

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 DON'T SETTLE TOO EARLY: SELF-REFLECTIVE RE-MASKING FOR DIFFUSION LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

<https://github.com/iiiutch-ii/RemeDi>

ABSTRACT

Mask-based Diffusion Language Models (DLMs) struggle to revise incorrect tokens: once a token is generated, it typically remains fixed. The key challenge is to identify potential errors in the inputs. In this paper, we propose *Remasking-enabled Diffusion Language Model (RemeDi)*, a mask-based DLM that introduces *remasking* as another fundamental mechanism, enabling more flexible text refinement in diffusion-based text generation. To achieve this, RemeDi jointly predicts token distributions and per-token confidence scores at each step. The confidence scores determine which tokens to be unmasked after the current step, allowing the model to identify tokens with low quality and remask them. These remasked tokens can be resampled with richer context in subsequent steps. We design a remask-aware pipeline to train this ability, including supervised fine-tuning which teaches the model to detect and remask incorrect tokens in addition to predict mask tokens, and reinforcement learning which optimizes full generation trajectories toward higher rewards. Experiments show that RemeDi achieves the state-of-the-art results among open-source DLMs on multiple datasets.

1 INTRODUCTION

Diffusion Language Models (DLMs) have recently emerged as a promising alternative to autoregressive language models (Nie et al., 2025; Ye et al., 2025; Lou et al., 2024; Arriola et al., 2025). A DLM defines a forward process that gradually corrupts text into a noise prior, and learns a reverse process to recover clean text (Campbell et al., 2022; Lou et al., 2024). Unlike autoregressive models, DLMs do not commit to a fixed left-to-right order, offering greater flexibility in generation and an inherent ability to predict multiple tokens in parallel.

A dominant variant is the *mask-based* DLM (Nie et al., 2025; Ye et al., 2025), where the noise is represented by a special mask token. Under this formulation, the model learns to recover masked tokens during training, while assuming that once tokens are unmasked, they are supposed to be correct without having to clean them later. This assumption is problematic: the model may generate wrong tokens, which should be revealed and corrected in later steps when more contexts are available. However, most existing DLMs (Nie et al., 2025; Ye et al., 2025) keep already unmasked tokens fixed, preventing them from being revised by self-reflecting on errors.

To address this, several works have explored methods to revise generated tokens. von Rütte et al. (2025) defines a new noise schedule by interpolating between masking and uniform noise, enabling revision of wrong tokens on small-scale models. Wang et al. (2025a) applies predictor-corrector samplers by stochastically remasking a subset of tokens only at inference time, where the remasking is performed randomly without training the model how to find and remask incorrect tokens. For large-scale DLMs, Seed Diffusion (Song et al., 2025) allows all tokens to be resampled at every step. However, it lacks a mechanism to ensure that the number of mask tokens decreases monotonically — it is a key feature for diffusion models to ensure decreasing noise levels over steps (Guo et al., 2025), so that the mask tokens will eventually vanish at the final step to complete the generation.

In this paper, we propose a self-reflective remasking approach to train DLMs. As illustrated in Fig. 1a, it aims to train DLMs with the ability of finding wrong tokens and turning them back to mask ones so that they can be resampled with richer context in later steps. Based on this, we introduce *Remasking-enabled Diffusion Language Model (RemeDi)*, a mask-based DLM that incorporates self-reflective remasking to revise already generated but incorrect tokens. RemeDi jointly predicts

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

denoising steps

... Calculate how many grapes were left for the pies...
 ... Calculate how many grapes were left for the pies...
 ... Calculate how many grapes were left for the pies...
 ... Calculate how many grapes were used for the pies...

(a) Green and Red tokens are already unmasked in the input sequence of the current step. Red tokens are remasked. Blue tokens are unmasked in the outputs. Gray tokens keep masked in the outputs, and we display the token with the highest probability at these positions.

Figure 1: (a) Illustration of remasking for quality improvement: RemeDi initially predicts the token “left”, but later identifies a semantic mismatch in the phrase “left for the pies”. The model then remasks this token and corrects it to the more appropriate “used”. (b) RemeDi outperforms existing DLMs in various tasks, including math, code and general benchmark.

token distributions and per-token confidence scores. At each diffusion step, high-confidence tokens are unmasked while low-confidence ones are (re-)masked, regardless of whether they have been previously unmasked.


The key challenge is to *train* the model how to remask incorrect tokens in a self-reflective manner. To this end, we design a remask-aware training pipeline in two stages: 1) **Remask SFT**, where the model learns to identify and remask incorrect tokens, while predicting masked tokens. We construct an input sequence for Remask SFT by randomly masking its tokens or replacing them with random alternatives to simulate the noise. The noise schedule deciding how many tokens are masked or randomly replaced is designed to follow the criterion that the noise level should monotonically decrease over steps. The model is then trained to remask and revise incorrect tokens over noisy input sequences. 2) **Remask RL**, where the model is further fine-tuned with outcome-based reinforcement learning. It seeks to optimize the entire generation trajectories toward final outputs with higher rewards by considering how to remask and predict tokens in each step.

As shown in Fig. 1b, RemeDi achieves the state-of-the-art performance among open-source DLMs, achieving competitive results on various benchmark datasets, including math problems (89.1% on GSM8K (Cobbe et al., 2021), 52.9% on MATH (Hendrycks et al., 2021)), code generation (73.2% on HumanEval (Chen et al., 2021), 59.4% on MBPP (Austin et al., 2021)), and general tasks (24.5% on AlpacaEval (Dubois et al., 2024), 85.4% on IFEval (Zhou et al., 2023), and 87.7% on ARC-C (Clark et al., 2018)).

2 RELATED WORK

2.1 MASK-BASED DIFFUSION LANGUAGE MODELS

Diffusion language models (DLMs) have emerged as promising alternatives to auto-regressive (AR) models for text generation. Among them, mask-based DLMs (Nie et al., 2025; Ye et al., 2025; Zheng et al., 2023; Ou et al., 2024) dominate, generating text by progressively denoising mask tokens. Recent studies(Arriola et al., 2025; Fathi et al., 2025; Huang & Tang, 2025; Sahoo et al., 2025; Wang et al., 2025b; Gat et al., 2025) have increasingly explored the fusion of AR and diffusion models, often through an iterative block-wise decoding strategy: inference proceeds by iteratively appending a block of mask tokens to the input sequence and denoising it, repeating until the EOS token is generated. This paradigm inherits the strengths — flexible generation order and parallel decoding from DLMs, and cache efficiency from AR — yielding faster inference without sacrificing



(b) Radar plot comparing the performance of RemeDi with other DLMs across various evaluation benchmarks.

108 quality. In our work, we adapt LLaDA-8B-Instruct (Nie et al., 2025) to variable-length block-by-
 109 block generation, serving as the backbone for our remasking mechanism.
 110

111 **2.2 REVISING ERRORS IN DIFFUSION LANGUAGE MODELS**
 112

113 A key limitation of mask-based DLMs is their inability to revise tokens once unmasked, even if they
 114 are incorrect. Existing efforts to address this fall into two categories. The first category (Campbell
 115 et al., 2022; Wang et al., 2025a) applies predictor-corrector samplers without training, for example
 116 by stochastically remasking a subset of tokens during inference. These methods lack a mechanism
 117 to identify which tokens are actually wrong. As a result, they have to rely on many extra sampling
 118 steps to take effect, which are inefficient and hard to optimize. The second category modifies the
 119 diffusion process to enable revision during the reverse diffusion process, e.g., combining mask diffu-
 120 sion process with either the uniform diffusion process (von Rütte et al., 2025) or edit-based diffusion
 121 process (Havasi et al., 2025; Song et al., 2025).

122 In short, none of these approaches provides a principled way to detect and selectively correct errors
 123 during generation. In contrast, RemeDi fulfills self-reflection by identifying and remasking error-
 124 prone tokens through a two-stage learning pipeline, and jointly training the model to resample the
 125 remasked tokens in later steps.
 126

127 **3 METHODS**
 128

129 **3.1 PRELIMINARIES: MASK-BASED DIFFUSION LANGUAGE MODELS**
 130

131 Diffusion Language Models (DLMs) aim to model text generation by approximating the probabil-
 132 ity distribution p_{data} over a finite vocabulary $\mathcal{V} = \{1, 2, \dots, V\}$. They define a discrete diffusion
 133 process in which the unknown data distribution p_{data} at $t = 0$ gradually evolves into a simple prior
 134 distribution p_{prior} at $t = T$ (Lou et al., 2024). At intermediate times t , we denote the distribution as
 135 p_t . Formally, this diffusion process can be described by a linear ODE involving a diffusion matrix
 136 Q_t :

$$\frac{dp_t}{dt} = Q_t p_t, \quad p_0 = p_{\text{data}}, \quad p_T = p_{\text{prior}}. \quad (1)$$

137 While t can be defined continuously as in Eq. (1), in practice we work with discrete timesteps $t_{0:N}$.
 138

139 In this paper, we focus on mask-based DLMs, where p_{prior} is a distribution that puts all its mass
 140 on the mask state, denoted as [M]. Given a clean sequence $x_0 \sim p_{\text{data}}$, a corrupted sequence x_t is
 141 obtained by randomly replacing part of the tokens with the mask token [M]. The model is trained to
 142 recover x_0 by predicting each mask token x_t^i with the output probability $p_\theta^i(x_0^i | x_t)$. The objective
 143 is:
 144

$$L_{\text{diffusion}}(\theta) = \mathbb{E}_{t, x_0, x_t} \left[-\frac{1}{t} \sum_{i=1}^L \mathbf{1}(x_t^i = [\text{M}]) \log p_\theta^i(x_0^i | x_t) \right], \quad (2)$$

145 where x_t is a sequence of length L , sampled from the forward process, $\mathbf{1}(\cdot)$ is an indicator function
 146 ensuring that the loss is computed only on mask tokens, following Nie et al. (2025).
 147

148 During inference, the reverse diffusion process begins with a sequence of only mask tokens and
 149 proceeds for N steps at monotonically decreasing timesteps $t_{0:N}$. At step t_{n-1} , the model takes the
 150 partially masked sequence $x_{t_{n-1}}$ as input and predicts all mask tokens simultaneously. A subset of
 151 tokens is unmasked to obtain x_{t_n} according to the noise schedule and the unmasking policy (e.g.,
 152 unmasking tokens with the highest confidence), while the remaining predictions are remasked and
 153 deferred to later steps.
 154

155 A limitation of this paradigm is that once a token is unmasked, it remains fixed in subsequent steps.
 156 In early stages, limited unmasked tokens often lead to unreliable predictions, resulting in errors
 157 that persist through the remainder of the generation process. As generation progresses, additional
 158 context may reveal these errors, but current paradigm offers no way to correct them. This motivates
 159 the ability to *remask* tokens, allowing the model to remask earlier predictions back to the mask token
 160 and predict them again using richer context in later steps.
 161




Figure 2: The structure of RemeDi, including Unmasking Policy Stream (UPS) to predict confidences h_θ for selecting the set of unmasking tokens \mathcal{U}_n , and Token Prediction Stream (TPS) to predict the token value when unmasking a masked position.

