
Reinforcing the Diffusion Chain of Lateral Thought with Diffusion Language Models

Zemin Huang^{1,2,*} Zhiyang Chen^{2,4,*} Zijun Wang^{1,2} Tiancheng Li^{1,2} Guo-Jun Qi^{2,3,4,†}

¹Zhejiang Univeristy

²MAPLE Lab, Westlake University

³Matterwave Intelligence

⁴Institute of Advanced Technology, Westlake Institute for Advanced Study

{huangzemin, chenzyang, wangzijun63, litiancheng}@westlake.edu.cn, guojunq@gmail.com

<https://github.com/maple-research-lab/LLaDOU>

Abstract

We introduce the *Diffusion Chain of Lateral Thought (DCoLT)*, a reasoning framework for diffusion language models. DCoLT treats each intermediate step in the reverse diffusion process as a latent "thinking" action and optimizes the entire reasoning trajectory to maximize the reward on the correctness of the final answer with outcome-based Reinforcement Learning (RL). Unlike traditional Chain-of-Thought (CoT) methods that follow a causal, linear thinking process, DCoLT allows bidirectional, non-linear reasoning with no strict rule on grammatical correctness amid its intermediate steps of thought. We implement DCoLT on two representative Diffusion Language Models (DLMs). First, we choose SEDD as a representative continuous-time discrete diffusion model, where its concrete score derives a probabilistic policy to maximize the RL reward over the entire sequence of intermediate diffusion steps. We further consider the discrete-time masked diffusion language model – LLaDA, and find that the order to predict and unmask tokens plays an essential role to optimize its RL action resulting from the ranking-based Unmasking Policy Module (UPM) defined by the Plackett-Luce model. Experiments on both math and code generation tasks show that using only public data and 16 H800 GPUs, DCoLT-reinforced DLMs outperform other DLMs trained by SFT or RL or even both. Notably, DCoLT-reinforced LLaDA boosts its reasoning accuracy by +9.8%, +5.7%, +11.4%, +19.5% on GSM8K, MATH, MBPP, and HumanEval.

1 Introduction

To enable complex reasoning, most large language models (LLMs) [15, 20] learn to decompose problems into simpler sub-steps and generate intermediate reasoning in natural language. Chain-of-Thought (CoT) [40] first reveals that step-by-step reasoning facilitates language models, as the outputs from previous steps could be rationales for more accurate next step prediction. Based on that, OpenAI’s PRM [23] supervises these intermediate reasoning steps with progressive rewards, to ensure the correctness of each single step. More recently, DeepSeek-R1 [15] eliminates the need of the reward model and verifies only the correctness of the final answer, relaxing the constraints on

*Equal contribution.

†Corresponding author.

This project was initiated and supported by MAPLE Lab at Westlake University.

only public data and 16 H800 GPUs, the DCoLT-reinforced LLaDA model achieves 88.1% on GSM8K [8], 44.6% on MATH [16], 51.6% on MBPP [4] and 59.1% on HumanEval [7] in the challenging zero-shot setting. Even compared with auto-regressive models that are trained with significantly more proprietary data and fully annotated CoT reasoning processes, it still demonstrates competitive performances.

2 Related Works

Diffusion Language Models Diffusion models have achieved impressive results in image modeling [18, 21, 33, 38, 44], with recent efforts [2, 3, 11–13, 24, 27, 29, 37, 43, 48, 49] extending to language tasks. Depending on the types of diffused distributions over text tokens, diffusion language models can be categorized into continuous diffusion models [13] and discrete diffusion models [3, 11, 24], with our work focusing primarily on the latter due to the discrete nature of language.

Among discrete diffusion models, masked diffusion models [24, 27, 43] emerge as a promising approach due to their superior performances. Recent advances simplify model training and design by adopting the straight cross-entropy loss [29, 37, 49] and removing the time embeddings [48]. The resultant models [27, 43] achieve the state-of-the-art performance among existing diffusion language models.

Particularly, we consider two representative examples of discrete diffusion language models in this paper: SEDD [24] – a discrete diffusion model with a continuous diffusion time, and LLaDA [27] – a masked diffusion model with discrete diffusion timesteps. We will reinforce the entire sequence of their reverse diffusion processes containing non-linear text generations to release their lateral thinking ability. In contrast, DoT (Diffusion of Thought) [42] uses annotated step-by-step CoT data for *supervised fine-tuning* with existing diffusion losses such as the score entropy loss [24] or the noise prediction loss [13]. Thus, it still encourages a conventional stepwise reasoning process as in *vertical thinking* [17].

Reinforcement Learning for Language Models Reinforcement Learning (RL) helps language models better align with human preference [30] or verifiable knowledge [22] (e.g. rewards on math/code tasks). Recently, there appear various algorithms for RL optimization [1, 25, 35, 36]. These methods enhance model’s ability to produce aligned outputs and show strong potential in inference-time scaling for challenging problems [15, 20].

However, existing RL approaches are primarily applied to auto-regressive language models [15, 30]. For DLMS, previous works explored policy gradient methods within the concrete score-matching framework [46], or estimated the log-probability of the sampled responses to perform reinforcement learning [47]. However, both models merely focused on the generated text response in the final diffusion step as an action to optimize, ignoring the role of intermediate diffusion steps as lateral thought to reinforce amid the reverse diffusion process.

Instead, in this paper, we seek to explicitly reinforce the reverse diffusion process to enable the training of lateral reasoning. We demonstrate in experiments that even though these intermediate diffusion steps perform non-linear rather than step-by-step reasoning, reinforcing them can eventually lead to correct answers at the final step that outperforms many CoT models [10, 40].

3 Methods

3.1 Diffusion Chain of Lateral Thought

To model the probability distribution p_{data} over a finite vocabulary $\mathcal{V} = \{1, \dots, V\}$ for text generation tasks, discrete diffusion processes model how the unknown data distribution p_{data} at $t = 0$ gradually evolves into a prior distribution p_{prior} at $t = T$ [6]. The distribution at an intermediate diffusion time t is denoted by p_t .

Then a generation process is realized by reversing this diffusion process. Practically, it begins by sampling x_0 from the prior p_{prior} and then iteratively denoising through a sequence of diffusion time $t_{0:N}$. At each step n , the model θ estimates the diffused distribution p_{θ, t_n} at time t_n , from which an intermediate sample x_n is drawn, as shown in Eq. 1. The diffusion time t_n decreases progressively as

n increases. After completing all denoising steps, reaching $t_N = 0$, the final response x_N is obtained.

$$x_n \sim p_{\theta, t_n}(\cdot | x_{n-1}) \quad (1)$$

This naturally generates a series of intermediate results $x_{0:N-1}$ before arriving at the final output x_N . This behavior is analogous to the Chain-of-Thought (CoT) technique [40]. However, unlike the vertical and causal structure of CoT, the diffusion process enables the model to produce intermediate content that facilitates reaching final answers, aligning more closely with the concept of lateral thinking—solving problems through indirect and creative approaches. Thus, we define the *entire* sequence of all intermediate steps $x_{0:N}$ as the *Diffusion Chain of Lateral Thought (DCoLT)*, as shown in Fig. 1, and seek to apply reinforcement learning to train it based on the given reward function.

Compared to standard CoT, DCoLT is distinguished with several notable features:

- **Bidirectional Reasoning:** CoT generates tokens sequentially in a causal, irreversible manner; once derived, earlier thought remains fixed in the context. Instead, DCoLT adopts a bidirectional structure: each token is influenced by both preceding and succeeding content with bidirectional self-attention masks, allowing global refinement throughout generations.
- **Format-Free Reasoning:** CoT typically adheres strictly to natural language format with complete linguistic structures. DCoLT, however, relaxes this constraint, allowing early stage of intermediate steps not necessarily to be complete or correct in format, thus enabling more divergent and creative patterns of thought before finally converging to complete responses.
- **Nonlinear Generation:** CoT generates tokens one-by-one linearly in an auto-regressive manner from left to right. In contrast, DCoLT allows nonlinear generation of tokens at various positions. This aligns with how human develops ideas by beginning with keywords or critical points and then refining the details around them.

Fig. 7-Fig. 9 in Appendix B show some examples of intermediate responses from DCoLT-reinforced DLMs, demonstrating the above features with bidirectional and nonlinear generations throughout reverse diffusion processes from incomplete masked phases to complete unmasked responses.

To enhance the lateral thinking, we adopt final-outcome rewarded reinforcement learning to train the model, promoting reasoning trajectories that can lead to correct final responses x_N . For this, we treat the generation of $x_{1:N}$ as a sequence of actions to optimize *as a whole*. At each denoising step n , a distribution $\pi_{\theta, n}(\cdot | x_{n-1})$ over possible outputs is defined, which serves as the policy for sampling x_n to calculate action probabilities during RL training. The reward r is assigned based on verifiable correctness of final results. We do not impose any explicit supervision on the intermediate steps, promoting the model’s lateral thinking ability to explore diverse reasoning strategies to maximize the reward. This could yield useful patterns of thinking processes – for example, as shown in Fig. 3 of Appendix A, an easy-to-hard progressive generation of responses emerges from the DCoLT-trained DLM, in contrast to its counterpart supervisedly trained with the conventional diffusion loss.

We summarize the algorithm in Alg. 1. Details about training specific diffusion language models such as SEDD [24] and LLaDA [27] will be elaborated on in the following two subsections. Considering the demanding computation graph that expands multi-step generations in memory, we back-propagate the gradients at each step and accumulate them after the whole reverse diffusion process (c.f. Line 26 and 28 in Alg. 1). We use the GRPO [36] for fair comparisons with other RL-based models [36, 47], while alternative RL training approaches such as PPO [35] and RLOO [1] can also be adopted.

