *arXiv.org*, abs/2508.15487, 2025. 1, 9, 17
657
- 658 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong
659 Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi
660 Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiase Chen, Jiangjie Chen, Chengyi
661 Wang, Hongli Yu, Weinan Dai, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying
662 Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. DAPO: an open-source
663 LLM reinforcement learning system at scale. *arXiv.org*, abs/2503.14476, 2025a. 18
- 664 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong
665 Liu, Lingjun Liu, Xin Liu, et al. Dapo: An open-source llm reinforcement learning system at scale.
666 *arXiv.org*, 2025b. 9
- 667 Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerl-
668 zoo: Investigating and taming zero reinforcement learning for open base models in the wild.
669 *arXiv.org*, 2025. 9
670
- 671 Siyan Zhao, Devaansh Gupta, Qinqing Zheng, and Aditya Grover. d1: Scaling reasoning in diffusion
672 large language models via reinforcement learning. *arXiv.org*, 2504.12216, 2025a. 7, 9
- 673 Siyan Zhao, Mengchen Liu, Jing Huang, Miao Liu, Chenyu Wang, Bo Liu, Yuandong Tian, Guan
674 Pang, Sean Bell, Aditya Grover, and Feiyu Chen. Inpainting-guided policy optimization for
675 diffusion large language models. *arXiv.org*, abs/2509.10396, 2025b. 9, 18
676
- 677 Rui-Jie Zhu, Tianhao Peng, Tianhao Cheng, Xingwei Qu, Jinfa Huang, Dawei Zhu, Hao Wang,
678 Kaiwen Xue, Xuanliang Zhang, Yong Shan, Tianle Cai, Taylor Kergan, Assel Kembay, Andrew
679 Smith, Chenghua Lin, Binh Nguyen, Yuqi Pan, Yuhong Chou, Zefan Cai, Zhenhe Wu, Yongchi
680 Zhao, Tianyu Liu, Jian Yang, Wangchunshu Zhou, Chujie Zheng, Chongxuan Li, Yuyin Zhou,
681 Zhoujun Li, Zhaoxiang Zhang, Jiaheng Liu, Ge Zhang, Wenhao Huang, and Jason Eshraghian. A
682 survey on latent reasoning, 2025. 7
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

Appendix

Table of Contents

A	Experimental Settings	14
A.1	Advantage Estimation	14
A.2	KL Divergence Approximation	14
A.3	Hyperparameter Setting & Other Technical Details	15
B	Model Analysis	15
B.1	Training-Inference Discrepancy	15
B.2	Remasking Scheduling Function	16
B.3	Block Size in Semi-AR	16
B.4	LLaDA and Dream	17
C	Additional Discussion on Experimental Results	18
C.1	Why over-denoising is an effective data filter?	18
D	The Use of Large Language Models	18

A EXPERIMENTAL SETTINGS

A.1 ADVANTAGE ESTIMATION

In Eq. (6), we define the score of each generation as

$$s_{t,k} = r(x_{t,k}) - r(x_{t+1,k}) + \frac{1}{t} \sum_{j=0}^{t-1} r(x_{j,k}) \quad (11)$$

to compute the advantage. The intuition behind this is the first part of the definition $r(x_{t,k}) - r(x_{t+1,k})$ incentivizes immediate improvement over the previous, noisier prediction and the second part $\frac{1}{t} \sum_{j=0}^{t-1} r(x_{j,k})$ encourages predictions that steer the trajectory toward cleaner states that score well overall. We highlight that the design of advantage estimation plays a pivotal role in determining the final performance, since poorly structured estimators can be easily hacked, allowing the model to overfit to superficial reward patterns rather than acquiring the intended behavior.

A.2 KL DIVERGENCE APPROXIMATION

We approximate the KL divergence between the trained model and the reference model $\mathbb{D}_{KL}[p_\theta || p_{\text{ref}}]$ in Eq. (4) following Schulman (2020). Define $d = \frac{p_{\text{ref}}(x)}{p_\theta(x)}$, three estimators are defined in Schulman (2020):

$$k_1 = -\log d \quad k_2 = \frac{(\log d)^2}{2} \quad k_3 = (d - 1) - \log d \quad (12)$$

where k_1 is an unbiased but high-variance estimator, as it is negative for half of the samples, while the KL divergence should always be positive. Using such an estimator inside Eq. (4) where the KL divergence is deducted from the clipped advantage may artificially reward models for ‘negative divergences’, producing biased optimization dynamics. Therefore, we seek an estimator that is always non-negative, for robustness and theoretical consistency with the definition of KL. k_2 and k_3 are both non-negative, and k_2 is more biased than k_3 , according to Schulman (2020). In our preliminary experiments, we found small differences in the performance between using k_2 and k_3 . Therefore, we choose k_2 as the final KL approximation in our current implementation for the simplicity of computation.