3.2 REMEDI

We propose **RemeDi**, a DLM that can identify and remask low-confidence tokens during generation to enable iterative self-reflection. We extend the standard transformer into a dual-stream transformer architecture as shown in Fig. 2, which comprises:

- **Token Prediction Stream (TPS):** A stack of transformer blocks that predict probabilities $p_\theta^i(\cdot|x_t)$ for masked tokens as in a typical DLM (Nie et al., 2025).
- **Unmasking Policy Stream (UPS):** Another stack of transformer blocks that output token-wise confidence score h_θ^i . It represents the model’s confidence over the output tokens, indicating if they should be unmasked with high confidence. Otherwise, if the confidence is too low for a token, it should be kept masked or remasked so that it could be sampled or resampled later.

The two streams run in parallel. During inference, UPS is inserted periodically and receives hidden states from TPS as input, producing an auxiliary representation f_{UPS} for confidence scoring. The two streams perform bidirectional feature sharing: UPS layers are conditioned on f_{TPS} , and their outputs also feed back into TPS to enrich its representations. At the final layer, p and h are produced simultaneously using two independent linear heads applied to f_{TPS} and f_{UPS} , respectively. More details about the model structure for these two streams can be found in Appendix B.1.

The token generation proceeds through iterative denoising steps. Given $x_{t_{n-1}}$ as the input, UPS first predicts a confidence score $h_{\theta,n}^i$ at each position i , and select a subset of positions \mathcal{U}_n to unmask at the current step. Then, for the positions selected to be unmasked, if they have already been unmasked in $x_{t_{n-1}}$, they remain unchanged; otherwise they are sampled from $p_\theta^i(\cdot|x_{t_{n-1}})$ predicted by TPS. Unlike existing mask-based DLMs where tokens are fixed once being unmasked, RemeDi re-decides a token to be unmasked or (re-)masked at each step by its trained confidence score. Thus, it is possible that an already generated token is assigned with a low confidence and remasked, allowing it to be resampled in later steps. A noise schedule controls that the total number of unmasked tokens increases linearly from 0 to L (Nie et al., 2025), so that the number of mask tokens approaches zero at the final step.

In the following sections, we elaborate on how to train RemeDi with Remask SFT and Remask RL algorithms.

216 3.2.1 REMASK SFT
217

218 Traditional mask-based DLMs conduct SFT with randomly masked input sequences (Nie et al.,
219 2025; Lou et al., 2024), while RemeDi needs to detect and remask possible incorrect tokens that
220 arise during the reverse diffusion process, so they can be resampled in later steps. To achieve this,
221 in SFT we treat such incorrect tokens as a second noise type in addition to the first noise type of
222 mask tokens in mask-based DLMs, and train the model to recover mask tokens as well as identify
223 unmasked tokens that should be remasked.

224 To simulate inference inputs at a diffusion time t , we construct training samples x_t from clean text
225 x_0 by applying two types of noise: given a randomly sampled diffusion time $t \in (0, 1)$, we set the
226 corresponding mask ratio $\rho_{t,\text{mask}}$, alongside the incorrect token ratio $\rho_{t,\text{incorrect}}$. With both ratios, we
227 randomly mask tokens with $\rho_{t,\text{mask}}$. Then, among the remaining unmasked positions, we sample a
228 subset with the ratio $\rho_{t,\text{incorrect}}$ and replace each selected token with a random alternative to simulate
229 the incorrect tokens that may occur in the reverse diffusion process.

230 As aforementioned, during the reverse diffusion process, the noise level, defined as the number of
231 mask tokens, should decrease monotonically (Guo et al., 2025). Since all incorrect tokens in an
232 input sequence of length L must be remasked as designed below for the SFT, we require:

$$233 \lceil \rho_{t,\text{incorrect}} \cdot (1 - \rho_{t,\text{mask}}) \cdot L \rceil < \lceil \rho_{t,\text{mask}} \cdot L \rceil \quad (3)$$

235 to ensure a monotonically decreasing number of mask tokens as outputs. Otherwise, remasking all
236 incorrect tokens would increase the total number of masks in the next step, violating the principle
237 that the number of mask tokens should decrease at each diffusion step.

238 Considering the above inequality, we choose $\rho_{t,\text{mask}} = t$ and $\rho_{t,\text{incorrect}} = 4r \cdot t(1-t)$ (r is a constant)
239 following (Nie et al., 2025; von Rütte et al., 2025). We set $r = 0.1$ in our experiments, under which
240 it is not hard to see that the inequality 3 always holds on $t \in [0, 1]$.

241 **Remask SFT Algorithm.** During training, in addition to the typical diffusion loss in Eq. 2, we
242 supervise the unmasking score h_θ with a binary cross-entropy (BCE) objective across all token
243 positions. We construct the training label y based on different token types:

- 245 • A clean token ($i \in \mathcal{S}_{\text{clean}} = \{i \mid x_t^i = x_0^i\}$) receives a positive unmask label $y^i = 1$,
246 indicating they should remain unmasked.
- 248 • An incorrect token ($i \in \mathcal{S}_{\text{incorrect}} = \{i \mid x_t^i \neq x_0^i, x_t^i \neq [\text{M}]\}$) receives a negative unmask
249 label $y^i = 0$, indicating that they should be remasked.
- 250 • A mask token ($i \in \mathcal{S}_{\text{mask}} = \{i \mid x_t^i = [\text{M}]\}$) is assigned a soft unmask label $y^i = p_\theta^i(x_0^i | x_t)$,
251 equal to the predicted probability of the ground-truth token x_0^i . A higher probability indi-
252 cates a higher likelihood that the predicted token is correct and thus should be unmasked.

253 With unmask labels assigned above, we seek to minimize

$$255 \mathcal{L}_{\text{UPS}}(\theta) = \sum_i \text{BCE}(\sigma(h_\theta^i), y^i), \quad (4)$$

258 where $\sigma(\cdot)$ is the sigmoid function. Thus, the overall Remask SFT objective is:

$$259 \mathcal{L}(\theta) = \mathcal{L}_{\text{diffusion}}(\theta) + \lambda_{\text{UPS}} \mathcal{L}_{\text{UPS}}(\theta), \quad (5)$$

261 where λ_{UPS} balances the two losses.

262 Finally, we summarize the Remask SFT in Algorithm 1, where we elaborate on how to construct the
263 input sequence and calculate the loss function for the Remask SFT.

265 3.2.2 REMASK RL
266

267 After training with Remask SFT, we further fine-tune the model with outcome-based reinforcement
268 learning (RL) to optimize the full generation trajectory (Huang et al., 2025). Specifically, we rein-
269 force the generation process with N denoising steps, beginning from an all-mask prior $x_{t_0} \sim p_{\text{prior}}$
at $t_0 = 1$ and proceeding through timesteps $t_{0:N}$.

Algorithm 1 Input sequence construction and loss calculation in Remask SFT

270
 271
 272 **Require:** Clean sequence $x_0 = [x_0^1, \dots, x_0^L]$ of length L . Model \mathcal{M} with learnable parameters θ .
 273 1: Sample $t \in (0, 1)$ according to the noise schedule, obtaining $\rho_{t,\text{mask}}$ and $\rho_{t,\text{incorrect}}$.
 274 2: **Construct noisy input x_t :**
 275 3: For each position i , replace x_0^i with $[\text{M}]$ w.p. $\rho_{t,\text{mask}}$
 276 4: Among remaining positions, replace x_0^i with a random alternative token w.p. $\rho_{t,\text{incorrect}}$
 277 5: Define index sets:
 278
$$\mathcal{S}_{\text{mask}} = \{i \mid x_t^i = [\text{M}]\}, \quad \mathcal{S}_{\text{incorrect}} = \{i \mid x_t^i \neq x_0^i \wedge x_t^i \neq [\text{M}]\}, \quad \mathcal{S}_{\text{clean}} = \{i \mid x_t^i = x_0^i\}$$

 279 6: Get model outputs: $[p_\theta, h_\theta] = \mathcal{M}(x_t; \theta)$
 280 7: Calculate the diffusion loss, on mask tokens only: $\mathcal{L}_{\text{diffusion}}(\theta) = -\frac{L}{|\mathcal{S}_{\text{mask}}|} \sum_{i \in \mathcal{S}_{\text{mask}}} \log p_\theta^i(x_0^i | x_t)$
 281 8: Get labels for UPS: $y^i = \begin{cases} 1 & i \in \mathcal{S}_{\text{clean}} \\ 0 & i \in \mathcal{S}_{\text{incorrect}} \\ \text{stopgrad}(p_\theta^i(x_0^i | x_t)) & i \in \mathcal{S}_{\text{mask}} \end{cases}$
 282 9: UPS BCE loss:
$$\mathcal{L}_{\text{UPS}}(\theta) = -\frac{1}{L} \sum_{i=1}^L \left(y^i \log \sigma(h_\theta^i) + (1 - y^i) \log (1 - \sigma(h_\theta^i)) \right)$$
 $\triangleright \sigma(\cdot)$ represents the sigmoid function
 283 10: Total loss: $\mathcal{L}(\theta) = \mathcal{L}_{\text{diffusion}}(\theta) + \lambda_{\text{UPS}} \mathcal{L}_{\text{UPS}}(\theta)$

 290
 291
 292 At each step t_n , RemeDi generates x_{t_n} from $x_{t_{n-1}}$ by invoking two coupled policies: an *unmasking policy* that chooses a subset of positions $\mathcal{U}_n = [u_n(1), \dots, u_n(K_n)]$ to unmask, and a *token prediction policy* that samples tokens at the chosen positions. Unlike standard DLMs, which never remask revealed tokens, RemeDi allows previously unmasked tokens to be remasked, enabling revision of earlier predictions.
 293
 294
 295
 296
 297 **Unmasking policy.** The UPS produces a per-token confidence score h_θ^i , indicating how strongly the model believes token at position i is correct (if unmasked) or predictable (if masked). At inference, we rank tokens by their confidence scores and prioritize high-confidence ones to unmask. The number of unmasked tokens K_n at each diffusion step is determined by linearly increasing from 0 to L . During RL training, we construct an unmasking policy to sample $\mathcal{U}_n = [u_n(1), \dots, u_n(K_n)]$ using the Plackett–Luce model (Plackett, 1975): based on h_θ , we use a multinomial distribution and sequentially sample K_n positions from $\{1, \dots, L\}$ without replacement. Formally, the probability of sampling \mathcal{U}_n is:
 298
 299
 300
 301
 302
 303
 304

$$\pi_{\theta,n}^{\text{unmask}}(\mathcal{U}_n \mid x_{t_{n-1}}) = \prod_{k=1}^{K_n} \frac{\exp(h_{\theta,n}^{u_n(k)})}{\sum_{j \notin \{u_n(1), \dots, u_n(k-1)\}} \exp(h_{\theta,n}^j)}. \quad (6)$$