3.2 A Continuous-Time DLM Case: DCoLT-reinforced SEDD

To define a discrete diffusion process, continuous-time diffusion language models such as SEDD [24] evolve a family of distributions p_t according to a continuous-time Markov process, which can be represented by the following linear ordinary differential equation.

$$\frac{dp_t}{dt} = Q_t p_t, \quad p_0 = p_{\text{data}}, \quad p_T = p_{\text{prior}} \quad (2)$$

Here, Q_t is the transition rate matrix defining the forward process. We may first consider the simple single-token case ($Q_t \in \mathbb{R}^{V \times V}$). Conversely, to generate a sample from p_{prior} , this process has a

Algorithm 1 A General Framework for Training DCoLT

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

```

1: while  $\theta$  not converged and maximum epochs not reached do
2:   Sample questions  $q \sim \mathcal{D}$ 
3:   for  $g = 1$  to  $G$  do ▷ Generate a group of  $G$  trajectories
4:     Initialize  $x_0^g$  with  $q$  and mask tokens.
5:     for  $n = 1$  to  $N$  do ▷  $N$  denotes the number of denoising steps
6:       if training SEDD then
7:         Sample  $x_n^g \sim p_{\theta, t_n}(\cdot | x_{n-1}^g)$ 
8:       else if training LLaDA then
9:         Calculate the ranking score  $h_{\theta, n}$  for each token
10:        Sample  $K$  tokens to unmask in this step:  $\mathcal{U}_n \sim \text{Plackett-Luce}(h_{\theta, n}, K)$ 
11:        Sample  $x_n^{g, i} \sim p_{\theta, n}^i(\cdot | x_{n-1}^g)$ ,  $\forall i \in \mathcal{U}_n$ 
12:      end if
13:    end for
14:     $r^g = \text{reward\_func}(q, x_N^g)$  ▷ Compute the rewards
15:  end for
16:  for  $g = 1$  to  $G$  do ▷ Compute the advantages
17:     $A^g = \frac{r^g - \text{mean}(r^{1:G})}{\text{std}(r^{1:G})}$ 
18:  end for
19:  for  $n = 1$  to  $N$  do ▷ Compute  $\pi_\theta$  and losses for each denoising step
20:    if training SEDD then
21:       $\pi_{\theta, n}(x_n^g | x_{n-1}^g) = \prod_{i=1}^{|x_n^g|} p_{\theta, t_n}(x_n^{g, i} | x_{n-1}^g)$  ▷ see Eq. 5
22:    else if training LLaDA then
23:       $\pi_{\theta, n}(x_n^g | x_{n-1}^g) = \pi_{\theta, n}^{\text{unmask}}(\mathcal{U}_n | x_n) \cdot \pi_{\theta, n}^{\text{token}}(x_n | x_{n-1}, \mathcal{U}_n)$  ▷ see Eq. 9
24:    end if
25:     $\mathcal{L}_{\theta, n} = -\frac{1}{G} \sum_{g=1}^G \frac{\pi_{\theta, n}(x_n^g | x_{n-1}^g)}{\pi_{\text{old}, n}(x_n^g | x_{n-1}^g)} A^g$ 
26:    Calculate the gradient  $\nabla_\theta \mathcal{L}_{\theta, n}$ 
27:  end for
28:  Update  $\theta$  with accumulated gradients  $\sum_{n=1}^N \nabla_\theta \mathcal{L}_{\theta, n}$  along the descent direction
29: end while
  
```

corresponding reverse process, defined by a reverse transition rate matrix \bar{Q}_t :

$$\frac{dp_{T-t}}{dt} = \bar{Q}_{T-t} p_{T-t}, \quad \bar{Q}_t(y, x) = \frac{p_t(y)}{p_t(x)} Q_t(x, y), \quad \bar{Q}_t(x, x) = - \sum_{y \neq x} \bar{Q}_t(y, x) \quad (3)$$

By Euler's method, we have the transition probability p_{t_n} at each diffusion step t_n , resulting in an iterative formula for multi-step generations from $t_0 = T$ to $t_N = 0$ with $n = 0, \dots, N$,

$$p_{t_n}(x_n = y | x_{n-1} = x) = \delta_{xy} + (t_{n-1} - t_n) \bar{Q}_{t_{n-1}}(y, x) \quad (4)$$

with x_n denoting x_{t_n} to avoid notational clutter, i.e., the diffused sample x_t at $t = t_n$.

Specifically, we consider a representative discrete diffusion model, SEDD [24]. This model learns to approximate the concrete score, i.e. $s_\theta(x, t)_y \approx \frac{p_t(y)}{p_t(x)}$ for any $y \neq x$ to represent the probability to transfer to other tokens. Thus, we may replace $\bar{Q}_{t_{n-1}}(y, x)$ with the model-estimated $s_\theta(x, t_{n-1})_y \cdot Q_{t_{n-1}}(x, y)$ in Eq. 4, making the transition probability tractable.

When scaling to sequences, we may apply τ -leaping to define the probability of action at each timestep as the product of transition probabilities across all tokens:

$$\pi_{\theta, n}(x_n | x_{n-1}) = \prod_{i=1}^{|x_n|} p_{\theta, t_n}(x_n^i | x_{n-1}) \quad (5)$$

$$p_{\theta, t_n}(x_n^i | x_{n-1}) = \begin{cases} s_\theta(x_{n-1}, t_{n-1})_{i, x_n^i} \cdot (t_{n-1} - t_n) \cdot Q_{t_{n-1}}(x_{n-1}^i, x_n^i), & x_{n-1}^i \neq x_n^i, \\ 1 - \sum_{y \neq x_n^i} s_\theta(x_{n-1}, t_{n-1})_{i, y} \cdot (t_{n-1} - t_n) \cdot Q_{t_{n-1}}(x_{n-1}^i, y), & x_{n-1}^i = x_n^i, \end{cases} \quad (6)$$

where $|x_n|$ denotes the token length of the sequence x_n . Note that the predicted p_{t_n} is dependent on the model parameters θ . During training, we may update θ to reinforce or suppress the probability of

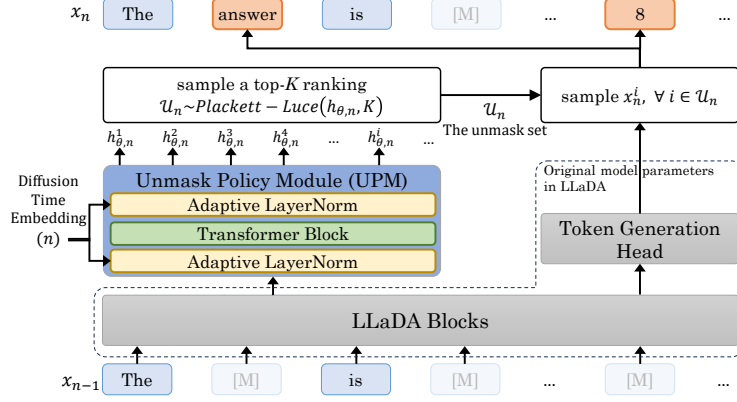


Figure 2: The structure of LLaDOU. It first predicts the token set to unmask \mathcal{U}_n according to the ranking score $h_{\theta,n}$ by the UPM, and then samples those unmasked tokens in \mathcal{U}_n by LLaDA blocks.

all actions along a trajectory of thoughts $[x_0, x_1, \dots, x_N]$, according to the reward associated with each completion. Unlike traditional diffusion model training, which typically optimizes each individual denoising step independently with ground truth responses, DCoLT jointly optimizes the entire reverse process of thought from $t_0 = T$ to $t_N = 0$, promoting the emergence of lateral thought leading to correct answers.

3.3 A Discrete-Time DLM Case: DCoLT-reinforced LLaDA

Some other diffusion language models instead operate at discrete timesteps [27, 43]. They can be viewed as Discrete-Time Diffusion Language Models with a multi-step generation process. For these models, we will figure out the probability of actions for each discrete step.

Among them is LLaDA [27] which is a typical Discrete-Time Masked Diffusion Language Model. It generates text by progressively unmasking a sequence initially prefilled entirely with mask tokens. At each generation step, the model receives a partially masked sequence as input. Then, it chooses and predicts a subset of masked tokens to reveal as clean text. As it progresses, the number of masked tokens gradually decreases, so that the model ultimately yields a fully generated sequence.

Specifically, we can define the model’s action at each step n in two parts: 1) determining the set \mathcal{U}_n of tokens to unmask, and 2) predicting the values of these tokens to obtain the new sequence x_n over the unmasked part \mathcal{U}_n .

The action of selecting which tokens to unmask at each step can be decided by ranking masked tokens with a score function under the current state. To this end, we introduce an Unmask Policy Module (UPM), which predicts a score value $h_{\theta,n}^i$ for each token i at the current diffusion step n . Based on these scores, we define a policy to sample a top- K ranked list $\mathcal{U}_n = [u_n(1), \dots, u_n(K)]$ from a Plackett–Luce model [28, 32]: a multinomial distribution is formed from the predicted scores, and K tokens are sequentially sampled without replacement, such that the corresponding scores are non-increasingly ordered $h_{\theta,n}^{u_n(1)} \geq \dots \geq h_{\theta,n}^{u_n(K)}$ with high probability.

Formally, let \mathcal{M}_n denote the set of tokens that remain masked after the n -th step, i.e., $\mathcal{M}_{n-1} \setminus \mathcal{U}_n$. Then, the probability of sampling a specific ranking list \mathcal{U}_n is given by:

$$\pi_{\theta,n}^{\text{unmask}}(\mathcal{U}_n | x_{n-1}) = \prod_{k=1}^K \frac{\exp(h_{\theta,n}^{u_n(k)})}{\sum_{j=k}^K \exp(h_{\theta,n}^{u_n(j)}) + \sum_{j \in \mathcal{M}_n} \exp(h_{\theta,n}^j)} \quad (7)$$

Specifically, UPM takes the hidden states from the last hidden layer as inputs at each denoising step n , and predicts a ranking score $h_{\theta,n}^i$ for i -th token. It contains only one transformer block, introducing marginal computation cost. Moreover, the step index n and the mask indicator of each token in x_{n-1} can benefit the model training. Thus, we embed both of them into the Unmask Policy Module (UPM) via adaptive layer normalization. The module structure is shown in Fig. 2. For convenience, we name the DCoLT-trained LLaDA with UPM by **LLaDOU** (LLaDA with Ordered Unmasking).

As shown in Fig. 10 of Appendix B, the learned scores $h_{\theta,n}$ can be viewed as the ranked confidences over the token predictions at the current step n . Higher scores indicate that the UPM predicts lower levels of diffusion noises may be present in the generated tokens, which are less likely to contain errors and thus could be unmasked with higher priorities at the current step. Fig. 7 shows the resulting unmasking orders from the learned ranking scores for some prompts.

Once the unmask set \mathcal{U}_n is determined, the model predicts their token values based on the output distribution by LLaDA blocks over the vocabulary. Viewing this prediction as a second-stage action, the probability of generating x_n given x_{n-1} and \mathcal{U}_n can be written as

$$\pi_{\theta,n}^{\text{token}}(x_n|x_{n-1},\mathcal{U}_n) = \prod_{i \in \mathcal{U}_n} p_{\theta,n}(x_n^i|x_{n-1}). \quad (8)$$

Thus, the probability of the complete policy for transitioning from x_{n-1} to x_n is the product of those of the unmask policy and the token prediction policy as used in Line 23 of Alg. 1:

$$\pi_{\theta,n}(x_n|x_{n-1}) = \pi_{\theta,n}^{\text{unmask}}(\mathcal{U}_n|x_{n-1}) \cdot \pi_{\theta,n}^{\text{token}}(x_n|x_{n-1},\mathcal{U}_n). \quad (9)$$

4 Experiments

We conduct DCoLT experiments on two DLMs – SEDD 400M and LLaDA 8B, each representing one of the two categories mentioned in Sec. 3.

4.1 Experiments on SEDD+DCoLT

4.1.1 Settings

We first conduct experiments on SEDD [24] for two different tasks: 1) Sudoku 4×4 task and 2) GSM8K-Aug math problems [10], both assessing math reasoning performance. These two tasks allow us to perform direct comparisons with existing reasoning algorithms, including both supervised fine-tuning [10, 42] and reinforcement learning approaches [36]. For all experiments, we use the rule-based reward function to compute the reward for each sample: a reward of 1 is assigned if the solution is correct, and 0 otherwise. We choose SEDD as the base model to conduct DCoLT training, which is of medium size around 400 M. For detailed settings please refer to Appendix C.1.

4.1.2 Results

From the results in Tab. 1, our method remarkably outperforms both (implicit) CoT [10] and DoT [42] on both tasks. While both CoT and DoT perform vertical thinking, DoT also uses the SEDD as its base model to simulate causal reasoning. In contrast, our approach only uses RL to train the model without any CoT data for supervised training. For CoT, we adopt an auto-regressive language model – GPT2 [31] as the base model, which has the similar model size of about 400M to the SEDD model.