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

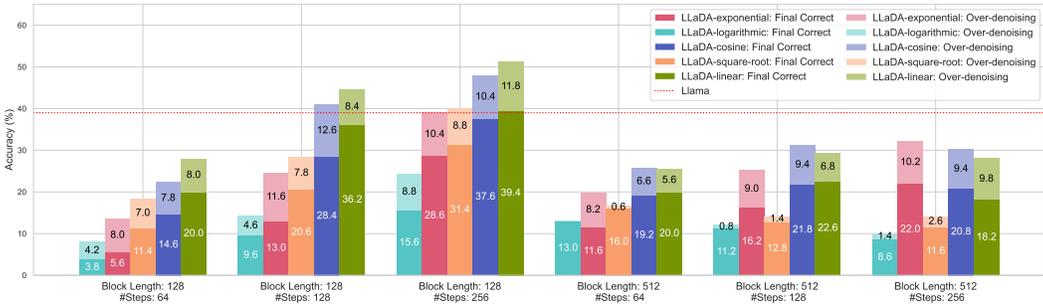


Figure 5: Comparison of different re-masking strategies introduced by MaskGIT. The performance of LLaDA with changing re-masking strategies is not consistent with MaskGIT, a image generation model.

A.3 HYPERPARAMETER SETTING & OTHER TECHNICAL DETAILS

Here we list other technical details that are necessary to reproduce our experimental results. We set the clipping parameter $\epsilon = 0.2$ in the objective in Eq. (5) and the KL divergence scaling factor $\beta = 0.08$ in Eq. (4). In our experiments, we observe that a larger β prevents the model from overfitting to the training data, given that we are supervising the intermediate trajectories and the model receives multiple supervisions for each data sample. Another implementation choice we make to avoid overfitting is to reduce the number of intermediate denoising steps used to create the gradients. Despite the expectation is over all the denoising steps as defined in Eq. (4), we sort the denoising steps according to their absolute advantage values and pick the largest c steps. This not only prevents the model from overfitting by reducing the number of the same data samples are trained on, but also reduces the computation time. c is usually set to 8 in our experiments. The intuition behind sorting advantages with absolute values is to supervise the model with both good and bad denoising. Learning rate is set to $7e-7$.

In Section 3.3, we find that using only over-denoising samples is more data-efficient and produces better results. Although this leads to overhead in pre-computing the over-denoising samples before training, it is still computationally efficient in practice. The reason is that for each inference setting, we only need to infer over the whole training set once and use the over-denoising data samples repetitively for sampling rollouts in the online RL in one training or over multiple training sessions. As the computation cost of inference is much less than completing forward and backward propagation in training, the overhead created by inference is negligible compared to the computation cost for training on all data.

B MODEL ANALYSIS

B.1 TRAINING-INFERENCE DISCREPANCY

In Section 2.4, we introduce the two remasking strategies from LLaDA, random remasking and low-confidence remasking (LCR). As the tokens are also randomly masked in pre-training, and LCR is usually used as the de-facto standard for inference (Nie et al., 2025), analyzing the difference between random remasking and low-confidence remasking provides us a closer look at the training-inference divide. Table 3 shows concrete examples to demonstrate the difference between random masking and LCR.

We observe that LCR consistently produces more coherent and complete masked intermediate completions compared to random remasking at the same masking ratio. At high masking ratios (e.g., 0.75 at Step 16), random remasking often fails to preserve the problem structure, leading to fragmented or nonsensical outputs, whereas LCR retains recognizable mathematical expressions and approaches the ground truth earlier. As the masking ratio decreases (Steps 32 and 48), the gap narrows, while LCR still maintains a higher degree of structural integrity.

810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

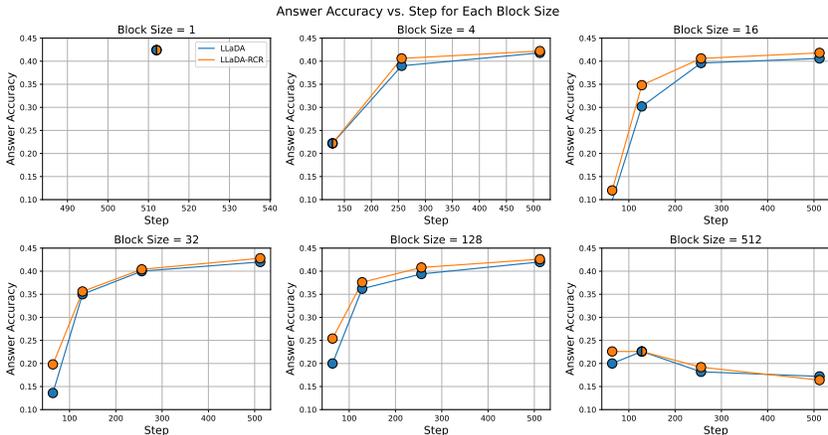


Figure 6: Interpolation of block sizes and number of time steps to investigate the effect of these two factors on the inference performance. Generation length is fixed to 512.