305
 306 **Token prediction policy.** For each position $i \in \mathcal{U}_n$, if $x_{t_{n-1}}^i = [\text{M}]$, the model samples token from
 307 $p_\theta^i(\cdot | x_{t_{n-1}})$; otherwise, the token remains unchanged as in the input. The probability of generating
 308 x_{t_n} given $x_{t_{n-1}}$ and \mathcal{U}_n is:
 309
 310
 311

$$\pi_{\theta,n}^{\text{token}}(x_{t_n} \mid x_{t_{n-1}}, \mathcal{U}_n) = \prod_{i \in \mathcal{U}_n, x_{t_{n-1}}^i = [\text{M}]} p_\theta^i(x_{t_n}^i \mid x_{t_{n-1}}). \quad (7)$$

312 **Joint policy.** Thus, the probability of transitioning from $x_{t_{n-1}}$ to x_{t_n} is the product of the unmasking
 313 probability and the token prediction probability:
 314
 315

$$\pi_{\theta,n}(x_{t_n} \mid x_{t_{n-1}}) = \pi_{\theta,n}^{\text{unmask}}(\mathcal{U}_n \mid x_{t_{n-1}}) \cdot \pi_{\theta,n}^{\text{token}}(x_{t_n} \mid x_{t_{n-1}}, \mathcal{U}_n). \quad (8)$$

316 With the probability defined in Eq. 8, we apply outcome-based reinforcement learning to encourage
 317 generation trajectories $x_{t_{0:N}}$ that lead to correct final responses x_{t_N} . Specifically, we adopt GRPO
 318 (Shao et al., 2024), a scalable RL paradigm for language models. The reward is defined according
 319 to task type: verifiable correctness for math and code, and reward-model evaluation for open-ended
 320 questions. Further details on datasets and reward design are provided in Appendix B.2.4.
 321
 322
 323

As shown in Fig. 11 of Appendix A, after Remask SFT and RL training, the learned h_θ serves as a reliable indicator to assess the quality of input tokens. Tokens already unmasked in the input typically receive high confidence scores. However, when certain tokens are assigned low confidence, they are more likely to be inadequate and are remasked for re-prediction in subsequent steps. It suggests that the UPS-predicted confidence scores provide a reliable estimate of per-token quality for the unmasking policy.

4 EXPERIMENTS

RemeDi enables remasking on a DLM capable of variable-length block-wise generation (Arriola et al., 2025) to support *variable-length generation*, a key feature for enabling the real-world DLM to generate an unfixed number of blocks (see Appendix B.2.2 for details). Since there are no open-source large-scale variable-length block-wise DLMs, we adapt our model from LLaDA, a widely used benchmark DLM. Starting from LLaDA’s model weights as initialization, RemeDi undergoes two stages of supervised fine-tuning and RL. We detail the training configurations in Appendix B.2, and the evaluation metrics in Appendix B.3.2.

Table 1: Model performance on math and code generation benchmarks. We highlight the best-performing model among compared DLMs in **bold**. “-” indicates unknown cases not mentioned in original papers.

Method	Math			Code	
	GSM8K	MATH	GPQA	HumanEval	MBPP
Diffusion Language Models					
Dream (Ye et al., 2025)	82.1	49.6	30.6	59.8	59.6
LLaDA (Nie et al., 2025)	78.3	38.9	28.1	45.7	39.0
LLaDA + ReMDM (Wang et al., 2025a)	81.4	38.5	-	44.5	37.8
d1-LLaDA (Zhao et al., 2025)	82.1	-	-	37.8	44.7
wd1-LLaDA (Tang et al., 2025)	82.3	-	-	-	-
LLaDA 1.5 (Zhu et al., 2025)	83.3	42.6	36.9	52.4	42.8
LLaDOU (Huang et al., 2025)	88.1	44.6	-	59.1	51.6
RemeDi (+ Remask SFT)	86.3	51.4	32.6	71.3	57.8
RemeDi (++) Remask RL	89.1	52.9	29.5	73.2	59.4
Auto-regressive Models					
LLaMA2 7B (Touvron et al., 2023)	14.6	2.5	28.4	12.8	20.8
MetaMath 7B (Yu et al., 2023)	66.5	19.8	-	-	-
CodeLLaMA 7B (Roziere et al., 2023)	-	-	-	34.8	44.4
Deepseek 7B (Bi et al., 2024)	63.0	15.8	-	48.2	35.2
DeepseekMath 7B (Shao et al., 2024)	88.2	51.7	-	-	-
DeepseekCoder 7B (Guo et al., 2024)	-	-	-	66.1	65.4
LLaMA3 8B (Dubey et al., 2024)	78.3	29.6	31.9	59.8	57.6
Gemma2 9B (Team, 2024)	76.7	44.3	32.8	68.9	74.9

4.1 RESULTS

To evaluate various capabilities of RemeDi in different aspects, we conducted detailed comparisons against existing large language models of comparable scale in Tab. 1 and Tab. 2, including both DLMs and auto-regressive models. We select nine popular benchmarks across general tasks, mathematics, coding, and human preference domains.

After Remask SFT, RemeDi demonstrates high performance on almost all these benchmarks. It not only achieves the state of the art performance among existing DLMs, but also outperforms auto-regressive models of similar model size. On math benchmarks, RemeDi after Remask SFT achieves 86.3% on GSM8K and 51.4% on MATH, surpassing MetaMath with math-specific instruction tuning. It is even on par with DeepseekMath using math-specific reinforcement learning. On code generation benchmarks, RemeDi achieves 71.3% on HumanEval, outperforming CodeLLaMA and Deepseek Coder. For general natural language tasks, RemeDi also demonstrates strong performance

378 Table 2: Model performance on general tasks. We highlight the best-performing model among
 379 compared DLMs in **bold**. “-” indicates unknown cases not mentioned in original papers.
 380

Method	Hellaswag	ARC-C	IFEval	AlpacaEval
Diffusion Language Models				
Dream (Ye et al., 2025)	70.3	79.2	67.5	5.9
LLaDA (Nie et al., 2025)	69.7	83.9	70.0	11.2
LLaDA 1.5 (Zhu et al., 2025)	70.5	83.5	73.5	13.9
RemeDi (+ Remask SFT)	71.1	85.2	81.9	12.5
RemeDi (++ Remask RL)	72.2	87.7	85.4	24.8
Auto-regressive Models				
LLaMA2 7B (Touvron et al., 2023)	51.5	57.3	-	-
Deepseek 7B (Bi et al., 2024)	68.5	49.4	-	-
LLaMA3 8B (Dubey et al., 2024)	75.5	82.4	-	-

388 ...A statistical model is a mathematical representation that explains how data is generated, often
 389 in a simplified and idealized way. It forms the foundation for understanding the data, making
 390 hypotheses about, and making and and and population...
 391

392 **step 420**

393 ...A statistical model is a mathematical representation that explains how data is generated, often
 394 in a simplified and idealized way. It forms the foundation for understanding the data, ~~making~~ tests
 395 and estimators, and making the inference inference...
 396

397 **step 424**

402 ...A statistical model is a mathematical representation that explains how data is generated, often
 403 in a simplified and idealized way. It forms the foundation for understanding the data, developing
 404 tests and estimators, and making the basis statistical...
 405

406 **step 425**

407 Figure 3: An example of the step-by-step generation process. **Green** and **Red** are already unmasked
 408 in the inputs. **Red** tokens are remasked. **Blue** tokens are unmasked in the outputs. **Gray** tokens
 409 remain masked, and we display the token with the highest probability at these positions. More
 410 examples can be found in AppendixA.
 411

412 in common knowledge answering (85.2% on ARC-C) and instruction following (81.9% on IFEval)
 413 tasks. It also aligns well with human preference (12.5% on AlpacaEval), outperforming other DLMs
 414 such as Dream and LLaDA.

415 After Remask RL, RemeDi achieves further improvements across a wide range of math, coding and
 416 general tasks. For example, the accuracies on GSM8K and MATH reach 89.1% and 52.9% respec-
 417 tively, outperforming all compared DLMs and AR models. Among all the benchmarks, RemeDi
 418 achieves its most substantial improvement on the AlpacaEval (Dubois et al., 2024) benchmark, with
 419 a +12.3% gain over the Remask SFT model. This demonstrates the effectiveness of our approach in
 420 post-training the model’s ability for a broad range of tasks.
 421

4.2 VISUALIZATION AND ANALYSIS

424 We visualize how remasking improves text generation in RemeDi in Fig. 3. The model initially gen-
 425 erated the token “making.” After generating the object “tests and estimators,” it found that “making”
 426 is not the proper verb in this verb-object structure. Thus, the model remasks it and opts for the
 427 more appropriate “developing.” This example shows RemeDi’s ability to iteratively refine its output
 428 content. We provide more examples in Appendix A, demonstrating that RemeDi is able to perform
 429 a variety of operations such as replacing, inserting and deleting with the remask mechanism.

430 To provide a quantitative analysis, we calculate the frequencies of remasking in a block of length
 431 32 on MATH-500 (Lightman et al., 2023), HumanEval (Chen et al., 2021), and AlpacaEval (Dubois
 432 et al., 2024). In Fig. 4, we can see that remasking occurs most frequently in code generation, fol-




Figure 4: Distribution of remasking frequencies per block across different evaluation datasets. The numbers in parentheses indicate the mean and standard deviation for each dataset.

Table 3: Statistics of the remasking frequencies per block (block size is fixed to 32) when generating responses to questions with different difficulty levels in MATH-500.

Difficulty Level	Remasking Frequencies / Block	Accuracy
1	9.13 ± 9.54	86.04%
2	8.91 ± 7.29	80.21%
3	10.13 ± 8.64	64.48%
4	13.91 ± 11.44	50.00%
5	13.95 ± 11.12	19.25%

lowed by mathematical reasoning, and general tasks. This pattern may be attributed to differences in structural constraints: code requires strict syntactic correctness, and mathematical solutions demand formally structured derivations, whereas responses to open-ended problems allow more flexibility.

We also analyzed remasking frequencies across different difficulty levels on MATH-500, as shown in Tab. 3. RemeDi tends to remask more frequently as the difficulty increases, rising from about 9 tokens per block at level 1–2 to nearly 14 tokens at level 4–5. This pattern suggests that iterative refinement becomes increasingly necessary for harder problems.