Particularly, SEDD + DCoLT achieves an accuracy of 96.2% on Sudoku 4×4 task, the best performance among all the methods. Using the same SEDD model, DCoLT greatly exceeds DoT (79.4%), even though the latter introduce step-by-step annotated CoT data for supervised training. For GPT2 model, both CoT and Implicit CoT, no matter if being post-trained with SFT or RL, perform worse than our model clearly. On the GSM8K-Aug dataset, SEDD + DCoLT also reaches 57.0% in accuracy, which is notably higher than both DoT and CoT counterparts. We will further analyze the generation process of SEDD + DCoLT in Appendix A, demonstrating how tokens are generated in a lateral thinking manner.

4.2 Experiments on LLaDOU (LLaDA+DCoLT)

4.2.1 Settings

We further apply DCoLT to LLaDA [27], a SOTA discrete-time masked-based DLM with 8B parameters. During generation, it starts with a masked sequence of length 256, and performs a 256-step reverse diffusion process. We adopt a block-wise unmasking strategy as in LLaDA [27], and divide the sequence into blocks of length 8. In Tab. 2, both LLaDA and LLaDOU are evaluated with this setting for fair comparison. We conduct experiments with 16 H800 GPUs to jointly train both UPM and LLaDA parts. For implementation details, please refer to Appendix C.2.

Table 1: Experimental Results on SEDD. All models listed below are of medium size, with approximately 400M parameters. On the Sudoku 4×4 dataset, we report the accuracy on a test set of 2,000 samples each with 9 empty cells. On GSM8K-Aug, we use the test set from GSM8K[8] for evaluation following [10].

Models	Post-Training	Sudoku 4×4	GSM8K-Aug
GPT2 + CoT[10]	SFT	71.5	43.9
GPT2 + Implicit CoT[10]	SFT	-	21.9
GPT2 + CoT	RL	74.6	-
SEDD + DoT [42]	SFT	79.4	53.5
SEDD + DCoLT	RL	96.2	57.0

Table 2: Model accuracies on math and code generation benchmarks. The "Post-Training" column indicates what kinds of post-training (SFT or RL) phase the model goes through. [†] indicates the model uses additional proprietary training data for the post-training phase. The numbers in parentheses represent the number of shots for the in-context learning, with “-” indicating unknown cases not mentioned in original papers. The results denoted with * are evaluated with the prompt templates in Appendix C.2 for fair comparison, while others are reported in original papers. The results denoted with [‡] are evaluated on a subset MATH-500 instead of MATH. We highlight the best-performing model among compared DLMs in **bold**.

Method	Post-Training	Math		Code	
		GSM8K	MATH	HumanEval	MBPP
Diffusion Language Models					
Dream 7B* [43]	baseline	81.1 (0)	42.9 (0)	51.8 (0)	49.6 (0)
LLaDA 8B* [27]	baseline	78.3 (0)	38.9 (0)	39.6 (0)	40.2 (0)
+ SFT [47]	+ SFT	81.1 (0)	34.8 [‡] (0)	-	-
+ diffu-GRPO [47]	+ RL	81.9 (0)	39.2 [‡] (0)	-	-
d1-LLaDA (SFT + diffu-GRPO)	+ SFT + RL	82.1 (0)	40.2 [‡] (0)	-	-
LLaDOU (LLaDA + DCoLT) 8B*	+ RL	88.1 (0)	44.6 (0)	59.1 (0)	51.6 (0)
Auto-regressive Models					
LLaMA2 7B [39]	baseline	14.6 (0)	2.5 (0)	12.8 (0)	20.8 (3)
MetaMath 7B [45]	+ SFT [†]	66.5 (0)	19.8 (0)	-	-
CodeLLaMA-Instruct 7B [34]	+ SFT [†]	-	-	34.8 (0)	44.4 (3)
Deepseek 7B [5]	baseline	63.0 (0)	15.8 (0)	48.2 (0)	35.2 (3)
DeepseekMath-Instruct 7B [36]	+ SFT [†]	82.9 (-)	46.8 (-)	-	-
DeepseekMath-RL 7B [36]	+ SFT [†] + RL [†]	88.2 (-)	51.7 (-)	-	-
DeepseekCoder-Instruct 7B [14]	+ SFT [†]	-	-	66.1 (-)	65.4 (-)

4.2.2 Results

As shown in Tab. 2, among all compared DLMs, LLaDOU consistently achieves the best performance across all benchmarks. On GSM8K, LLaDOU reaches 88.1%, significantly higher than other methods. On the more challenging MATH dataset, LLaDOU achieves 44.6%, outperforming baseline models such as LLaDA 8B (+5.7%) and Dream 7B (+1.7%), as well as models enhanced with post-training by SFT, RL or both. Notably, d1-LLaDA [47] uses 1K questions paired with detailed reasoning traces [26] in SFT to boost reasoning capabilities. In contrast, LLaDOU achieves superior performance – +6.0% on GSM8K and +4.4% on MATH – without any reasoning supervision, relying solely on reward signals based on the final answers.

Even when compared with auto-regressive models trained with a significantly larger amount of proprietary data, LLaDOU remains highly competitive. For example, DeepseekMath [36] gathers 776K questions with reasoning process annotations for SFT and 144K questions for RL, while LLaDOU is trained on just 15K public training samples from GSM8K and MATH. Despite using nearly two orders of magnitude fewer data for RL training, LLaDOU matches DeepseekMath’s performance on GSM8K (88.1%), highlighting its data efficiency and the effectiveness of reward-driven reasoning via diffusion-based lateral reasoning. In Appendix B, we provide a detailed analysis of the thinking process generated by LLaDOU, including both visualization examples and quantitative evaluations. Fig. 12 illustrates the reward curve during DCoLT training in the math domain, showing an increasing trend in the outcome-based reward over training iterations.

Table 3: Accuracy of LLaDOU on MATH subsets across difficulty levels (5: the hardest and 1: the easiest) and generation lengths. The result at $L = 512$ is obtained from the model fine-tuned with this length, while results at other lengths are directly evaluated using LLaDOU trained with $L = 256$.

Length	Level 1	Level 2	Level 3	Level 4	Level 5
128	80.8%	61.0%	45.6%	29.5%	13.2%
256	83.3%	65.1%	52.3%	35.9%	18.7%
384	82.4%	66.6%	54.9%	39.0%	20.2%
512	82.6%	69.7%	56.9%	40.2%	21.5%

Similar trends are observed in code generation benchmarks. Unlike models that rely on valuable ground-truth code for supervised training, LLaDOU requires no code for supervised training at all. Instead, it only needs several test cases to provide a simple outcome-based reward for reinforcement learning: The model is rewarded only if the generated code passes all unit tests. Despite this code-free reward signal, LLaDOU achieves 51.6% on MBPP and 59.1% on HumanEval, outperforming other DLMs and auto-regressive models, except for DeepseekCoder [14] that benefits from 2B tokens of high-quality instruction-tuning data—approximately $150\times$ more than we used (about 12M tokens). Details of post-training datasets used by all compared methods are provided in the Appendix C.3.

4.2.3 Ablation Studies

We provide some ablation studies to reveal the role of some model design and hyper-parameter settings in LLaDOU. Unless otherwise specified, all experiments are conducted on LLaDOU with $N = 64$ diffusion steps, and the model is trained for 150 iterations with a batch size of 32.

Better performance with the UPM In Tab. 4, we ablate the model design in LLaDOU. Even if we freeze the original model parameters in LLaDA part and only train the UPM part, it achieves a significant improvement from 47.27% to 69.24%. This result indicates that the unmasking policy plays a crucial role in our model. By training the LLaDA part with RL together, the accuracy further improves to 81.06%.

Table 4: Ablation on the Unmasking Policy Module

Trained parameters		GSM8K Acc.
UPM	LLaDA	
×	×	47.27 (Baseline)
w/ AdaLN	×	69.24
w/o AdaLN	✓	80.53
w/ AdaLN	✓	81.06

We use the adaptive layernorm in the UPM. Removing it degrades accuracy to 80.53%. This demonstrates that incorporating the embeddings of diffusion step n and the mask indicators also benefits the training of the unmasking policy.

Extension to longer generation length Though the models in Tab. 2 is trained to generate completions with a fixed generation length 256 (i.e., the length of initialized mask tokens), it can benefit from generating longer sequences without further RL training on longer generations. This is especially effective when answering difficult questions on MATH dataset, in which solving problems sometimes require longer reasoning. Tab. 5 further summarizes the overall results under different generation lengths and models. It suggests that longer generations in inference lead to better accuracies. For example, a generation length of 384 achieves an accuracy of 46.22%, which is +1.62% higher than the default length of 256. This improvement vanishes as we further increase the generation length. However, if we further tune the model on a longer generation length, for example $L = 512$, the accuracy further improves to 47.3%.

In Tab. 3, we report the accuracy of LLaDOU on MATH subsets across difficulty levels. The accuracy on level-1 problems saturates with length 256, suggesting that shorter generations are sufficient for simpler questions. In contrast, the performance on harder problems (levels 2–5) continues to improve with longer generations, indicating that complex reasoning benefits from longer responses. These results suggest the potential of how scaled generation lengths may improve the performance of DCoLT-trained DLMs, as shown in Fig. 13 of Appendix D.

Table 5: Model accuracies with extended generation lengths on MATH without further RL-finetuning the model on these longer generations. The results denoted with * are taken from [47] and evaluated on a subset MATH-500 instead of MATH.

Model	Generation Length				
	256	320	384	448	512
LLaDA 8B [27]	38.9	40.1	41.5	42.3	42.5
+ SFT* [47]	38.6	-	-	-	40.2
+ diffu-GRPO* [47]	37.2	-	-	-	39.2
d1-LLaDA (SFT+diffu-GRPO)* [47]	38.6	-	-	-	40.2
LLaDOU (LLaDA + DCoLT) 8B	44.6	45.7	46.2	45.7	45.9

5 Conclusion

In this paper, we propose a new paradigm of model reasoning ability, Diffusion Chain of Lateral Thought (DCoLT), for diffusion language models. DCoLT considers the intermediate steps in the reverse diffusion process as the thinking actions, and optimizes this process with final-outcome rewarded Reinforcement Learning (RL). We implement DCoLT on two representative diffusion language models, SEDD and LLaDA. On SEDD, we derive the corresponding RL optimization over the distribution chain of thinking actions via the predicted concrete scores at denoising steps. On the masked-based diffusion language model LLaDA, we regard the order in which tokens are unmasked at each diffusion step as its action, and propose an Unmask Policy Module (UPM) to optimize over the order distribution with the Plackett-Luce model. The experiments indicate that DCoLT outperforms other SFT- or RL-based training algorithms, and demonstrate its effectiveness on a wide range of tasks and benchmarks, including math and code generation.

Limitations First, due to limited training data and compute, our model’s performances on Math and code generation tasks still have much rooms to improve. Prior works [14, 34, 36, 45] demonstrate that proprietary data often significantly improves the model performances. Also, scaling the token length of sequences during training can also enhance reasoning. We will seek to boost the model ability along these directions when computing resources become available. Second, currently DCoLT is only validated on tasks having a verifiable reward function. We may need a reward model to cover more general tasks. We will develop these further in our future research.