B.2 REMASKING SCHEDULING FUNCTION

The remasking scheduling function ϕ introduced in Section 2.4 defines how many tokens should be masked at every timestep, which can also influence the generation performance. Though the choices of ϕ is not discussed in the LLaDA paper, multiple design choices of the ϕ are studied in MaskGIT (Chang et al., 2022) for image generation tasks. We list the definitions of these functions as following

$$\phi(t, T) = \begin{cases} \frac{e - e^{\frac{t}{T}}}{e - 1}, & \text{Exponential} \\ 1 - \left(\frac{t}{T}\right)^3, & \text{Cubic} \\ 1 - \left(\frac{t}{T}\right)^2, & \text{Square} \\ \frac{1 + \cos\left(\pi \frac{t}{T}\right)}{2}, & \text{Cosine} \\ 1 - \frac{t}{T}, & \text{Linear} \\ 1 - \sqrt{\frac{t}{T}}, & \text{Square Root} \\ 1 - \frac{\ln\left(1 + \frac{t}{T}\right)}{\ln(2)}, & \text{Logarithmic} \end{cases} \quad (13)$$

Out of these functions, LLaDA uses the linear strategy that the same amount of tokens are *not* remasked at each step. We conduct evaluation with the representative ones: exponential (exp), logarithmic (log), cosine, square root, and linear, on the MATH-500 task. Results in Fig. 5 show that the linear strategy works the best, which is not consistent with MaskGIT on image synthesis, where the concave functions like cosine that follows a less-to-more generation procedure works better.

B.3 BLOCK SIZE IN SEMI-AR

As discussed in Section 3.1, a key factor that affects the performance of MDLMs is the block size. When block size $B = L$ where L is the generation length, the model denoises all the tokens simultaneously which we refer to as *pure-Diff*. And we refer to $B < L$ as *semi-AR* where the sequence is divided into several blocks and generated from left to right. Within each block, tokens are denoised in parallel.

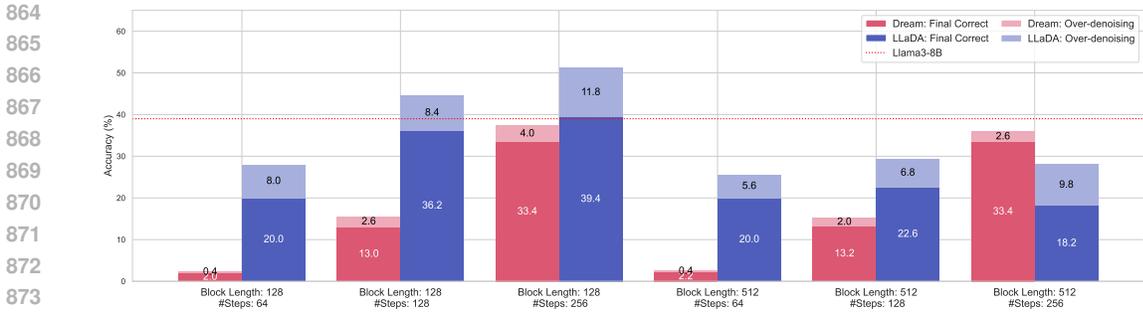


Figure 7: Comparison of LLaDA and Dream on MATH. Results indicate that the performance of Dream is highly dependent on the number of denoising steps, while LLaDA benefits more from the semi-AR setting (with block length 128).

To better understand the interplay between block size and the number of denoising steps, we conduct experiments with LLaDA-Instruct on the MATH-500 benchmark, varying $B \in \{1, 4, 16, 32, 128, 512\}$ and the number of steps $T \in \{32, 64, 128, 256, 512\}$. The results in Fig. 6 reveal several trends: (1) Small block sizes (e.g., $B = 4$) show slower initial gains but benefit more from increasing T , indicating that finer-grained autoregression allows iterative refinement over more steps; (2) Medium block sizes (e.g., $B = 16, 32, 128$) achieve the highest accuracy at moderate T , suggesting a favorable balance between parallelism and left-to-right dependency modeling; (3) pure-Diff (i.e., $B = 512$) consistently underperforms, and accuracy even drops at higher T , implying that excessive parallel denoising hinders the model’s ability to progressively refine solutions on math tasks. However, as discussed earlier following Table 2, pure-Diff is more flexible in dealing with tasks that require answers with varied length, e.g., Countdown. Therefore, improving the performance of pure-Diff is as well crucial to the wider application of masked diffusion language models.