4.3 ABLATION STUDIES

Remask SFT We compare the improvement brought by the Remask SFT (introduced in Sec. 3.2.1) with that of vanilla SFT, under the same training configuration detailed in Appendix B.2.5. We start from a baseline model that has already completed the warm-up phase tuning for variable-length block-wise generation, and perform training on the full code-category dataset and the open2math-1M-gpt-4.1-mini dataset mentioned in Appendix B.2.1. As shown in Tab. 4, Remask SFT outperforms vanilla SFT on all benchmarks, especially on MATH-500 (+2.6%) and HumanEval (+1.8%), demonstrating that Remask SFT is an effective training method to improve DLM’s performance.

Remask RL We compare Remask RL with LLaDOU RL (Huang et al., 2025), another algorithm that also reinforces the whole generation trajectories in the reverse diffusion process. Since LLaDOU RL is developed on LLaDA, we also implement Remask RL on LLaDA for the sake of fair comparison. All experiments are conducted on GSM8K with a generation length of 256, 64

Table 4: Experiment results after supervised tuning with different algorithms. The baseline model is already tuned to be a variable-length block-wise generation DLM (see Appendix B.2.2).

Method	GSM8K	MATH-500	HumanEval	MBPP
Baseline	80.3	34.7	41.5	42.6
Vanilla SFT	83.1	40.1	48.2	43.4
Remask SFT	83.6	42.7	50.0	44.0

486
487
488 Table 5: GSM8K pass@1 accuracy comparison between Remask and LLaDOU RL
489
490
491
492
493
494

Training Steps	Remask RL	LLaDOU RL
50	80.00%	77.58%
100	81.40%	78.86%
150	81.59%	80.00%
200	83.33%	82.35%

495 denoising steps, a block length of 64, and temperature 0.7, while all other hyperparameters follow
496 the LLaDOU setup (see Appendix B.2.5).

497 Remask RL demonstrates advantages in both convergence speed and performance. As shown in
498 Tab. 5, Remask RL achieves a higher final accuracy of 83.33%, with a particularly noticeable im-
499 provement in early training stages (e.g., 80.00% vs. 77.58% at step 50). This indicates that the more
500 flexible remask process contributes to both faster convergence and stronger model performance.
501

502 5 CONCLUSION

503 In this paper, we introduce the Remasking-enabled Diffusion Language Model (RemeDi), a new
504 self-reflective remasking mechanism to address the limitation of existing mask-based DLMs that
505 they cannot revise generated tokens. In RemeDi, remasking is achieved by predicting a confidence
506 score to identify noisy tokens, allowing them to be remasked and then resampled with richer context
507 in later steps.

508 Through a two-stage training pipeline of Remask SFT and Remask RL, RemeDi achieves the state-
509 of-the-art performance among open-source DLMs. Our analysis further shows that the learned con-
510 fidence scores provide a reliable signal of per-token quality during generation. RemeDi opens a
511 promising direction for self-reflective text generation, further releasing the full potentials of DLMs
512 to solve complex tasks with higher quality.

513 6 REPRODUCIBILITY STATEMENT

514 We provide an anonymous link containing the inference code and model weights, details of the
515 datasets and configurations used in both Remask SFT and RL in Appendix B.2, and the evaluation
516 settings in Appendix B.3.

517 REFERENCES

518 ajibawa 2023. Maths-college. Hugging Face Datasets. URL <https://huggingface.co/datasets/ajibawa-2023/Maths-College>.

519 Marianne Arriola, Aaron Gokaslan, Justin T Chiu, Zhihan Yang, Zhixuan Qi, Jiaqi Han, Sub-
520 ham Sekhar Sahoo, and Volodymyr Kuleshov. Block diffusion: Interpolating between auto-
521 regressive and diffusion language models. *arXiv preprint arXiv:2503.09573*, 2025.

522 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan,
523 Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language
524 models. *arXiv preprint arXiv:2108.07732*, 2021.

525 Akhiad Bercovich, Itay Levy, Izik Golan, Mohammad Dabbah, Ran El-Yaniv, Omri Puny, Ido Galil,
526 Zach Moshe, Tomer Ronen, Najeeb Nabwani, Ido Shahaf, Oren Tropp, Ehud Karpas, Ran Zil-
527 berstein, Jiaqi Zeng, Soumye Singhal, Alexander Bukharin, Yian Zhang, Tugrul Konuk, Ger-
528 ald Shen, Ameya Sunil Mahabaleshwaran, Bilal Kartal, Yoshi Suhara, Olivier Delalleau, Zijia
529 Chen, Zhilin Wang, David Mosallanezhad, Adi Renduchintala, Haifeng Qian, Dima Rekesh,
530 Fei Jia, Somshubra Majumdar, Vahid Noroozi, Wasi Uddin Ahmad, Sean Narenthiran, Alek-
531 sander Ficek, Mehrzad Samadi, Jocelyn Huang, Siddhartha Jain, Igor Gitman, Ivan Moshkov,
532 Wei Du, Shubham Toshniwal, George Armstrong, Branislav Kisacanin, Matvei Novikov, Daria

540 Gitman, Evelina Bakhturina, Jane Polak Scowcroft, John Kamalu, Dan Su, Kezhi Kong, Markus
 541 Kliegl, Rabeeh Karimi, Ying Lin, Sanjeev Satheesh, Jupinder Parmar, Pritam Gundecha, Bran-
 542 don Norick, Joseph Jennings, Shrimai Prabhumoye, Syeda Nahida Akter, Mostofa Patwary,
 543 Abhinav Khattar, Deepak Narayanan, Roger Waleffe, Jimmy Zhang, Bor-Yiing Su, Guyue
 544 Huang, Terry Kong, Parth Chadha, Sahil Jain, Christine Harvey, Elad Segal, Jining Huang,
 545 Sergey Kashirsky, Robert McQueen, Izzy Putterman, George Lam, Arun Venkatesan, Sherry
 546 Wu, Vinh Nguyen, Manoj Kilaru, Andrew Wang, Anna Warna, Abhilash Somasamudramath,
 547 Sandip Bhaskar, Maka Dong, Nave Assaf, Shahar Mor, Omer Ullman Argov, Scot Junkin, Olek-
 548 sander Romanenko, Pedro Larroy, Monika Katariya, Marco Rovinelli, Viji Balas, Nicholas Edel-
 549 man, Anahita Bhiwandiwalla, Muthu Subramaniam, Smita Ithape, Karthik Ramamoorthy, Yut-
 550 ing Wu, Suguna Varshini Velury, Omri Almog, Joyjit Daw, Denys Fridman, Erick Galinkin,
 551 Michael Evans, Katherine Luna, Leon Derczynski, Nikki Pope, Eileen Long, Seth Schneider,
 552 Guillermo Siman, Tomasz Grzegorzek, Pablo Ribalta, Monika Katariya, Joey Conway, Trisha
 553 Saar, Ann Guan, Krzysztof Pawelec, Shyamala Prayaga, Oleksii Kuchaiev, Boris Ginsburg,
 554 Oluwatobi Olabiyi, Kari Briski, Jonathan Cohen, Bryan Catanzaro, Jonah Alben, Yonatan Geif-
 555 man, Eric Chung, and Chris Alexiuk. Llama-nemotron: Efficient reasoning models, 2025. URL
 556 <https://arxiv.org/abs/2505.00949>.

557 Xiao Bi, Deli Chen, Guanting Chen, Shanhua Chen, Damai Dai, Chengqi Deng, Honghui Ding,
 558 Kai Dong, Qiushi Du, Zhe Fu, et al. Deepseek llm: Scaling open-source language models with
 559 longtermism. *arXiv preprint arXiv:2401.02954*, 2024.

560 Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and
 561 Arnaud Doucet. A continuous time framework for discrete denoising models. *Advances in Neural*
 562 *Information Processing Systems*, 35:28266–28279, 2022.

563 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared
 564 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri,
 565 Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan,
 566 Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian,
 567 Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios
 568 Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex
 569 Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders,
 570 Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec
 571 Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob Mc-
 572 Grew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large
 573 language models trained on code. 2021.

574 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and
 575 Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge.
 576 *ArXiv*, abs/1803.05457, 2018.

577 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
 578 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to
 579 solve math word problems, 2021. URL <https://arxiv.org/abs/2110.14168>, 9, 2021.

580 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
 581 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.
 582 *arXiv e-prints*, pp. arXiv–2407, 2024.

584 Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. Length-controlled al-
 585 pacaeval: A simple way to debias automatic evaluators. *arXiv preprint arXiv:2404.04475*, 2024.

586 Nima Fathi, Torsten Scholak, and Pierre-André Noël. Unifying autoregressive and diffusion-based
 587 sequence generation. *arXiv preprint arXiv:2504.06416*, 2025.

588 Itai Gat, Heli Ben-Hamu, Marton Havasi, Daniel Haziza, Jeremy Reizenstein, Gabriel Synnaeve,
 589 David Lopez-Paz, Brian Karrer, and Yaron Lipman. Set block decoding is a language model
 590 inference accelerator. *arXiv preprint arXiv:2509.04185*, 2025.

592 Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao
 593 Bi, Yu Wu, YK Li, et al. Deepseek-coder: When the large language model meets programming—the
 rise of code intelligence. *arXiv preprint arXiv:2401.14196*, 2024.