Acknowledgements

This work was supported by National Natural Science Foundation of China under Grant No. 92467104, and Zhejiang Leading Innovative and Entrepreneur Team Introduction Program (2024R01007).

References

- [1] Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce style optimization for learning from human feedback in llms. *arXiv preprint arXiv:2402.14740*, 2024.
- [2] Marianne Arriola, Aaron Gokaslan, Justin T Chiu, Zhihan Yang, Zhixuan Qi, Jiaqi Han, Subham Sekhar Sahoo, and Volodymyr Kuleshov. Block diffusion: Interpolating between autoregressive and diffusion language models. *arXiv preprint arXiv:2503.09573*, 2025.
- [3] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in neural information processing systems*, 34:17981–17993, 2021.
- [4] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- [5] Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiusi Du, Zhe Fu, et al. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*, 2024.

- [6] Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. *Advances in Neural Information Processing Systems*, 35:28266–28279, 2022.
- [7] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- [8] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems, 2021. URL <https://arxiv.org/abs/2110.14168>, 9, 2021.
- [9] DeepSeek-AI. Deepseek-v3 technical report, 2024. URL <https://arxiv.org/abs/2412.19437>.
- [10] Yuntian Deng, Kiran Prasad, Roland Fernandez, Paul Smolensky, Vishrav Chaudhary, and Stuart Shieber. Implicit chain of thought reasoning via knowledge distillation. *arXiv preprint arXiv:2311.01460*, 2023.
- [11] Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37:133345–133385, 2024.
- [12] Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An, Peilin Zhao, Wei Bi, Jiawei Han, et al. Scaling diffusion language models via adaptation from autoregressive models. *arXiv preprint arXiv:2410.17891*, 2024.
- [13] Ishaan Gulrajani and Tatsunori B Hashimoto. Likelihood-based diffusion language models. *Advances in Neural Information Processing Systems*, 36:16693–16715, 2023.
- [14] Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao 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.
- [15] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- [16] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*, 2021.
- [17] James S Hernandez and MBBS Prathibha Varkey. Vertical versus lateral thinking. *Physician executive*, 34(3):26, 2008.
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [19] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in neural information processing systems*, 34:12454–12465, 2021.
- [20] Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024.
- [21] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35: 26565–26577, 2022.
- [22] Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, et al. Training language models to self-correct via reinforcement learning. *arXiv preprint arXiv:2409.12917*, 2024.

- [23] Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*, 2023.
- [24] Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Forty-first International Conference on Machine Learning*, 2024.
- [25] Yu Meng, Mengzhou Xia, and Danqi Chen. Simpo: Simple preference optimization with a reference-free reward. *Advances in Neural Information Processing Systems*, 37:124198–124235, 2024.
- [26] Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time scaling. *arXiv preprint arXiv:2501.19393*, 2025.
- [27] Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv preprint arXiv:2502.09992*, 2025.
- [28] Shuzi Niu, Yanyan Lan, Jiafeng Guo, and Xueqi Cheng. A new probabilistic model for top-k ranking problem. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pages 2519–2522, 2012.
- [29] Jingyang Ou, Shen Nie, Kaiwen Xue, Fengqi Zhu, Jiacheng Sun, Zhenguo Li, and Chongxuan Li. Your absorbing discrete diffusion secretly models the conditional distributions of clean data. *arXiv preprint arXiv:2406.03736*, 2024.
- [30] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- [31] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.
- [32] Stephen Ragain and Johan Ugander. Choosing to rank. *arXiv preprint arXiv:1809.05139*, 2018.
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [34] Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, et al. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*, 2023.
- [35] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [36] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- [37] Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis Titsias. Simplified and generalized masked diffusion for discrete data. *Advances in neural information processing systems*, 37:103131–103167, 2024.
- [38] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [39] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

- [40] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- [41] Zhangchen Xu, Yang Liu, Yueqin Yin, Mingyuan Zhou, and Radha Poovendran. Kodcode: A diverse, challenging, and verifiable synthetic dataset for coding. *arXiv preprint arXiv:2503.02951*, 2025.
- [42] Jiacheng Ye, Shansan Gong, Liheng Chen, Lin Zheng, Jiahui Gao, Han Shi, Chuan Wu, Xin Jiang, Zhenguo Li, Wei Bi, et al. Diffusion of thought: Chain-of-thought reasoning in diffusion language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [43] Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng Kong. Dream 7b, 2025. URL <https://hkunlp.github.io/blog/2025/dream>.
- [44] Zilyu Ye, Zhiyang Chen, Tiancheng Li, Zemin Huang, Weijian Luo, and Guo-Jun Qi. Schedule on the fly: Diffusion time prediction for faster and better image generation, June 2025. URL <https://arxiv.org/abs/2412.01243>.
- [45] Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models. *arXiv preprint arXiv:2309.12284*, 2023.
- [46] Oussama Zekri and Nicolas Boullé. Fine-tuning discrete diffusion models with policy gradient methods. *arXiv preprint arXiv:2502.01384*, 2025.
- [47] Siyan Zhao, Devaansh Gupta, Qinqing Zheng, and Aditya Grover. d1: Scaling reasoning in diffusion large language models via reinforcement learning. *arXiv preprint arXiv:2504.12216*, 2025.
- [48] Kaiwen Zheng, Yongxin Chen, Hanzi Mao, Ming-Yu Liu, Jun Zhu, and Qincheng Zhang. Masked diffusion models are secretly time-agnostic masked models and exploit inaccurate categorical sampling. *arXiv preprint arXiv:2409.02908*, 2024.
- [49] Lin Zheng, Jianbo Yuan, Lei Yu, and Lingpeng Kong. A reparameterized discrete diffusion model for text generation. *arXiv preprint arXiv:2302.05737*, 2023.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading “NeurIPS Paper Checklist”,**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification: The proposed Diffusion Chain of Lateral Thought is introduced in Section 3, and the results highlighted in the abstract and introduction can be found in Section 4.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We have discussed the limitations in Section 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Please follow Appendix C for both train and evaluation details. Moreover, we will release our model later.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We will release the code later.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: Please refer to Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[No\]](#)

Justification: Due to the high computational cost of training large language models, we omit error bars in our reported results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: Please refer to Appendix C for the computer resources we use.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: The research conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Please refer to Appendix G.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [Yes]

Justification: See Appendix H.

Guidelines:

- The answer NA means that the paper poses no such risks.

- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [\[Yes\]](#)

Justification: See Appendix F.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[NA\]](#)

Justification: We do not release new assets in the submission phase.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[NA\]](#)

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: The paper describe the CoT dataset created by LLM in Appendix C.1

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Analysis of Thinking Process in SEDD+DCoLT

For the Sudoku dataset, we aim to visualize the thinking process by analyzing the generation step of different cells. We categorize all 4×4 cells into three groups: (1) *given cells* – those initially provided by the puzzle; (2) *easy cells* – those that can be directly inferred using basic Sudoku rules (each number 1–4 must appear exactly once per row, column, and 2×2 subgrid); and (3) *hard cells* – the remaining positions that require more complex reasoning, often involving the resolution of dependencies among other non-given cells.

In Fig. 3, we plot the distribution of generation step for these three cell types, the left side displays statistics for the SEDD + SFT model, while the right side shows our further RL-trained SEDD + DCoLT model. The right figure reveals that the our SEDD + DCoLT significantly prioritizes

generating *easy cells* in earlier steps than *hard cells*, as these cells can be directly determined from the initial *given cells*, indicating that our model learns a progressive generation strategy from easy to hard. Notably, in the left figure, we observe almost no difference between generation orders across various difficulty levels for SEDD + SFT model. This comparison reveals that the DCoLT changes its SFT-trained counterpart’s reasoning behavior to more progressive generation aligned naturally with how humankind handles problems in a real world.

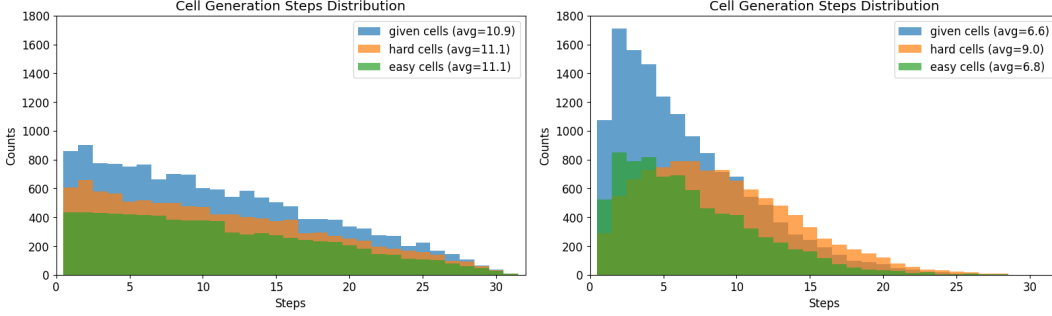


Figure 3: This figure shows the model’s inference patterns on Sudoku dataset, with the left side displaying the pattern of SEDD + SFT model and the right side showing that of our SEDD + DCoLT model. We plot the total number of generated tokens for these three cell types over diffusion steps on Sudoku 4×4 dataset, with the number of sampling steps set to 32.

To further demonstrates the progressive generation behavior on sample level, we visualize the generation process of our SEDD + DCoLT on the Sudoku dataset in Fig. 4. We show the predicted \hat{x}_0 at each timestep, obtained by selecting the token with the highest probability (excluding the mask token) at each token position. As shown in the top row, positions (0, 0) and (0, 1) can be regarded as *hard cells*, because the value can’t be directly inferred from *given cells* denoted by black borders. Initially, our model assigns the highest probabilities to incorrect answers in these red cells due to insufficient contextual information. However, as the denoising process progresses, the model gradually refines its predictions by leveraging information from newly unmasked cells (shown in dark blue), eventually converging to the correct values. This demonstrates the importance of nonlinear generation for Sudoku 4×4 tasks.

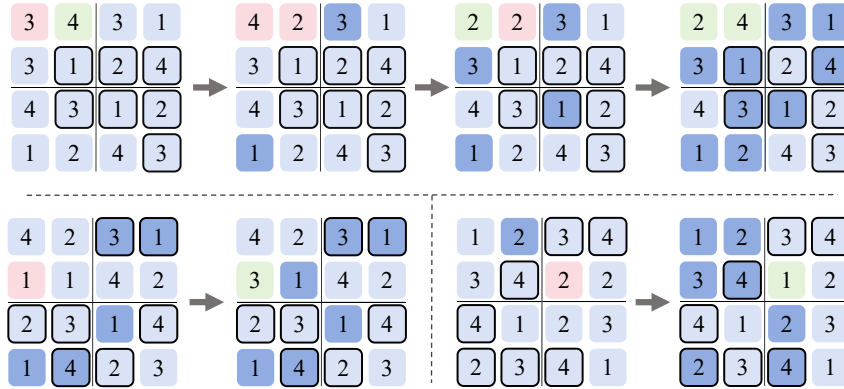


Figure 4: Visualization of predicted tokens \hat{x}_0 by SEDD + DCoLT on Sudoku 4×4 : those that are still masked appear in light blue, with unmasked ones in dark blue, incorrect predictions in red, and corrected ones in green. Black borders indicate the *given cells*.