B.4 LLADA AND DREAM

As an extended analysis, we choose another pre-trained MDLM Dream 7B (Ye et al., 2025) to compare with LLaDA. Note that Dream 7B builds on a different algorithm to LLaDA and initializes from an auto-regressive (AR) checkpoint Qwen2.5 7B (Yang et al., 2024). Though the algorithm still follows the absorbing discrete diffusion framework to transit tokens to and from the predefined absorbing state (masking token), Dream differs in predicting the next token and filling the next token with the prediction if the next token is masked. This technical difference originates from the fact that Dream is adapted from an AR model which is trained on next token prediction. We observe a clear behavioral difference between the pre-trained from scratch LLaDA and AR-adapted Dream, on MATH task shown in Fig. 7.

The first observation is that the performance of Dream is highly dependent on the denoising steps, but not the block length. In the semi-AR setting (Block Length of 128 in Fig. 7), which substantially improves the performance of LLaDA, Dream achieves downstream performance similar to the pure-Diff setting. We hypothesize that this arises because the model is adapted from an AR model, where earlier tokens consistently receive higher confidence. Consequently, pure-Diff effectively converges toward semi-AR, since tokens in early positions typically obtain higher confidence scores and are rarely remasked. Evidence for this convergence is shown in Fig. 8, where Dream predominantly assigns higher confidence to earlier masked tokens.

Our second observation is that Dream exhibits a smaller degree of over-denoising compared to LLaDA. We hypothesize that this difference also stems from its adaptation from AR models. While we see over-denoising as a symptom of the training–inference divide inherent in masked diffusion, Dream is pre-trained on a mixture of objectives, including both next-token prediction and masked diffusion. In particular, the pre-training of Dream on the masked diffusion objective consumes 580 billion tokens, whereas the original initialization Qwen 2.5 pre-training covers 18 trillion tokens. We argue that this imbalance in the scale of pre-training tokens between the two stages reduces the extent to which Dream experiences the training–inference divide, making its behavior more similar to that of an AR model. As a result, the discrepancy is less pronounced in Dream, and applying our proposed methods might yield only marginal improvements.

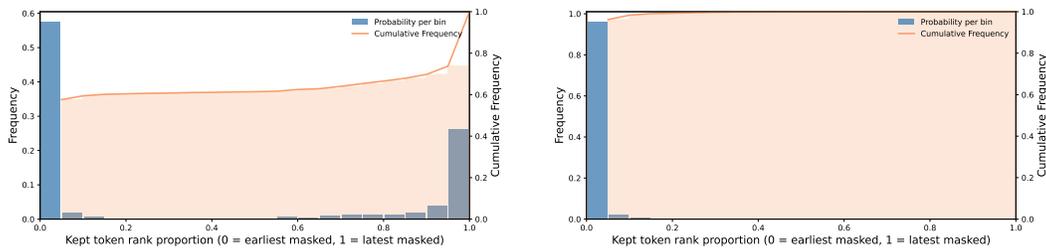


Figure 8: Quantitative analysis of predicted and ‘kept’ (not remasked) tokens in LLaDA and Dream. We apply confidence-based remasking with pure-Diff to both models and plot the distribution of retained tokens by their rank among masked positions at respective denoising steps across trajectories. The results show that Dream tends to assign higher confidence to earlier masked tokens, resembling AR models, whereas LLaDA exhibits a two-sided pattern, giving higher confidence to both early and late masked tokens.

C ADDITIONAL DISCUSSION ON EXPERIMENTAL RESULTS

C.1 WHY OVER-DENOISING IS AN EFFECTIVE DATA FILTER?

In Section 3.3, we demonstrate that training solely on over-denoising samples leads to significant performance improvements. Our intuition for this advantage is twofold. (1) Over-denoising samples yield more informative gradients because predictions often flip back and forth between correct and incorrect answers. This dynamic allows the model to learn from denoising steps that steer predictions in both directions, thereby stabilizing its outputs over time. (2) Advantages are mostly non-zero if we use over-denoising samples for online sampling. As we only use the binary rewards in our settings, advantages are either 0 or 1 for any denoising step. For denoising steps that fail to produce correct answers, rewards are uniformly zero, and exploration collapses if all rollouts in a group obtain zero rewards. This is precisely the well-known zero-advantage problem in training LLMs with RL (Yu et al., 2025a; Zhao et al., 2025b). To solve this issue, Zhao et al. (2025b) proposed a novel inpainting method for MDLMs to steer exploration toward promising trajectory spaces, thereby mitigating the zero-advantage issue and improving sample efficiency. We view open-denoising filtering as an alternative approach to addressing the same problem.

D THE USE OF LARGE LANGUAGE MODELS

Large language models are used in our early-stage writing for wording and grammar checking, as well as searching for missing literature. LLMs are not involved in the later iterations of the paper writing. Therefore, we do not consider LLMs as significant contributors to this paper.