594 Zhehao Guo, Jiedong Lang, Shuyu Huang, Yunfei Gao, and Xintong Ding. A comprehensive review
 595 on noise control of diffusion model. *arXiv preprint arXiv:2502.04669*, 2025.
 596

597 Marton Havasi, Brian Karrer, Itai Gat, and Ricky TQ Chen. Edit flows: Flow matching with edit
 598 operations. *arXiv preprint arXiv:2506.09018*, 2025.
 599

600 Zhiwei He, Tian Liang, Jiahao Xu, Qiuwei Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian Yu,
 601 Zhenwen Liang, Wenxuan Wang, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and
 602 Dong Yu. Deepmath-103k: A large-scale, challenging, decontaminated, and verifiable mathemat-
 603 ical dataset for advancing reasoning. 2025. URL <https://arxiv.org/abs/2504.11456>.
 604

605 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 606 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *NeurIPS*,
 607 2021.
 608

609 Chihang Huang and Hao Tang. CtrlDiff: Boosting large diffusion language models with dynamic
 610 block prediction and controllable generation. *arXiv preprint arXiv:2505.14455*, 2025.
 611

612 Siming Huang, Tianhao Cheng, Jason Klein Liu, Jiaran Hao, Liuyihan Song, Yang Xu, J. Yang,
 613 J. H. Liu, Chenchen Zhang, Linzheng Chai, Rui Feng Yuan, Zhaoxiang Zhang, Jie Fu, Qian Liu,
 614 Ge Zhang, Zili Wang, Yuan Qi, Yinghui Xu, and Wei Chu. Opencoder: The open cookbook for
 615 top-tier code large language models. 2024. URL <https://arxiv.org/pdf/2411.04905.pdf>.
 616

617 Zemin Huang, Zhiyang Chen, Zijun Wang, Tiancheng Li, and Guo-Jun Qi. Reinforcing the diffusion
 618 chain of lateral thought with diffusion language models. *arXiv preprint arXiv:2505.10446*, 2025.
 619

620 Jaeyeon Kim, Kulin Shah, Vasilis Kontonis, Sham Kakade, and Sitan Chen. Train for the
 621 worst, plan for the best: Understanding token ordering in masked diffusions. *arXiv preprint
 622 arXiv:2502.06768*, 2025.
 623

624 Hynek Kydlíček. Math-verify: A robust mathematical expression evaluation system. <https://github.com/huggingface/Math-Verify>, 2025.
 625

626 Jia LI, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Costa Huang,
 627 Kashif Rasul, Longhui Yu, Albert Jiang, Ziju Shen, Zihan Qin, Bin Dong, Li Zhou, Yann
 628 Fleureau, Guillaume Lample, and Stanislas Polu. Numinamath. [<https://huggingface.co/AI-MO/NuminaMath-CoT>] (https://github.com/project-numina/aimo-progress-prize/blob/main/report/numina_dataset.pdf), 2024.
 629

630 Jijie Li, Li Du, Hanyu Zhao, Bowen Zhang, Liangdong Wang, Boyan Gao, Guang Liu, and Yonghua
 631 Lin. Infinity instruct: Scaling instruction selection and synthesis to enhance language models,
 632 2025. URL <https://arxiv.org/abs/2506.11116>.
 633

634 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
 635 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth
 636 International Conference on Learning Representations*, 2023.
 637

638 Chris Yuhao Liu, Liang Zeng, Jiacai Liu, Rui Yan, Jujie He, Chaojie Wang, Shuicheng Yan, Yang
 639 Liu, and Yahui Zhou. Skywork-reward: Bag of tricks for reward modeling in llms. *arXiv preprint
 640 arXiv:2410.18451*, 2024.
 641

642 Chris Yuhao Liu, Liang Zeng, Yuzhen Xiao, Jujie He, Jiacai Liu, Chaojie Wang, Rui Yan, Wei
 643 Shen, Fuxiang Zhang, Jiacheng Xu, Yang Liu, and Yahui Zhou. Skywork-reward-v2: Scaling
 644 preference data curation via human-ai synergy. *arXiv preprint arXiv:2507.01352*, 2025.
 645

646 Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios
 647 of the data distribution. In *Forty-first International Conference on Machine Learning*, 2024.
 648

649 Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Tang, Manan Roongta, Colin Cai,
 650 Jeffrey Luo, Tianjun Zhang, Erran Li, Raluca Ada Popa, and Ion Stoica. Deepscaler: Surpassing
 651 o1-preview with a 1.5b model by scaling rl, 2025. Notion Blog.
 652

648 Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing
 649 Ma, Qingwei Lin, and Dixin Jiang. Wizardcoder: Empowering code large language models with
 650 evol-instruct, 2023.

651

652 Konstanty Marczak. science_qa. Hugging Face Datasets, 2023. URL https://huggingface.co/datasets/KonstantyM/science_qa.

653

654 Arindam Mitra, Hamed Khanpour, Corby Rosset, and Ahmed Awadallah. Orca-math: Unlocking
 655 the potential of slms in grade school math, 2024.

656

657 mlfoundations dev. open2math-1m-gpt-4.1-mini. Hugging Face Datasets, 2025. URL <https://huggingface.co/datasets/mlfoundations-dev/open2math-1M-gpt-4.1-mini>.

658

659

660 Dhruv Nathawani, Shuoyang Ding, Vitaly Lavrukhin, Igor Gitman, Somshubra Majumdar,
 661 Evelina Bakhturina, Boris Ginsburg, and Jane Polak Scowcroft. Nemotron-Post-Training-
 662 Dataset-v2, August 2025. URL <https://huggingface.co/datasets/nvidia/Nemotron-Post-Training-Dataset-v2>.

663

664

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

668

669 Jingyang Ou, Shen Nie, Kaiwen Xue, Fengqi Zhu, Jiacheng Sun, Zhenguo Li, and Chongxuan
 670 Li. Your absorbing discrete diffusion secretly models the conditional distributions of clean data.
 671 *arXiv preprint arXiv:2406.03736*, 2024.

672

673 Robin L Plackett. The analysis of permutations. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 24(2):193–202, 1975.

674

675 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien
 676 Dirani, Julian Michael, and Samuel R. Bowman. Gpqa: A graduate-level google-proof q&a
 677 benchmark, 2023.

678

679 Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi
 680 Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, et al. Code llama: Open foundation models for
 681 code. *arXiv preprint arXiv:2308.12950*, 2023.

682

683 Subham Sekhar Sahoo, Zhihan Yang, Yash Akhauri, Johnna Liu, Deepansha Singh, Zhoujun Cheng,
 684 Zhengzhong Liu, Eric Xing, John Thickstun, and Arash Vahdat. Esoteric language models. *arXiv preprint arXiv:2506.01928*, 2025.

685

686 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 687 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical
 688 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

689

690 Yuxuan Song, Zheng Zhang, Cheng Luo, Pengyang Gao, Fan Xia, Hao Luo, Zheng Li, Yuehang
 691 Yang, Hongli Yu, Xingwei Qu, et al. Seed diffusion: A large-scale diffusion language model with
 692 high-speed inference. *arXiv preprint arXiv:2508.02193*, 2025.

693

694 Xiaohang Tang, Rares Dolga, Sangwoong Yoon, and Ilija Bogunovic. wd1: Weighted policy optimi-
 695 zation for reasoning in diffusion language models. *arXiv preprint arXiv:2507.08838*, 2025.

696

697 Gemma Team. Gemma. 2024. doi: 10.34740/KAGGLE/M/3301. URL <https://www.kaggle.com/m/3301>.

698

699 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Niko-
 700 lay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open founda-
 701 tion and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

702

703 Dimitri von Rütte, Janis Fluri, Yuhui Ding, Antonio Orvieto, Bernhard Schölkopf, and Thomas
 704 Hofmann. Generalized interpolating discrete diffusion. *arXiv preprint arXiv:2503.04482*, 2025.

702 Guanghan Wang, Yair Schiff, Subham Sekhar Sahoo, and Volodymyr Kuleshov. Remasking discrete
 703 diffusion models with inference-time scaling. *arXiv preprint arXiv:2503.00307*, 2025a.
 704

705 Xu Wang, Chenkai Xu, Yijie Jin, Jiachun Jin, Hao Zhang, and Zhijie Deng. Diffusion llms can do
 706 faster-than-ar inference via discrete diffusion forcing. *arXiv preprint arXiv:2508.09192*, 2025b.
 707

708 Yunhui Xia, Wei Shen, Yan Wang, Jason Klein Liu, Huifeng Sun, Siyue Wu, Jian Hu, and Xiaolong
 709 Xu. Leetcodedataset: A temporal dataset for robust evaluation and efficient training of code llms.
 710 *arXiv preprint arXiv:2504.14655*, 2025.

711 Tengyu Xu, Eryk Helenowski, Karthik Abinav Sankararaman, Di Jin, Kaiyan Peng, Eric Han, Shao-
 712 liang Nie, Chen Zhu, Hejia Zhang, Wenxuan Zhou, et al. The perfect blend: Redefining rlhf with
 713 mixture of judges. *arXiv preprint arXiv:2409.20370*, 2024.

714 Zhangchen Xu, Yang Liu, Yueqin Yin, Mingyuan Zhou, and Radha Poovendran. Kodcode: A
 715 diverse, challenging, and verifiable synthetic dataset for coding. *arXiv preprint arXiv:2503.02951*,
 716 2025.

717 Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng
 718 Kong. Dream 7b, 2025. URL <https://hkunlp.github.io/blog/2025/dream>.

719

720 Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhen-
 721 guo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions
 722 for large language models. *arXiv preprint arXiv:2309.12284*, 2023.

723

724 Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen.
 725 Mammoth: Building math generalist models through hybrid instruction tuning. *arXiv preprint
 726 arXiv:2309.05653*, 2023.

727

728 Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a ma-
 729 chine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association
 730 for Computational Linguistics*, 2019.

731 Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image
 732 diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*,
 733 pp. 3836–3847, 2023.

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

736

737 Lin Zheng, Jianbo Yuan, Lei Yu, and Lingpeng Kong. A reparameterized discrete diffusion model
 738 for text generation. *arXiv preprint arXiv:2302.05737*, 2023.

739

740 Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou,
 741 and Le Hou. Instruction-following evaluation for large language models, 2023. URL <https://arxiv.org/abs/2311.07911>.

742

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

746

747 A GENERATION PROCESS OF REMEDI

748

749 To better understand how RemeDi leverages the remasking mechanism, we visualize intermediate
 750 steps when solving math, code, and open-ended problems. Since the full responses are usually
 751 long, we focus on the token segment where critical remasking occurs. In the following figure, **green**
 752 boxes indicate already generated tokens, **blue** boxes represent tokens unmasked in this step, **red**
 753 boxes denote tokens remasked in this step, and **gray** boxes represent tokens that remain masked,
 754 showing the token with the highest probability. Key tokens are highlighted with bounding boxes.

755

We find that remasking enables diverse forms of revision beyond simple correction, including:

- **Correcting calculation errors:** Remask can correct calculation errors. As shown in Fig. 5, the model initially predicted “\div” as the most probable operator and generated “0”. However, since the actual operator generated was “\mod”, the model remasked the previous “0” and regenerated “5” as the correct result.
- **Refining text quality:** Remasking allows more precise wording. In Fig 6, the initially generated phrase “methyl group” is not adequate when answering this problem. RemeDi replaces them with more precise term “secondary carbon” by remasking.
- **Merging adjacent tokens:** When two consecutive tokens correspond to a single vocabulary token, RemeDi may remask them and merge into one, thereby freeing a slot for subsequent generation. In Fig. 7, the separate tokens “,” and “\” were remasked and merged into the single token “,\”, releasing one token slot.
- **Splitting tokens:** Conversely, the model can split a token into smaller parts to fill idle positions, ensuring that no tokens remain unused and the denoising process can complete. In Fig. 8, to fill in the slot before “Mish”, RemeDi remasks “Mish” and regenerates it as two tokens, “M” and “ish”.
- **Inserting tokens:** Remasking also supports insertion. In Fig. 9, to add the word “again” before “bounces”, the model first remasked the two tokens “b” and “ounces”, and then regenerated the sequence with the insertion.
- **Deleting tokens:** Finally, remasking can delete tokens and replace them with nothing or control symbols. In Fig. 10, the phrase “per hour →” was removed and replaced with a line break.