For GSM8K-Aug dataset, since it’s hard to determine which token position is more difficult to predict, we only show the generation process and the generation step statistics for each position. Fig. 5 demonstrates the contrasting generation behaviors between SEDD + DCoLT and GPT2 + CoT on the GSM8K-Aug dataset. While GPT2 + CoT consistently follows a left-to-right generation order (where earlier generation steps correspond to lower token positions), SEDD + DCoLT shows minimal positional bias in generation order across token positions. Furthermore, our results in Fig. 6 reveal

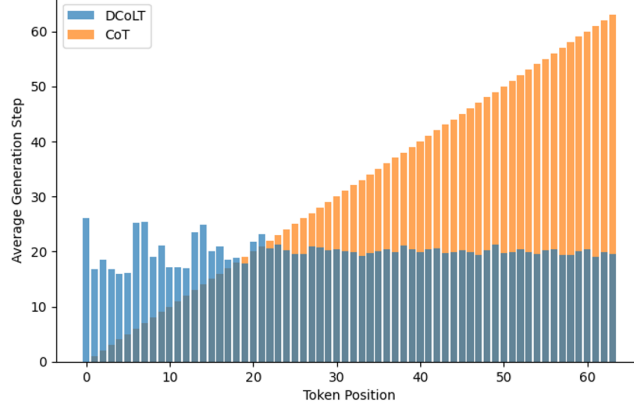


Figure 5: Average generation step for each token position in GSM8K-Aug. Different from CoT, SEDD + DCoLT generates in a non-linear way.

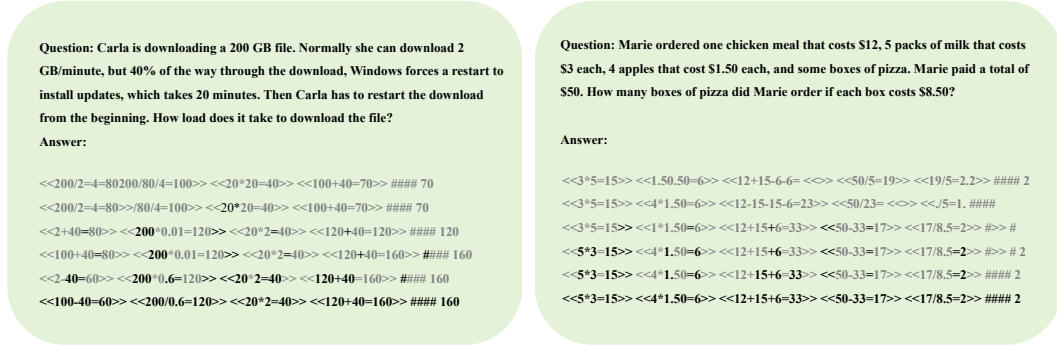


Figure 6: Our model’s generation process on GSM8K-Aug, each line displaying the predicted- \hat{x}_0 tokens across different steps, with masked tokens shown in gray and unmasked ones shown in black.

that SEDD + DCoLT’s generation order is sample-dependent, maintaining a nonlinear generation throughout the process.

B Analysis of Thinking Process in LLaDOU (LLaDA+DCoLT)

We illustrate how LLaDOU generates a 64-token answer to a GSM8K question in Fig. 7. Tokens shown in darker shades are generated in later diffusion steps. As the figure shows, key numbers and symbols tend to emerge early in the generation process, while surrounding textual elements are filled in later to ensure grammatical correctness and fluency.

Fig. 8 and Fig. 9 further show intermediate steps of the output when generating responses of length 256 with 64 steps, in which key numbers and symbols are generated to gradually form the structure of thinking processes before grammatically correct sentences are completed.

Fig. 10 visualizes the predicted ranking scores $h_{\theta,n}$ to unmask tokens during the generation process. The examples show that at each step, the unmasking score evaluates how likely the predicted token is correct at each position. For example, in Step 20 of case (a), the token “3” prior to the selected token “0” has a smaller value of the unmasking score, suggesting that it is more likely than “0” to be wrong. This token is corrected to “0” before being unmasked at a later step.

M ig uel uses 2 x 30 = 60 sheets of paper per week . If there are 4 weeks in a month , then Miguel will use 60 x 4 = 240 sheets of paper every month . Here is the solution : 240 <|eot_id|>
<|endoftext|>

(a) Miguel uses 2 pads of paper a week for his drawing. If there are 30 sheets of paper on a pad of paper, how many sheets of paper does he use every month?

If each train travels 80 miles west and then 150 miles north , each train travels a total of 80 + 150 = 230 miles . Therefore , each train covers 230 miles in the two days . Result : 230 <|eot_id|>
<|endoftext|>

(b) Two trains leave San Rafael at the same time. They begin traveling westward, both traveling for 80 miles. The next day, they travel northwards, covering 150 miles. What's the distance covered by each train in the two days?

She pays 500 + 800 = \$1300 She pays 1300 * .10 = \$130 for insurance That means in total she paid 1300 + 130 = \$1430 The answer is 1430 <|eot_id|> <|endoftext|>

(c) Janet buys a brooch for her daughter. She pays \$500 for the material to make it and then another \$800 for the jeweler to construct it. After that, she pays 10% of that to get it insured. How much did she pay?

The child already has \$4 savings , so he needs to save an additional \$12 - \$4 = \$8 . If he promises to save \$2 daily , then it will take him \$8 / \$2 = 4 days to save the remaining amount of money . Solution : 4 <|eot_id|> <|endoftext|>

(d) An 8-year old child wants to buy a toy car which costs \$12. He already has \$4 savings. How many days will it take him to save the remaining amount of money if he promises to save \$2 daily from his allowance?

Figure 7: Examples of 64-token responses generated by LLaDOU on GSM8K. Token color reflects the generation order, progressing from light (early steps) to dark (later steps). The corresponding questions are provided in the sub-captions. We observe from the results that many tokens bearing key information to final answers such as the numbers, units and factual entities tend to be unmasked at earlier steps, while those tokens for function words and symbols such as “of” and “the” tend to be unmasked at later steps.

[illegible][illegible]

To determine the total number of bolts used, we follow the following steps:

Step 1: Identify the amount of blue fiber used.
 The robe takes 2 bolts of blue fiber.

Step 2: Determine the amount of white fiber used.
 The robe takes half as much white fiber as blue fiber. Since there are 2 bolts of blue fiber, the amount of white fiber is:

$$\frac{1}{2} \times 2 = 1$$
 bolt of white fiber.

Step 3: Calculate the total number of bolts.
 Add the total number of bolts from the blue fiber and white fiber:

$$2 + 1 = 3$$
 bolts.

Therefore, the total number of bolts used is 3.

[illegible]

To determine the total number of bolts used to make the robe, we follow follow follow steps:

****Step 1: Identify the amount of blue fiber used.****
The robe takes 2 bolts of blue fiber.

****Step 2: Determine the amount of white fiber used.****
The robe uses half as much white fiber as blue fiber. Since there are 2 bolts of blue fiber, the amount of white fiber is:
$$\sqrt{\frac{1}{2}} \times 2 = 1 \text{ bolt}$$

****Step 3: Calculate the total number of bolts used.****
Add the total number of bolts from the blue fiber and the white fiber:
$$\sqrt{2 + 1} = 3 \text{ bolts}$$

Therefore, the total number of bolts used to make the robe is $\sqrt{\boxed{3}}$.

Table 6: Evaluation Hyperparameters for Sudoku and GSM8K-Aug

Models	Sudoku			GSM8K-Aug		
	Generation Tokens	NFEs	Temp.	Generation Tokens	NFEs	Temp.
GPT2 + CoT	256	256	0.0	256	256	-
SEDD + DoT	256	512	0.5	256	64	0.5
SEDD + DCoLT	19	32	0.5	64	64	0.0

Models	Post-Training	Post-Training Steps	Learning Rate	Batch Size
GPT2 + CoT	SFT	1,500	3e-4	1024
SEDD + DoT	SFT	1,500	1e-4	1024
GPT2 + CoT	RL	500	1e-5	64
SEDD + DCoLT	RL	500	1e-5	64

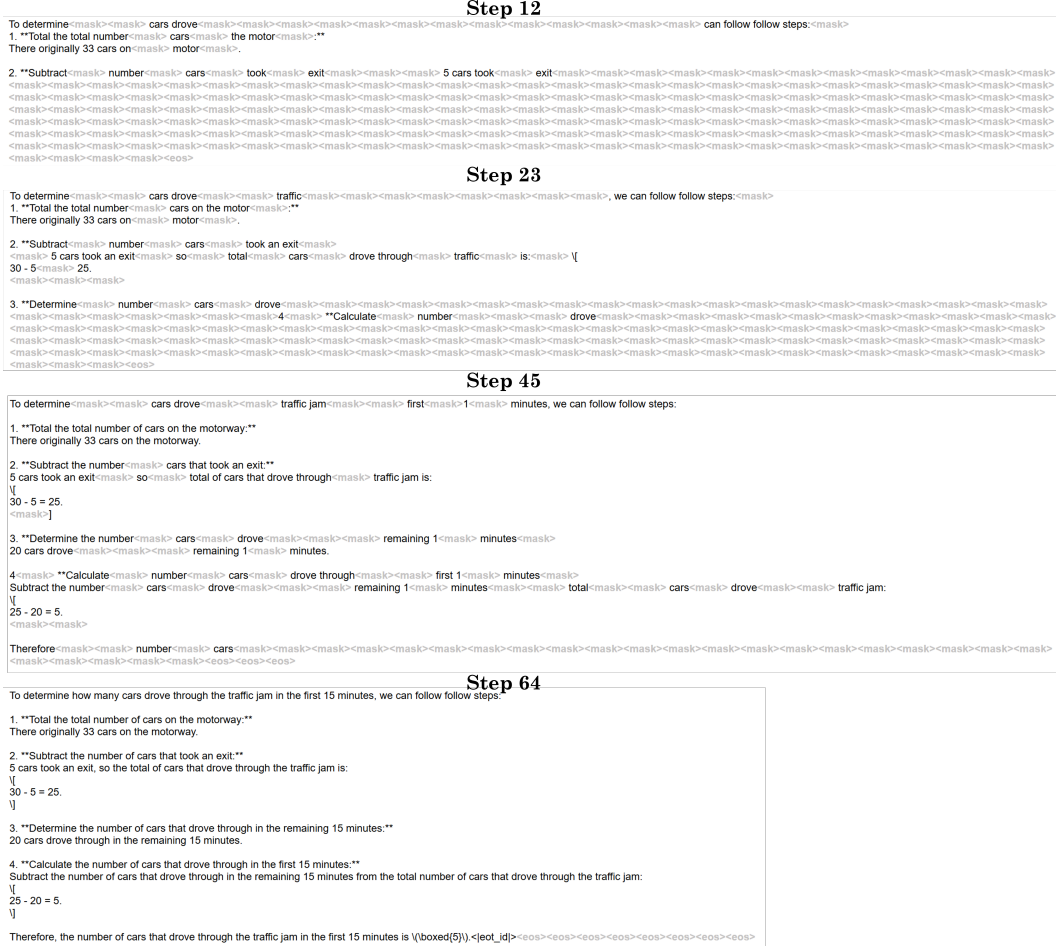


Figure 9: Intermediate step outputs from LLaDOU during the generation of a 256-token response. The question is from GSM8K: *Cars have lined up on the motorway. Some of the cars drive through in the first 15 minutes of the traffic jam, then 20 more cars drive through in the remaining 15 minutes of the jam. 5 cars from the line take an exit so they don't have to drive through the traffic jam. If there were originally 30 cars on the motorway, how many cars drove through the traffic jam in the first 15 minutes?*