These cases illustrate that remasking gives RemeDi considerable freedom to revise its outputs in multiple ways, greatly extending the flexibility of diffusion-based text generation.

Moreover, we illustrate the predicted confidence scores in Fig. 11. In general, unmasked tokens receive higher confidence scores h_θ than masked tokens, unless the model judges them as unreliable and decides to remask them. For example, see the tokens “say” in Fig. 11b and “in” in Fig. 11c. Interestingly, these tokens that are eventually remasked already exhibit relatively low confidence at the step when they were first predicted, as reflected by the lighter background green shading. This suggests that the model was uncertain about them from the start. Once more context is revealed in subsequent steps, RemeDi is able to revise such low-confidence tokens into more appropriate alternatives.

B EXPERIMENT DETAILS

B.1 DUAL-STREAM MODEL STRUCTURE

We construct TPS with the same transformer structure as LLaDA (Nie et al., 2025), comprising 32 transformer blocks. The model weights in this stream are initialized with LLaDA-8B-Instruct. For the UPS, we stack four transformer blocks with the same hidden dimension to construct a smaller network with random initialization. These two streams are weakly coupled via bi-directional connections at TPS blocks 1, 11, 21, and 31. At each connection point, the output from the previous TPS block is added to the current UPS feature to form the input for the next UPS block, while the output of the current UPS block is added to the output of the corresponding TPS block before it is passed onward. Thus, both streams enrich their representation with features from each other. To preserve the TPS’s original capability inherited from the pretrained weights, we add a zero-initialized projection (Zhang et al., 2023) on the connections from UPS to TPS. This ensures that the model’s token prediction behavior is unchanged at the beginning of training, and gradually learns to predict the confidence for unmasking tokens in a diffusion process. The model comprises a total of 8.9B parameters.

B.2 TRAINING CONFIGURATIONS AND DATASETS

B.2.1 DATASETS

We use both high quality public datasets and in-house data for training, including four major categories: mathematics, code, general conversation, and science. The mathematics category includes



834 Figure 5: An example of correcting calculation errors with remasking. Question: *A group of N students, where $N < 50$, is on a field trip. If their teacher puts them in groups of 8, the last group has 5 students. If their teacher instead puts them in groups of 6, the last group has 3 students. What is the sum of all possible values of N ?*



861 Figure 6: An example of refining text quality with remasking. Question: *What is the major outcome
862 of the reaction between 4,4-dimethylcyclopent-1-enol and bromine?*

Figure 7: An example of merging adjacent tokens. Question: *A quantum mechanical particle of mass m moves in two dimensions in the following potential, as a function of $(r, \theta) : V(r, \theta) = 1/2kr^2 + 3/2kr^2 \cos^2(\theta)$. Find the energy spectrum.*

	...	Let	's	calculate	the	total	cost	for	each	item	:	\n
step 25	**	Step		1	:	Calculate	the	cost	for	shorts	**	\n
	-	Mish	ka	bought		3	pairs	...				
step 26	**	Step		1	:	Calculate	the	cost	for	shorts	**	\n
	M	ish	ka	bought		3	pairs	...				
step 27	**	Step		1	:	Calculate	the	cost	for	shorts	**	\n
	M	ish	ka	bought		3	pairs	...				

Figure 8: An example of splitting tokens. Question: *Mishka bought 3 pairs of shorts, 3 pairs of pants, and 3 pairs of shoes. One pair of shorts costs \$16.50. One pair of pants costs \$22.50 and one pair of shoes costs \$42. How many dollars did Mishka spend on all the clothing items?*





Figure 10: An example of deleting tokens. Question: *When Billy was first hired, he was paid at a rate of \$10 per hour. After 2 months, he was given a raise of \$0.50 per hour. On his first anniversary at work, he was given a raise of \$1.00 per hour. Sally just started working at a different business, and her starting salary is \$0.50 more per hour than Billy's starting salary was. If both Billy and Sally work 20 hours, how much more money will Billy earn than Sally, in dollars?*

NuminaMath-CoT (LI et al., 2024), MetaMathQA (Yu et al., 2023), orca-math-word-problems-200k (Mitra et al., 2024), Maths-College (ajibawa 2023), DeepMath-103K (He et al., 2025), MathInstruct (Yue et al., 2023), and open2math-1M-gpt-4.1-mini (mlfoundations dev, 2025). The code category comprises evol-codealpaca-v1 (Luo et al., 2023), opc-sft-stage1 (Huang et al., 2024), and KodCode-V1-SFT-40 (Xu et al., 2025). General conversation data is drawn from open-perfectblend (Xu et al., 2024) and Infinity-Instruct (Li et al., 2025). The science category incorporates the science_qa (Marczak, 2023) dataset. In total, the public datasets provide roughly 18.8M samples. The in-house data, containing about 140K samples of prompt-response pairs, was primarily generated by GPT-4.1 and covers mathematics, code, science, and instruction-following tasks. These generated samples went through a human check to ensure overall quality.

B.2.2 VARIABLE-LENGTH BLOCK-WISE GENERATION

We first adapt LLaDA into a DLM capable of variable-length generation. The underlying architecture remains unchanged, but we fine-tune it for variable-length block-wise generation. During inference, generation proceeds block by block, where each block consists of $L = 32$ tokens. For each block, the model runs a full reverse diffusion process until the block is fully denoised, after which the completed block is appended to the context. The next block is then generated in the same manner, continuing until an `<eos>` token appears. Similar to auto-regressive (AR) models, we enforce causality using a block-wise causal mask in the self-attention layers: each token can attend to all tokens within its current block and all tokens in previously generated blocks, but never to tokens in future blocks.

For supervised finetuning, the response part is divided into blocks of length $L = 32$, with the last incomplete block padded by `<eos>` tokens. The training objective is to recover a noised version of each block conditioned on all previous clean blocks. Following Arriola et al. (2025), we concatenate both the clean blocks and their corresponding noised versions into a single input sequence, allowing all blocks to be trained jointly in one forward-backward pass.

In this stage, we train variable-length block-wise generation with learning rate 2×10^{-6} , batch size 160, and a gradient threshold of 1.0. The baseline model in Table 4 reports the results of the variable-length block-wise model on several datasets, alongside the comparisons when Remask SFT and Remask RL added to this baseline model.

1026
 1027
 1028
 1029
 1030
 1031
 1032
 1033 **step 80** ... Let C_J be the number of chairs Jenna has.
 1034 fewer22 than...
 1035 **step 81** ... Let C_J be the number of chairs Jenna has.
 1036 fewer2 than...
 1037 **step 85** ... Let C_J be the number of chairs Jenna has.
 1038 They knowia has fewer fewer than than...
 1039 **step 87** ... Let C_J be the number of chairs Jenna has.
 1040 They each have 2 fewer sofas than chairs:...
 1041 (a) Question from GSM8K: *Ophelia and Jenna are living in the same apartment building. They each have 2 fewer sofas than chairs. Jenna has 3 times as many chairs as Ophelia. If Ophelia has 20 sofas, calculate the total number of sofas and chairs that they have.*
 1042
 1043 **step 389** ...
 1044 **step 402** ...
 1045 **step 408** ...
 1046 **step 411** ...
 1047 (b) Question from Alpaca-Eval: *Hi, I need to find the area of a 20-sided die for my math homework. Can you help me do that?*
 1048
 1049
 1050 **step 100** ... Some individuals may feel comfortable and valued in platonic relationships, while others may
 1051 feel pressure or attraction in romantic romantic...
 1052 friendship friendship...
 1053 **step 108** ... Some individuals may feel comfortable and valued in platonic relationships, while others may
 1054 feel pressure or attraction in romantic situations.
 1055 **step 112** ... Some individuals may feel comfortable and valued in platonic relationships, while others may
 1056 feel pressure or attraction towards romantic relationships.
 1057 **step 115** ... Some individuals may feel comfortable and valued in platonic relationships, while others may
 1058 feel pressure or attraction towards romantic relationships.
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066 (c) Question from Alpaca-Eval: *Can a boy and girl ever just be best friends?*
 1067
 1068
 1069 Figure 11: Visualization of per-token confidence scores predicted by UPS. The darkness of the
 1070 background color indicates the value of h_θ^i — darker means larger. The font color indicates different
 1071 type of tokens: **Green** and **red** tokens are already unmasked in the input sequence of the current
 1072 step. **Red** tokens are remasked. **Blue** tokens are newly unmasked in the outputs. **Gray** tokens keep
 1073 masked in the outputs, and we display the token with the highest probability at these positions.
 1074
 1075
 1076
 1077
 1078
 1079

1080
1081

B.2.3 REMASK SFT

1082
1083
1084
1085
1086

After finetuning LLaDA for variable-length block-wise generation, we attach the Unmasking Policy Stream (UPS) to construct the model architecture shown in Fig. 2. We then further train RemeDi with Remask SFT with $\lambda_{UPS} = 0.3$. For optimization, we apply a learning rate of 2.0×10^{-5} to the newly introduced parameters in UPS, while keeping the learning rate for the original parameters at 2.0×10^{-6} .

1087

1088
1089
1090
1091

To enable effective RL training, we curated a dataset spanning mathematics, coding, STEM, instruction-following and preference-alignment tasks:

1092
1093
1094

- **Math:** GSM8K(Cobbe et al., 2021), MATH(Hendrycks et al., 2021), DeepScaleR(Luo et al., 2025)
- **Code:** KodCode-V1-SFT-R1(Xu et al., 2025), LeetCodeDataset(Xia et al., 2025)
- **General:** Skywork-Reward-Preference-80K-v0.2(Liu et al., 2024), Llama-Nemotron-Post-Training-Dataset-RL (instruct-following)(Bercovich et al., 2025), Nemotron-Post-Training-Dataset-v2 (stem)(Nathawani et al., 2025)

1095
1096
1097
1098

To ensure quality, we applied the following filters:

1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109

- **Length:** Since our RL training limits generation length to 1024 tokens, we discard any sample—question plus answer or question alone—exceeding this bound.
- **Verifiable:** For math data, When both a short answer and a detailed response are available, we keep the sample only if the two answers match; for code data, we require that the provided solution passes all test cases.
- **Deduplication:** Given the diverse sources, we perform global deduplication using Min-HashLSH .

1110
1111

Reward Design. Our reward function incorporates two distinct types of reward signals:

1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122

- **Verifiable Reward:** Verifiable rewards are widely used in mathematics, code, and STEM domains, where the answer is first extracted and then verifies it: math and STEM via Math-Verify (Kydlíček, 2025), and code via executing test cases and computing the pass rate. We also incorporated verifiable instruction-following samples with IFEval (Zhou et al., 2023) format to further improve the model’s ability to follow instructions.
- **Model-based Reward:** We incorporated the Skywork-Reward-V2-Llama-3.1-8B (Liu et al., 2025), which was trained on human preference data, to evaluate response quality. Each response is assigned a scalar reward score reflecting human preference, thereby enhancing the model’s ability to produce outputs that better align with human preference during RL training.

1123
1124
1125

We optimized the model using the AdamW optimizer with a learning rate of 5.0×10^{-6} , $\beta = (0.9, 0.999)$ and a maximum gradient norm of 1.0, for a total of 100 training steps.

1126
1127

B.2.5 ABLATION SETUP

1128
1129

We provide detailed configurations for ablation studies in Sec. 4.3.

1130
1131
1132
1133

Remask SFT Both the vanilla SFT model and the Remask SFT model are trained with identical hyper-parameters: block size 32, maximum generation length 1024, global batch size 80, and a total of 3000 training steps. We use the AdamW optimizer with a learning rate of 2.0×10^{-5} in newly added parameters of UPS and of 2.0×10^{-6} in original TPS parameters, $\beta = (0.9, 0.999)$, and a maximum gradient norm of 1.

1134 **Algorithm 2** Remask RL

1135 **Require:** Model parameters θ , a dataset \mathcal{D} , and reward_func.

1136 1: **while** θ not converged and maximum epochs not reached **do**

1137 2: Sample questions $q \sim \mathcal{D}$

1138 3: **for** $g = 1$ to G **do** ▷ Generate a group of G trajectories

1139 4: Initialize $x_{t_0}^g$ with q and mask tokens.

1140 5: **for** $n = 1$ to N **do** ▷ N denotes the number of denoising steps

1141 6: Calculate the ranking score $h_{\theta,n}$ for each token

1142 7: Sample K_n positions to unmask in this step: $\mathcal{U}_n \sim \text{Plackett-Luce}(h_{\theta,n}, K_n)$

1143 8: Sample $x_{t_n}^{g,i} \sim p_{\theta,n}^i(\cdot | x_{t_{n-1}}^g)$, $\forall i \in \mathcal{U}_n, x_{t_{n-1}}^{g,i} = [\text{M}]$

1144 9: **end for**

1145 10: $r^g = \text{reward_func}(q, x_{t_N}^g)$ ▷ Compute the rewards

1146 11: **end for**

1147 12: **for** $g = 1$ to G **do** ▷ Compute the advantages as in GRPO

1148 13: $A^g = \frac{r^g - \text{mean}(r^{1:G})}{\text{std}(r^{1:G})}$

1149 14: **end for**

1150 15: **for** $n = 1$ to N **do** ▷ Compute π_θ and losses for each denoising step

1151 16: $\pi_{\theta,n}(x_{t_n}^g | x_{t_{n-1}}^g) = \pi_{\theta,n}^{\text{unmask}}(\mathcal{U}_n^g | x_{t_{n-1}}^g) \cdot \pi_{\theta,n}^{\text{token}}(x_{t_n}^g | x_{t_{n-1}}^g, \mathcal{U}_n^g)$ ▷ see Eq. 8

1152 17: $\mathcal{L}_{\theta,n} = -\frac{1}{G} \sum_{g=1}^G \frac{\pi_{\theta,n}(x_{t_n}^g | x_{t_{n-1}}^g)}{\pi_{\text{old},n}(x_{t_n}^g | x_{t_{n-1}}^g)} A^g$

1153 18: Calculate the gradient $\nabla_{\theta} \mathcal{L}_{\theta,n}$

1154 19: **end for**

1155 20: Update θ with accumulated gradients $\sum_{n=1}^N \nabla_{\theta} \mathcal{L}_{\theta,n}$ along the descent direction

21: **end while**

1156

1157

1158 **Remask RL** Both Remask RL and LLaDOU RL are trained on LLaDA using identical RL hyper-parameters. Specifically, roll-outs are generated with temperature 0.7, generation length 256, block length 64, and $N = 64$ denoising steps. Each batch consists of 16 prompts, with each prompt generating $G = 16$ roll-outs, for a total of 200 training steps. We use the AdamW optimizer with a learning rate of 5.0×10^{-6} , $\beta = (0.9, 0.999)$, and a maximum gradient norm of 1. The complete training algorithm is elaborated in Alg. 2.

1163

1164 B.3 EVALUATION DETAILS

1165 B.3.1 INFERENCE SETTINGS

1166 **RemeDi** For evaluation, RemeDi uses a maximum generation length of 2048 on MATH and 1024 on all other datasets. At each step, only one token is unmasked, with a block size of 32 for generation. Both TPS and UPS adopt greedy sampling.

1171

1172 **LLaDA** The evaluation of LLaDA largely follows (Nie et al., 2025). On GSM8K and MATH, we set the generation length to 256 with a block length of 8, unmasking one token per step in a semi-autoregressive manner with greedy sampling. On HumanEval and MBPP, we use a generation length of 512 with a block length of 32, while keeping all other settings unchanged.

1176

1177 **LLaDA + ReMDM** We implemented ReMDM (Wang et al., 2025a) on top of LLaDA(Nie et al., 2025). Specifically, we adopted the “ReMDM-cap + Switch” configuration with $\eta_{\text{cap}} = 0.04$ and $t_{\text{switch}} = 0.55$. For evaluation, we set the generation length to 256/256/512/512 and the block length to 8/8/32/32 for GSM8K, MATH, HumanEval, and MBPP, respectively.

1181

1182 B.3.2 BENCHMARKS

1183 Here we provide the detailed input prompts and how the metrics are computed for different benchmarks:

1185

1186 **GSM8K** GSM8K evaluates multi-step mathematical reasoning in elementary problems (Cobbe et al., 2021). We illustrate below a zero-shot prompt used to evaluate the model. After generation, we

1188 extract the answer in “boxed{}” from the response, and check if it is equivalent to the ground truth
 1189 with the scripts developed by Hendrycks et al. (2021). We report the accuracy on this benchmark.
 1190

1191 Janet’s ducks lay 16 eggs per day. She eats three for breakfast every
 1192 morning and bakes muffins for her friends every day with four. She sells
 1193 the remainder at the farmers’ market daily for \$2 per fresh duck egg. How
 1194 much in dollars does she make every day at the farmers’ market? (Please
 1195 put the final answer in \boxed{} tag, i.e. \\$\boxed{answer here}\\$)
 1196

1197 **MATH** MATH contains 5,000 challenging competition mathematics problems (Hendrycks et al.,
 1198 2021). We evaluate the model in a zero-shot setting, with prompts like the one below. After gener-
 1199 ation, we extract the answer in “boxed{}” from the response, and verify if it is equivalent to the
 1200 ground truth with the scripts developed by Hendrycks et al. (2021).
 1201

1202 Convert the point \$(0,3)\$ in rectangular coordinates to polar coordinates.
 1203 Enter your answer in the form \$(r, \theta)\$, where \$r > 0\$ and \$0 \leq \theta < 2 \pi\$. (Please put the final answer in \boxed{} tag, i.e.
 1204 \\$\boxed{answer here}\\$)
 1205

1206 **GPQA** GPQA is a challenging multiple-choice benchmark for testing LLM’s complex scientific
 1207 reasoning and specialized knowledge domains (Rein et al., 2023). We used all 448 questions from
 1208 the main version of GPQA and evaluated the model in a zero-shot setting with the prompt shown
 1209 below. We select the token with the highest probability at the <mdm_mask> position as the final
 1210 answer, and report the pass@1 on this benchmark.
 1211

1212 <|startoftext|><|start_header_id|>user<|end_header_id|>
 1213 What is the correct answer to this question: A large gene has dozens of
 1214 exons, of which the central ones code for folded triple helical repeats
 1215 that connect the cytoskeleton with sarcolemma and extracellular space.
 1216 Each exon usually codes for one folded triple alpha helix. The most common
 1217 mutations of the gene are central exon deletions that create out-of-frame
 1218 peptides and progressive degenerative organ waste. A solution is to deliver
 1219 a Morpholino that recognizes the 5' end of the out-of-frame exon in pre-
 1220 mRNA. The molecule prevents binding of the spliceosome and creates exon
 1221 skipping and in-frame joining. Several missing exons are well tolerated by
 1222 an organism. Which structure below is not involved in the proposed therapy?
 1223 Choices:
 1224 (A) lariat
 1225 (B) R-loops
 1226 (C) antisense
 1227 (D) polyA tail
 1228 Answer: Your answer should be in the format 'The best answer is
 1229 [the_answer_letter]' where the [the_answer_letter] is one of (A), (B), (C)
 1230 or (D).<|eot_id|><|start_header_id|>assistant<|end_header_id|>

1231 The best answer is <mdm_mask>.
 1232

1233 **MBPP** MBPP consists of 500 python programming problems for entry level programmers (Austin
 1234 et al., 2021). We evaluate the model with the prompt below in a zero-shot setting. After generation,
 1235 we extract the python code from the response, and check if it passes all test cases associated with
 1236 this problem, and report pass@1 on this benchmark.
 1237

1238 You are an expert Python programmer. Your task is to complete the
 1239 implementation of a function named `remove_Occ`.
 1240

1241 **** TARGET FUNCTION ****
 1242

1243 Write a python function to remove first and last occurrence of a given
 1244 character from the string.
 1245

1246 **** UNIT TESTS ****

```

1242 Your code should pass unit tests like:
1243     assert remove_Occ("hello", "l") == "heo"
1244     assert remove_Occ("abcda", "a") == "bcd"
1245     assert remove_Occ("PHP", "P") == "H"
1246
1247 Here is the function to complete:
1248 ```python
1249 def remove_Occ(input_param_1, input_param_2):
1250     """Write a python function to remove first and last occurrence of a
1251     given character from the string."""
1252     ...
1253
1254 HumanEval HumanEval consists of 164 hand-written programming problems (Chen et al., 2021).
1255 We evaluate the model on it with the prompt below in a zero-shot setting. After generation, we
1256 extract the Python code from the response, and check whether the output function passes all the
1257 provided test cases; we then report the pass@1 on this benchmark.
1258
1259 You are an expert Python programmer, Python code should be placed between
1260 the line of ```python and the line of ``` for easy extraction later, and
1261 here is your task:
1262 ```python
1263 from typing import List
1264 def has_close_elements(numbers: List[float], threshold: float) -> bool:
1265     """Check if in given list of numbers, any two numbers are closer to
1266     each
1267     other than the given threshold.
1268
1269     Examples:
1270     >>> has_close_elements([1.0, 2.0, 3.0], 0.5)
1271     False
1272     >>> has_close_elements([1.0, 2.8, 3.0, 4.0, 5.0, 2.0], 0.3)
1273     True
1274     """
1275
1276 Hellaswag Hellaswag is a benchmark dataset specifically designed to evaluate machine common-
1277 sense reasoning capabilities (Zellers et al., 2019). It primarily assesses a model's ability to infer the
1278 most plausible subsequent event based on given contextual information in natural language under-
1279 standing tasks. We evaluate the model in a zero-shot setting on Hellaswag. We follow the approach
1280 of (Nie et al., 2025), incorporating Classifier-Free Guidance (CFG) and set the CFG weight to 0.5.
1281 Under CFG intervention, the model simultaneously computes conditional predictions (based on the
1282 given context) and unconditional predictions (absent specific context), guiding the generation pro-
1283 cess by scaling the difference between them. The final accuracy is calculated based on the model's
1284 normalized probability assigned to the correct option.
1285
1286 ARC-C ARC-C is a highly challenging benchmark dataset specifically designed to evaluate ma-
1287 chine abstract reasoning and scientific problem-solving capabilities (Clark et al., 2018). We evaluate
1288 the model in a zero-shot setting on ARC-C, with the prompt below. After generation, we extract the
1289 answer after 'The best answer is', and report the pass@1 rate on this benchmark.
1290
1291 Given the following question and four candidate answers (A, B, C and
1292 D), choose the best answer.
1293 Question: An astronomer observes that a planet rotates faster after a
1294 meteorite impact. Which is the most likely effect
1295 of this increase in rotation?
1296 A. Planetary density will decrease.
1297 B. Planetary years will become longer.
1298 C. Planetary days will become shorter.
1299 D. Planetary gravity will become stronger.
1300 Your response should end with "The best answer is [the_answer_letter]"
```