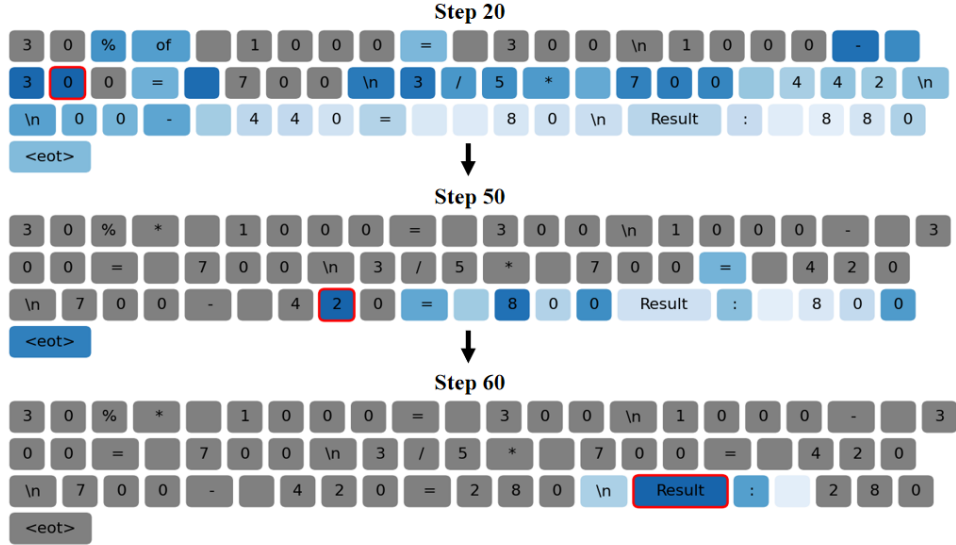
C Experiment Settings

C.1 SEDD+DCoLT

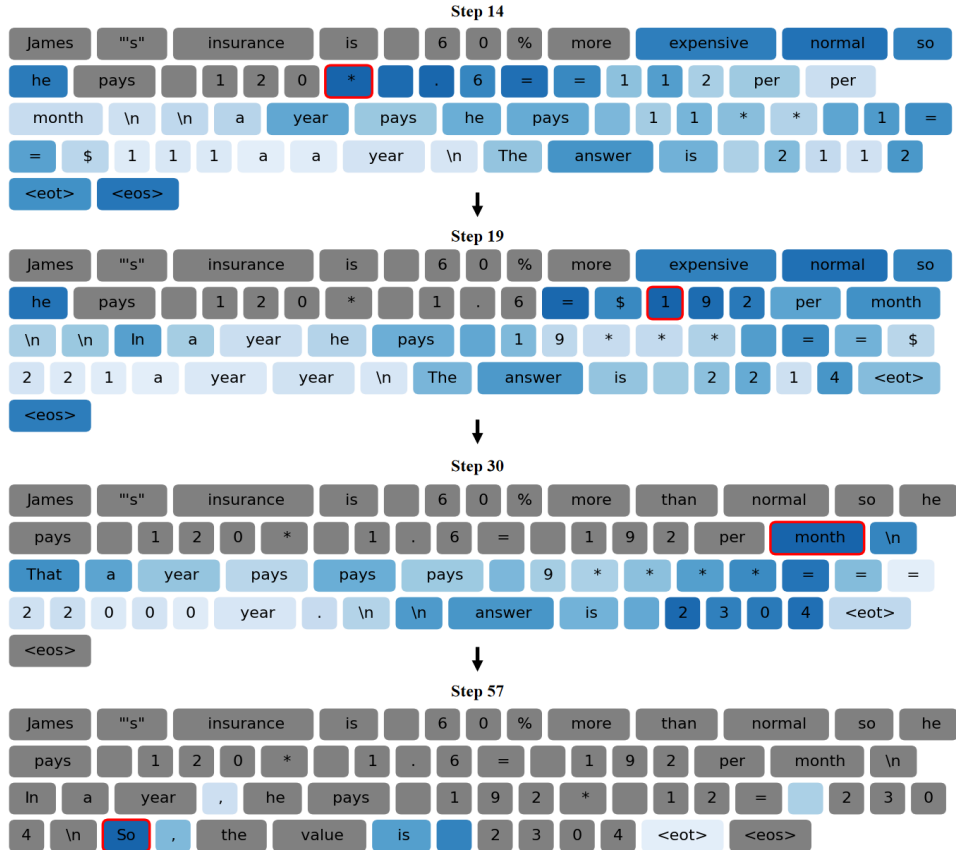
Sudoku 4×4 We first generated 50,000 sudoku puzzles of size 4×4 by uniformly replacing 1-9 cells with zeros. Using DeepSeek V3-0324[9], we produced CoT reasoning steps and filtered the results to remove repetitive responses and incorrect solutions. All training datasets were created from this cleaned dataset to ensure fair comparison. Dataset samples can be found in the Fig. 11.

We trained both DoT and CoT (with SFT) for 1,500 steps. For CoT (with RL), we first performed 1,000 steps of SFT to pretrain the CoT response, reaching 70.5% in accuracy, followed by another 500 steps of RL training using GRPO[36]. For a fair comparison, DCoLT was also pre-trained to reach 72.3% accuracy before applying 500 steps of RL training. Tab. 7 presents training hyperparameters for the Sudoku 4×4 dataset. For both SFT and RL post-training, we employ Adam optimizer with $(\beta_1, \beta_2) = (0.9, 0.999)$. During RL post-training, we set the group size of GRPO to 32, and the kl coefficient to 0.

For the Sudoku 4×4 task, the reward is determined by verifying whether the generated solution satisfies all the rules of a valid 4×4 Sudoku, namely that each row, each column, and each 2×2



(a) The number of students in a school hall was 1000. The hall had 3 entrances A, B, and C which also served as the exits. If after a meeting 30% of the students went out of the building through exit A, $\frac{3}{5}$ of the remaining went out through exit B, and the rest went out through exit C, calculate the number of students who went out through exit C.



(b) James has to buy insurance. Since he had an accident it was 60% more than normal. The normal cost is \$120 a month. How much does he pay a year?

Figure 10: Illustration of intermediate steps of generated responses by LLaDOU on GSM8K. Gray tokens represent positions that have already been unmasked prior to the current step. For each masked position, we show the predicted token with the highest generation probability. Token color indicates the unmasking score, with darker blue indicating higher probability of be unmasked. The token eventually selected to unmask at each step is bounded in a red box. The corresponding prompts are given in the sub-captions.

<p>Fill the position where value are 0 in a 4x4 Sudoku puzzle with digits 1-4, where each row, column, and 2x2 subgrid contains the numbers from 1 to 4 exactly once.</p> <p>Input:</p> <pre> 1 4 2 3 0 3 0 4 3 0 0 2 4 0 3 1 </pre>																							
<p>Output:</p> <table> <tr> <th>Step 1:</th><th>Step 2:</th><th>Step 3:</th><th></th></tr> <tr> <td>1 4 2 3</td><td>1 4 2 3</td><td>1 4 2 3</td><td></td></tr> <tr> <td>2 3 1 4</td><td>2 3 1 4</td><td>2 3 1 4</td><td></td></tr> <tr> <td>3 0 0 2</td><td>3 1 4 2</td><td>3 1 4 2</td><td></td></tr> <tr> <td>4 0 3 1</td><td>4 0 3 1</td><td>4 2 3 1</td><td></td></tr> </table>				Step 1:	Step 2:	Step 3:		1 4 2 3	1 4 2 3	1 4 2 3		2 3 1 4	2 3 1 4	2 3 1 4		3 0 0 2	3 1 4 2	3 1 4 2		4 0 3 1	4 0 3 1	4 2 3 1	
Step 1:	Step 2:	Step 3:																					
1 4 2 3	1 4 2 3	1 4 2 3																					
2 3 1 4	2 3 1 4	2 3 1 4																					
3 0 0 2	3 1 4 2	3 1 4 2																					
4 0 3 1	4 0 3 1	4 2 3 1																					

<p>Fill the position where value are 0 in a 4x4 Sudoku puzzle with digits 1-4, where each row, column, and 2x2 subgrid contains the numbers from 1 to 4 exactly once.</p> <p>Input:</p> <pre> 3 4 2 1 0 2 4 3 4 0 0 0 0 3 1 4 </pre>																							
<p>Output:</p> <table> <tr> <th>Step 1:</th><th>Step 2:</th><th>Step 3:</th><th>Step 4:</th></tr> <tr> <td>3 4 2 1</td><td>3 4 2 1</td><td>3 4 2 1</td><td>3 4 2 1</td></tr> <tr> <td>1 2 4 3</td><td>1 2 4 3</td><td>1 2 4 3</td><td>1 2 4 3</td></tr> <tr> <td>4 0 0 0</td><td>4 1 0 0</td><td>4 1 3 2</td><td>4 1 3 2</td></tr> <tr> <td>0 3 1 4</td><td>0 3 1 4</td><td>0 3 1 4</td><td>2 3 1 4</td></tr> </table>				Step 1:	Step 2:	Step 3:	Step 4:	3 4 2 1	3 4 2 1	3 4 2 1	3 4 2 1	1 2 4 3	1 2 4 3	1 2 4 3	1 2 4 3	4 0 0 0	4 1 0 0	4 1 3 2	4 1 3 2	0 3 1 4	0 3 1 4	0 3 1 4	2 3 1 4
Step 1:	Step 2:	Step 3:	Step 4:																				
3 4 2 1	3 4 2 1	3 4 2 1	3 4 2 1																				
1 2 4 3	1 2 4 3	1 2 4 3	1 2 4 3																				
4 0 0 0	4 1 0 0	4 1 3 2	4 1 3 2																				
0 3 1 4	0 3 1 4	0 3 1 4	2 3 1 4																				

Figure 11: Examples of CoT data used for training GPT2+CoT and SEDD+DoT.

subgrid must contain the digits 1–4 exactly once. A reward of 1 is assigned only when the solution is entirely correct; otherwise, the reward is 0.

GSM8K-Aug Our experiments are conducted on the GSM8K-Aug dataset (sourced from [10]), which contains 384,623 training samples with simplified CoT processes. After filtering out samples with too long answers (exceeding 64 tokens) and those with invalid formats, we retained 382,553 examples. For the pretraining of SEDD on GSM8K-Aug, we train it for 120K steps using the Adam optimizer with a learning rate of 3×10^{-4} , a batch size of 512, and a generation length of 64 tokens. In the RL training phase, we employ the Adam optimizer with a learning rate of 5×10^{-5} and a batch size of 32, and set the group size for GRPO to 64. This RL phase is trained for 1,000 steps.

All training was conducted on 8 H800 GPUs. For GSM8K-Aug pre-training, we spent 160 GPU hours on pretraining for 120K steps, and additional 168 GPU hours for the RL post-training. Evaluation hyperparameters for both datasets are provided in Tab. 6.