1296 where the [the_answer_letter] is one of A, B, C or D.
 1297

1298 **IFEval** IFEval evaluates the model’s instruction-following capability with verifiable instructions
 1299 (Zhou et al., 2023). We use the official evaluation code provided by the IFEval Benchmark, and
 1300 compute the model’s accuracy based on the loose metric.
 1301

1302 **AlpacaEval** AlpacaEval evaluates the model’s instruction-following capability with the LLM-
 1303 as-a-Judge methodology (Dubois et al., 2024). As officially recommended by the AlpacaE-
 1304 val benchmark, we use GPT4-1106-preview as the baseline/reference model and the
 1305 weighted_alpaca_eval_gpt4_turbo as the evaluator/annotator, and assess the win rate of
 1306 the responses generated by RemeDi under length-controlled conditions to eliminate the confounding
 1307 effect of response length.
 1308

1309 C MORE EXPERIMENTS

1311 C.1 COMPARISON WITH SEED DIFFUSION

1312 Since the official Seed Diffusion model and implementation are not publicly available, we do our
 1313 best to reproduce it. For a fair comparison, we train seed diffusion under the same base model and
 1314 identical training configuration as in Appendix B.2.5. As shown in Table 6, Remask SFT consis-
 1315 tently outperforms Seed Diffusion, demonstrating the advantage of learning an explicit remasking
 1316 policy during training.
 1317

1318 Table 6: Unified head-to-head comparison with other training algorithms under identical settings.
 1319

1320 Method	1321 GSM8K	1322 MATH-500	1323 HumanEval	1324 MBPP
1322 Baseline	1323 80.3	1324 34.7	41.5	42.6
1323 Vanilla SFT	1324 83.1	40.1	48.2	43.4
1324 Seed Diffusion	1325 63.9	28.0	5.4	9.8
1325 Remask SFT	1326 83.6	1327 42.7	1328 50.0	1329 44.0

1327 C.2 PREDICTOR–CORRECTOR VS. LEARNED REMASK POLICY

1328 We apply a representative predictor-corrector sampler, ReMDM (Wang et al., 2025a), to RemeDi-
 1329 Instruct. The evaluation setup is the same as in Appendix B.3.1. The results in Table 7 show that our
 1330 learned remasking policy (via Remask SFT) is more effective than the random remasking strategy
 1331 employed in predictor-corrector samplers.
 1332

1333 Table 7: Comparison between our learned remask policy in RemeDi and the ReMDM predictor-
 1334 corrector, both evaluated with RemeDi-Instruct.
 1335

1336 Method	1337 GSM8K	1338 MATH-500	1339 HumanEval	1340 MBPP
1338 RemeDi + predictor-corrector	1339 58.3	38.7	39.6	54.2
1339 RemeDi (Ours)	1340 86.3	1341 52.2	1342 71.3	1343 57.8

1341 C.3 EFFECT OF DIFFERENT SAMPLERS

1342 To isolate the effect of our multi-task objective and incorrect-token augmentation, we compared
 1343 vanilla sampler, adaptive sampler(Kim et al., 2025) and our remask sampler under the same Remask
 1344 SFT model in Table 8.
 1345

1346 C.4 MATCHED-COMPUTE ABLATION: EXTRA SFT VS. REMASK RL

1347 To directly address whether performance gains come from RL or merely additional training time,
 1348 we conduct a matched-compute ablation: we continue training the RemeDi-Instruct model for an
 1349

1350
1351
1352 Table 8: Effect of different samplers under the same Remask SFT model.
1353
1354
1355
1356

Sampler	GSM8K	MATH-500	HumanEval	MBPP	IFEval
Vanilla	86.3	38.3	43.9	55.4	69.2
Adaptive(Kim et al., 2025)	86.6	40.3	43.9	55.4	69.2
Remask (Ours)	86.3	52.2	71.3	57.8	81.9

1357
1358 extra 2,000 steps (which consumes approximately 32 H800-days, the same compute as used in the
1359 RL stage) and evaluate its performance. The results are summarized in Table 9:
13601361 Table 9: Matched-compute ablation between extra SFT training and Remask RL.
1362

	GSM8K	MATH-500	HumanEval	MBPP
RemeDi-Instruct	86.3	52.2	71.3	57.8
+ ~32 H800-days SFT training	83.6	52.6	62.8	57.8
+ ~32 H800-days RL training	89.1	53.2	73.2	59.4

1363
1364 C.5 UPS STRUCTURE ABLATIONS
1365
13661367
1368 Table 10 reports ablations on UPS components, showing that removing either the bi-residual
1369 connections or the zero-init bridge leads to clear performance degradation. We train all models under
1370 the same base model and identical training configuration as in Appendix B.2.5
1371
13721373 Table 10: UPS structure ablations. Removing either the bi-residual connections or the zero-init
1374 bridge degrades performance
1375


	GSM8K	MATH-500	HumanEval	MBPP
Baseline	83.6	42.7	50.0	44.0
w/o bi-residuals	76.6	43.1	45.7	42.2
w/o zero-init	75.2	42.7	45.1	44.8

1381
1382 C.6 EFFECTIVENESS OF RL
13831384 To assess the effect of RL, we compare the per-token remask frequencies between the SFT-only
1385 model and the RL-trained model. Interestingly, the RL-trained model performs **more** remasking on
1386 average. As shown in Table 11, higher remask frequency correlates with improved performance,
1387 suggesting that additional remasking provides more opportunities for RemeDi to detect and correct
1388 early-step errors.
13891390 D EXTRA VISUALIZATION
13911392 Fig. 12 compares the per-token remask frequency of the SFT-only model and the RL-trained model
1393 on GSM8K. We observe that RL training consistently increases the remask frequency, suggesting
1394 that the RL objective explicitly encourages the model to revise uncertain tokens more often, which
1395 correlates with the observed improvement in final answer quality.
13961397 Fig. 13 reports the remask ratios at different timesteps when inferencing on GSM8K. The ratio
1398 first increases and then steadily decreases as the denoising process converges. Since the number of
1399 unmasked tokens K_n is explicitly controlled at each step, the process naturally terminates without
1400 spikes of remasking in late stage.1401 Fig. 14 shows the throughput–performance trade-off of RemeDi compared to AR and DLM base-
1402 lines. By adjusting the number of tokens denoised per step, RemeDi achieves acceleration in decod-
1403 ing speed with only small drops in GSM8K accuracy. It forms a better Pareto front than LLaDA,
1404 Dream, and other auto-regressive models.

1404 Table 11: Average remask frequency (ARF) and performance across tasks. ARF measures how
 1405 many times each token is remasked on average during decoding.

1406

1407	Model	GSM8K		HumanEval		AlpacaEval	
		1408 Acc	1408 ARF	1409 Acc	1409 ARF	1410 Acc	1410 ARF
1411	RemeDi (+ Remask SFT)	86.3	0.16	71.3	0.070	12.5	0.012
1412	RemeDi (++ Remask RL)	89.1	0.56	73.2	0.89	24.8	0.086



1422 Figure 12: Comparison of per-token remask frequency between the SFT-only (left) and the RL-
 1423 trained (right) model on GSM8K. The RL-trained model performs remasking more frequently on
 1424 average, indicating that RL encourages more remasking behavior.

1425

1426 Fig. 15 shows the reward curve for Remask-RL and LLaDOU-RL on GSM8K. Due to the larger
 1427 action space introduced by the remask operation, Remask-RL starts with a lower initial reward.
 1428 However, it quickly surpasses LLaDOU-RL within the early stages of training and maintains a
 1429 consistently higher reward thereafter. This trend is consistent with the accuracy comparison in Table
 1430 5, where Remask-RL demonstrates both faster convergence and a higher final performance.

1431

1432

1433 THE USE OF LARGE LANGUAGE MODELS

1434

1435 We used large language models (LLMs) in two limited ways: (1) for minor English polishing of the
 1436 paper text, and (2) as required by the AlpacaEval benchmark, where LLMs are invoked for automatic
 1437 evaluation.

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

Figure 13: Average remask ratio across diffusion timesteps when inferencing on GSM8K. The remask ratio rises during early steps and then gradually declines as the process converges. Because the number of unmasked tokens K_n is explicitly defined at every diffusion step, the procedure ensures stable remask termination and avoids late-stage re-emergence of remasking.

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511




Figure 13: Average remask ratio across diffusion timesteps when inferencing on GSM8K. The remask ratio rises during early steps and then gradually declines as the process converges. Because the number of unmasked tokens K_n is explicitly defined at every diffusion step, the procedure ensures stable remask termination and avoids late-stage re-emergence of remasking.




Figure 14: Throughput–performance trade-off of RemeDi compared with other AR and DLM models. By increasing the number of denoised tokens per step, RemeDi provides a smooth quality–latency trade-off. All results are measured with batch size 1 and sequence length 1024 on a single H800 GPU.

1512

1513

1514

1515

1516

1517

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542


1543

1544

1545

1546

Figure 15: GSM8K training reward curves comparing Remask-RL and LLaDOU-RL. Solid lines show EMA-smoothed rewards, and faint lines denote raw step-wise values. Remask-RL exhibits faster early-stage improvement and ultimately reaches a higher reward, consistent with the performance gains summarized in Table 5



1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

1565