C.2 LLaDOU (LLaDA+DCoLT)

C.2.1 Implementation Details for Math Problems

Datasets & Benchmarks We consider two popular benchmarks, GSM8K [8] and MATH [16]. In our experiments, we follow the common train-test split on these datasets. For GSM8K, there are 7.5K questions for training and 1.32K questions for testing. For MATH, there are 7.5K questions for training and 5K questions for testing. We report the accuracy on their test set. During training, we extract the final answers from generated responses to compute the rewards.

Reward function We assign a hard reward for each completion, that it is 1 only if the final answer is equivalent to the ground truth [16]. For GSM8K and MATH, following [36], the answer must appear inside a `\boxed{}`; we extract the boxed content, parse it into a number or expression, and check symbolic equivalence against the reference answer.

Training configurations The model is trained with 64 prompts in a batch, each generating 16 completions to form a group for advantage calculation. We take an AdamW optimizer with a learning rate of 5×10^{-6} , and $(\beta_1, \beta_2) = (0.9, 0.999)$. We do not apply the KL penalty by default, as it provides marginal benefits in our experiments. The whole training lasts for 140 iterations on 16 H800 GPUs, which takes about 63 GPU days (i.e., about 4 days on wall clock with 16 GPUs). We visualize the reward curves during training in Fig. 12.

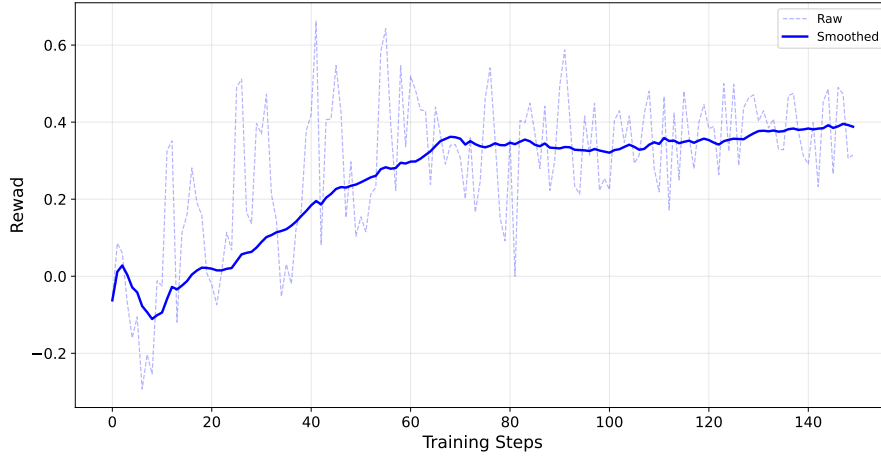


Figure 12: Reward curves on MATH and GSM8K during DCoLT training. Dashed lines denote the raw rewards, while solid lines represent the moving-average smoothed rewards, showing the overall trend.

C.2.2 Implementation Details for Code Generation

Dataset We filter the KodCode-V1-SFT-R1 [41] dataset to construct our training dataset. KodCode is a large fully-synthetic open-source dataset providing verifiable code solutions and test cases for coding tasks.

To ensure the quality and consistency of the training data, we retain only samples labeled with the "instruct" style and containing exactly one entry function for test with a non-empty docstring. We further exclude any samples where the provided solutions include class definitions or constructor methods, and ensure that the solutions contain exactly one function definition. Additionally, we only use samples from a group of subsets, including Prefill, Taco, Leetcode, Codeforces, Code Contests, and Filter, for their appropriate difficulty level.

After filtering, we obtain a training set with 48.9K samples (around 12M tokens). We then format the samples into the prompt template below:

You are an expert Python programmer. Your task is to complete the implementation of a function named ``<function_name>``.

```
** TARGET FUNCTION **
<docstring description>
```

```
** UNIT TESTS **
Your code should pass unit tests like:
    <assert statement 1>
    <assert statement 2>
    ...
```

```
Here is the function to complete:
```python
def <function_name>(<parameters>):
 """<docstring description>"""
...

```

**Reward function** We assign a hard 0/1 reward for each completion, that it is 1 only if the generated code block passes all provided test cases. The code blocks are extracted from the responses based on the markdown format.



**Training configurations** The training configurations are kept the same as for math. The training lasts for 240 iterations, and it takes about 127 GPU days on 16 H800 GPUs.

**Evaluation details** We evaluate the model in a zero-shot setting on MBPP [4] and HumanEval [7], and report the pass@1 metric. We show some examples of their input prompts below.

- **MBPP:**

You are an expert Python programmer. Your task is to complete the implementation of a function named `remove\_Occ`.

**\*\* TARGET FUNCTION \*\***

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

**\*\* UNIT TESTS \*\***

Your code should pass unit tests like:

```
assert remove_Occ("hello", "l") == "heo"
assert remove_Occ("abcda", "a") == "bcd"
assert remove_Occ("PHP", "P") == "H"
```

Here is the function to complete:

```
```python
def remove_Occ(input_param_1, input_param_2):
    """Write a python function to remove first and last occurrence of a
    given character from the string."""
    ...
```

- **Humaneval:**

You are an expert Python programmer. Your task is to complete the implementation of a function named `has_close_elements`.

Here is the function to complete:

```
```python
from typing import List
def has_close_elements(numbers: List[float], threshold: float) -> bool:
 """Check if in given list of numbers, any two numbers are closer to each
 other than the given threshold.
```

Examples:

```
>>> has_close_elements([1.0, 2.0, 3.0], 0.5)
False
>>> has_close_elements([1.0, 2.8, 3.0, 4.0, 5.0, 2.0], 0.3)
True
"""
...

```

### C.2.3 Direct Comparison with d1-LLaDA

A direct comparison with d1[47] using the same LoRA structure and dataset still reveals the superior performance by our method. Specifically, we train LLaDOU on GSM8K only, using LoRA with rank  $r = 128$  and scaling factor  $\alpha = 64$ . During inference, we set the generation length to 256, the number of diffusion steps to 128, and the block length to 32 — mirroring d1’s configuration. In the same setting, DCoLT with LoRA achieves an accuracy of 84.7%, which outperforms d1 with both diffu-GRPO (79.8%) and d1-LLaDA (81.1%) setups, as shown in Tab. 8. These results indicate that LLaDOU still performs better than d1 with the same LoRA structure and the training data.

### C.2.4 Applying DCoLT to Dream-7B on GSM8K

To further examine the generality of DCoLT, we apply it to the Dream-7B [43] model on GSM8K. Unlike LLaDA, which is a natively masked diffusion model, Dream-7B originates from an auto-regressive model.

Table 8: Comparison between d1 and LLaDOU with LoRA on GSM8K. Both methods only use the GSM8K to train the LLaDA model with their respective RL approach.

Model	Accuracy (%)
LLaDA 8B	78.3
+diffu-GRPO	79.8
d1-LLaDA (SFT + diffu-GRPO)	81.1
LLaDOU 8B (LoRA)	<b>84.7</b>

Using the same reinforcement learning configuration as in the main experiments (64 denoising steps and generation length  $L = 256$ ), DCoLT raises Dream-7B’s accuracy from 50.11% to 80.53%, an absolute gain of +30.42%. The improvement magnitude is comparable to that observed on LLaDA-8B, suggesting that DCoLT consistently enhances reasoning ability across heterogeneous diffusion architectures. The model is trained with a batch of 64 prompts, each generating 16 completions for advantage calculation, using AdamW with a learning rate of  $2.5 \times 10^{-6}$  and  $(\beta_1, \beta_2) = (0.9, 0.999)$ .

### C.3 Post-Training Data Used in Tab. 2

In Tab. 9, we list the data used in post-training for math and code generation tasks by different models. Note that some methods [14, 34, 36] involve a considerable large amount of proprietary math/code data for both pretraining and post-training of models, which makes crucial contributions to high accuracies.

Table 9: Post-training data for math / code generation used in Tab. 2. “-” means this type of post-training was not used for the model.

Method	SFT	RL
Auto-regressive Models		
LLaMA2 7B [39]	-	-
MetaMath 7B [45]	395K samples (MetaMathQA)	-
CodeLLaMA-Instruct 7B [34]	thousands of SFT + millions of RS examples	-
Deepseek 7B [5]	-	-
DeepseekMath-Instruct 7B [36]	776K (problem, solution) pairs	-
DeepseekMath-RL 7B [36]	776K (problem, solution) pairs	144K questions
DeepseekCoder-Instruct 7B [14]	2B tokens	-
Diffusion Language Models		
Dream 7B [43]	-	-
LLaDA 8B [27]	-	-
+ SFT [47]	1K samples from s1K	-
+ diffu-GRPO [47]	-	GSM8K: 7.5K / MATH: 7.5K
d1-LLaDA (SFT + diffu-GRPO)	1K samples from s1K	GSM8K: 7.5K / MATH: 7.5K
LLaDOU (LLaDA + DCoLT) 8B	-	Math: 7.5K GSM8K + 7.5K MATH / Code: 48K filtered from KodCode

## D Potential Scaling Law for Longer Generations

Fig. 13 illustrates some examples of generated responses of various lengths to the same prompt for the LLaDOU. In DLMs, the generation length tends to increase as more mask tokens are initialized at the beginning of the reverse diffusion process. This differs from auto-regressive language models, where the length of a generated response is determined by when the End of Text (EoT) token appears.

Although DLMs also use the EoT token to mark the end of the output, the token often appears later in the output sequence to generate a longer thinking process when more mask tokens are initialized.

**Answer (Generation Length=256):**

1. Set the equations equal to each other:  $(x^4 = 7x^2 - 10)$ .
2. Rearrange the equation:  $(x^4 - 7x^2 + 10 = 0)$ .
3. Let  $(u = x^2)$ , then  $(u^2 - 7u + 10 = 0)$ .
4. Solve the quadratic equation:  $(u = \frac{7 \pm \sqrt{49 - 40}}{2}) = \frac{7 \pm 3}{2}$ .  
-  $(u = 5)$  or  $(u = 2)$ .
5. Therefore,  $(x^2 = 5)$  or  $(x^2 = 2)$ .
6. The  $(x)$ -coordinates are  $(\pm \sqrt{5})$  and  $(\pm \sqrt{2})$ .
7. Identify  $(m = 5)$  and  $(n = 2)$ .
8. Calculate  $(m - n)$ :  $(m - n = 5 - 2 = 3)$ .

**Answer (Generation Length=384):**

1. Set the equations equal to each other:  
 $(x^4 = 7x^2 - 10)$
  2. Rearrange the equation:  
 $(x^4 - 7x^2 + 10 = 0)$
  3. Let  $(u = x^2)$ :  
 $(u^2 - 7u + 10 = 0)$
  4. Solve the quadratic equation:  
 $(u = \frac{7 \pm \sqrt{49 - 40}}{2}) = \frac{7 \pm 3}{2}$   
 $(u = \frac{7 \pm \sqrt{9}}{2})$   
 $(u = \frac{7 \pm 3}{2})$   
 $(u = 5 \text{ or } u = 2)$
  5. Substitute back to find  $(x)$ :  
 $(x^2 = 5 \text{ or } x^2 = 2)$   
 $(x = \pm \sqrt{5} \text{ or } x = \pm \sqrt{2})$
  6. Determine  $(m)$  and  $(n)$ :  
 $(m = 5 \text{ and } n = 2)$
  7. Calculate  $(m - n)$ :  
 $(m - n = 5 - 2 = 3)$
- 3**

**Answer (Generation Length=512):**

- Step 1: Set the equations equal to each other.  
 $(x^4 = 7x^2 - 10)$
- Step 2: Rearrange the equation to set it to zero.  
 $(x^4 - 7x^2 + 10 = 0)$
- Step 3: Let  $(u = x^2)$ . The equation becomes:  
 $(u^2 - 7u + 10 = 0)$
- Step 4: Solve the quadratic equation for  $(u)$ .  
Using the quadratic formula  $(u = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a})$ ,  
where  $(a = 1)$ ,  $(b = -7)$ , and  $(c = 10)$ :  
 $(u = \frac{7 \pm \sqrt{49 - 40}}{2}) = \frac{7 \pm 3}{2}$   
 $(u = \frac{7 \pm \sqrt{9}}{2})$   
 $(u = \frac{7 \pm 3}{2})$
- Step 5: Calculate the two solutions for  $(u)$ .  
 $(u = 5)$   
 $(u = 2)$
- Step 6: Find the corresponding  $(x)$  values.  
For  $(u = 5)$ :  
 $(x^2 = 5 \implies x = \pm \sqrt{5})$   
For  $(u = 2)$ :  
 $(x^2 = 2 \implies x = \pm \sqrt{2})$
- Step 7: Determine  $(m)$  and  $(n)$ .  
From  $(x^2 = 5)$ ,  $(m = 5)$ .  
From  $(x^2 = 2)$ ,  $(n = 2)$ .
- Step 8: Calculate  $(m - n)$ .  
 $(m - n = 5 - 2 = 3)$
- Final answer:  
**3**

**Answer (Generation Length=512 finetuned):**

- To find the points of intersection between the graphs of  $(y = x^4)$  and  $(y = 7x^2 - 10)$ , we set the equations equal to each other:  
 $(x^4 = 7x^2 - 10)$
- Rearrange the equation to form a quartic equation:  
 $(x^4 - 7x^2 + 10 = 0)$
- Let  $(u = x^2)$ . Then the equation becomes:  
 $(u^2 - 7u + 10 = 0)$
- We solve this quadratic equation using the quadratic formula  $(u = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a})$ :  
 $(u = \frac{7 \pm \sqrt{49 - 40}}{2}) = \frac{7 \pm 3}{2}$   
 $(u = \frac{7 \pm \sqrt{9}}{2})$   
 $(u = \frac{7 \pm 3}{2})$
- This gives us two solutions:  
 $(u = 5)$   
 $(u = 2)$
- Since  $(u = x^2)$ , we have:  
 $(x^2 = 5)$   
 $(x^2 = 2)$
- Thus, the  $(x)$ -coordinates are:  
 $(x = \pm \sqrt{5})$   
 $(x = \pm \sqrt{2})$
- We are given that  $(m > n)$ , so  $(m = 5)$  and  $(n = 2)$ . Therefore,  
 $(m - n = 5 - 2 = 3)$
- Thus, the answer is:  
**3**

(a) The graphs of  $y = x^4$  and  $y = 7x^2 - 10$  intersect at four points with  $x$ -coordinates  $\pm\sqrt{m}$  and  $\pm\sqrt{n}$ , where  $m > n$ . What is  $m - n$ ?

This offers a flexible way for DLMs to control various lengths of generated responses. As shown in Tab. 5, longer generations could improve performance, particularly when the model is fine-tuned by DCoLT with increasing generation length. This suggests the existence of a potential scaling law for longer generations.

Fig. 13b shows that some mistakes generated in the thinking process can be corrected in longer generations, thus achieving better results over shorter ones. We will study such a scaling phenomenon further in the future.

**Answer (Generation Length=256):**

First, we convert the numbers in base 5 to base 10 for easier calculation:  
 $-(4321_5 = 4 \cdot 5^3 + 3 \cdot 5^2 + 2 \cdot 5 + 1 = 276)$   
 $-(1234_5 = 1 \cdot 5^3 + 2 \cdot 5^2 + 3 \cdot 5 + 4 = 139)$   
 - Perform the subtraction:  $(276 - 139 = 137)$

Next, we convert 137 back to base 5:  
 $-(137 \div 5 = 27) \text{ remainder } (2)$   
 $-(27 \div 5 = 5) \text{ remainder } (2)$   
 $-(5 \div 5 = 1) \text{ remainder } (0)$   
 $-(1 \div 5 = 0) \text{ remainder } (1)$

Reading the remainders from bottom to top, we get  $(1042_5)$ .

Therefore, the answer in base 5 is  $(\boxed{3042_5})$

**Answer (Generation Length=384):**

Step 1: Convert both numbers to base 10.

For  $(4321_5)$ :

$$[4 \cdot 5^3 + 3 \cdot 5^2 + 2 \cdot 5 + 1 = 1000 + 75 + 10 + 1 = 1086]$$

For  $(1234_5)$ :

$$[1 \cdot 5^3 + 2 \cdot 5^2 + 3 \cdot 5 + 4 = 125 + 50 + 15 + 4 = 194]$$

Step 2: Perform the subtraction in base 10.

$$[1086 - 194 = 892]$$

Step 3: Convert the result back to base 5.

$$[892 \div 5 = 178 \text{ remainder } 2]$$

$$[178 \div 5 = 35 \text{ remainder } 3]$$

$$[35 \div 5 = 7 \text{ remainder } 0]$$

$$[7 \div 5 = 1 \text{ remainder } 2]$$

$$[1 \div 5 = 0 \text{ remainder } 1]$$

Reading the remainders from bottom to top, we get  $(1202_5)$ .

$(\boxed{3202_5})$

**Answer (Generation Length=512):**

Step 1: Convert  $(4321_5)$  to base 10.

$$[4321_5 = 4 \cdot 5^3 + 3 \cdot 5^2 + 2 \cdot 5 + 1 = 4 \cdot 125 + 3 \cdot 25 + 2 \cdot 5 + 1 = 500 + 75 + 10 + 1 = 586]$$

Step 2: Convert  $(1234_5)$  to base 10.

$$[1234_5 = 1 \cdot 5^3 + 2 \cdot 5^2 + 3 \cdot 5 + 4 = 1 \cdot 125 + 2 \cdot 25 + 3 \cdot 5 + 4 = 125 + 50 + 15 + 4 = 194]$$

Step 3: Subtract the base 10 equivalents.

$$[586 - 194 = 392]$$

Step 4: Convert 392 back to base 5.

$$[392 \div 5 = 78 \text{ remainder } 2]$$

$$[78 \div 5 = 15 \text{ remainder } 3]$$

$$[15 \div 5 = 3 \text{ remainder } 0]$$

$$[3 \div 5 = 0 \text{ remainder } 3]$$

Reading the remainders from bottom to top, we get  $(3032_5)$ .

$(\boxed{3032_5})$

**Answer (Generation Length=512 finetuned):**

First, we convert both numbers to base 10 for easier calculation.

$(4321_5)$ :

$$[4 \cdot 5^3 + 3 \cdot 5^2 + 2 \cdot 5 + 1 = 4 \cdot 125 + 3 \cdot 25 + 2 \cdot 5 + 1 = 500 + 75 + 10 + 1 = 586]$$

$(1234_5)$ :

$$[1 \cdot 5^3 + 2 \cdot 5^2 + 3 \cdot 5 + 4 = 1 \cdot 125 + 2 \cdot 25 + 3 \cdot 5 + 4 = 125 + 50 + 15 + 4 = 194]$$

Next, we perform the subtraction in base 10:

$$[586 - 194 = 392]$$

Now, we convert 392 back to base 5.

First, repeatedly divide by 5:

$$[392 \div 5 = 78 \text{ remainder } 2]$$

$$[78 \div 5 = 15 \text{ remainder } 3]$$

$$[15 \div 5 = 3 \text{ remainder } 0]$$

$$[3 \div 5 = 0 \text{ remainder } 3]$$

Reading the remainders from bottom to top, we get  $(3032_5)$ .

Thus, the final answer is:

$(\boxed{3032_5})$

(b) Find  $4321_5 - 1234_5$ . Express your answer in base 5.

Figure 13: Examples of how increasing generation lengths enable the DCoLT-reinforced model to unfold more complex reasoning processes. Incorrect steps are marked in red, while the corrected steps are highlighted in green. The generation length means the number of mask tokens initialized in the input sequence to the reverse diffusion process. Except for the model finetuned with 512 generation length, other models are only trained with 256 generation length as shown in experiments.

Table 10: Model accuracies with different block length on GSM8K dataset. LLaDOU (LLaDA + DCoLT) is trained with a generation length of 256 and 64 steps of reverse diffusion under respective block lengths.

Model	Block Length		
	8	32	256
LLaDA 8B	63.26	62.12	49.47
LLaDOU (LLaDA+DCoLT) 8B	83.17	83.11	82.27

## E Ablation on Various Block Lengths

LLaDA [27] divides the sequence into several blocks and generates responses blockwise from left to right in a semi-autoregressive manner. This policy introduces additional priors on the block-wise order of text generation. In Tab. 10, for the baseline LLaDA 8B model, setting block length to 8 achieves an accuracy of 63.26% on GSM8K, significantly higher than 49.47% without dividing blocks.

However, when we apply DCoLT to LLaDA with different block lengths, all LLaDOU models achieve similar performance, with the one with a block length 8 achieving slightly higher accuracy (83.17%) than the one with block length 256 without dividing blocks (82.27%). These results suggest that LLaDOU does not rely on the blocking prior as in the baseline model.

## F Licenses for existing assets

The code and models associated with this paper will be released to the public later. For all code and data assets, we refer to their licenses in Tab. 11.

Table 11: Reference assets and their licenses.

Asset	License	Utility
SEDD [24]	MIT	Code & Model
GSM8K-Aug [10]	-	Data
LLaDA [27]	MIT	Code & Model
MATH [16]	MIT	Data
GSM8K [8]	MIT	Data
KodCode [41]	CC BY-NC 4.0	Data

## G Broader Impacts

Our work introduces a novel perspective of lateral reasoning in language models, in contrast to the widely adopted CoT paradigms used in LLMs. By enabling non-linear, format-free reasoning through a reverse diffusion process, our approach offers new insights into the mechanisms of reasoning in generative models. This perspective has the potential to benefit a broad range of reasoning-intensive tasks, including embodied AI, autonomous agents, and complex decision-making systems, providing transparent thinking processes to minimize the hallucination risks that could emerge from inadequate reasoning of language models.

However, we acknowledge that DCoLT could pose potential unexpected societal impacts if misused, especially when handling misleading or adversarial prompts. In this work, we focus on math and code generation tasks, in which outputs are objectively verifiable. With such a reward design, we can reduce the influence of dataset biases and encourage the development of reasoning behaviors that are aligned with the factual and logical consistency. Handling subjective preferences in rewarding the models to train the DCoLT could be more challenging, and we leave it to our future works.

## **H Safeguards**

DCoLT is designed for math and code generation tasks that involve objectively verifiable outputs and well-defined correctness criteria. These domains present a relatively low risk of societal misuse compared to open-ended language generation tasks. To further reduce the potential for unintended use, we recommend deploying DCoLT alongside an input checker that ensures that the model only processes input relevant to its target domains. This approach helps mitigate the risks related to adversarial prompting or misuse beyond the intended scope.