

RFG: TEST-TIME SCALING FOR DIFFUSION LARGE LANGUAGE MODEL REASONING WITH REWARD-FREE GUIDANCE

Anonymous authors

Paper under double-blind review

ABSTRACT

Diffusion large language models (dLLMs) have shown great potential in large-scale language modeling, and there is an increasing interest in further improving the capacity to solve complex problems by guiding the reasoning process step by step. Common practice for autoregressive language models typically learns a process reward model with dense annotation for each intermediate step. However, this is challenging for dLLMs where the generation is in an any-order fashion and intermediate states are partially masked sentences. To this end, in this paper, we propose reward-free guidance (RFG), a principled method for guiding the reasoning trajectory of dLLMs without explicit process reward. The key idea of RFG is to parameterize the process reward by log-likelihood ratios of the enhanced and reference dLLMs, where the enhanced model can be easily obtained by any off-the-shelf dLLM that has been post-trained with reinforcement learning (RL) or supervised fine-tuning (SFT). We provide theoretical justification that RFG induces the reward-guided sampling distribution with no additional reward. We conduct comprehensive experiments on four challenging mathematical reasoning and code generation benchmarks using a diverse suite of dLLMs enhanced with various post-training methods. RFG consistently yields significant improvements across all tasks and model types, achieving accuracy gains of up to 9.2%. These findings establish RFG as a general training-free framework that scales test-time reasoning without reliance on external reward models.

1 INTRODUCTION

Recently, diffusion large language models (dLLMs) have achieved remarkable progress in language modeling (Sohl-Dickstein et al., 2015; Austin et al., 2021a; Campbell et al., 2022; Meng et al., 2022; Lou et al., 2024; Sahoo et al., 2024; Shi et al., 2024; Xu et al., 2025), ranging from mathematical reasoning and planning (Gong et al., 2023; Zhao et al., 2025) to expert coding (Gong et al., 2025b). By scaling up mask-predict pretraining on large-scale corpora through bidirectional computation, dLLMs have shown surprisingly competitive or even superior performance over autoregressive (AR) model baselines (Prabhudesai et al., 2025). Despite the impressive advancements, the current success of dLLMs is primarily limited to pre-training or continue-training on a specific domain, with limited exploration in test-time computation and alignment. The test-time scaling paradigm has shown great promise in AR-based LLMs (Brown et al., 2024; Muennighoff et al., 2025), which allows significant performance gain without further expensive supervised fine-tuning or reinforcement learning (RL). Therefore, developing an efficient and effective dLLM test-time computation method remains a promising but underexplored direction.

Test-time computation aims to mimic human reasoning behavior that thinks longer for more complicated problems to obtain more reliable results (Lightman et al., 2023; Wang et al., 2023; Snell et al., 2024). Formally, given a difficult question, we want to design algorithms that can effectively utilize additional computation purely at test time to improve the output accuracy. Common practices in AR LLMs typically involve an additional *reward* model, among them, the most native approach is the *best-of-N* sampling (Cobbe et al., 2021; Lightman et al., 2023). This approach first samples N responses in parallel from the LLM and then selects the best one, based on the highest reward rating provided by the reward model. Later approaches improve the method by training a fine-grained

process reward for each reasoning step, and then guide the reasoning processes by sequentially alternating between generating one reasoning step from the base LLM and searching against the reward model (Wang et al., 2023; Lu et al., 2024). However, there has been limited effort in developing advanced test-time guided computation for dLLMs. While the community has studied naive solutions analogous to the best-of-N algorithm (Wang et al., 2025a; Singhal et al., 2025; Dang et al., 2025), it is challenging to adopt process reward models (PRMs) for dLLMs and guide the denoising process step by step. The main obstacle arises from the inherent any-order sampling nature of dLLMs, resulting in intermediate steps that are incomplete text with partially masked sentences.

In this paper, we view the test-time scaling for dLLMs as a guided sampling problem, where we apply guidance on each denoising step towards accurate reasoning traces. To this end, we propose reward-free guidance (RFG), a principled framework to guide the dLLM reasoning with an implicitly parameterized PRM. Our key innovation is that, instead of training an explicit PRM on incomplete intermediate generations, we specially parameterize an outcome reward model (ORM) with pretrained dLLMs and freely decompose it to PRMs for each diffusion step. Specifically, we formulate the reward model as the log-likelihood ratio of the policy and reference models $r_\theta(\mathbf{x}) = \beta \log \frac{p_\theta(\mathbf{x})}{p_{\text{ref}}(\mathbf{x})}$, with both implemented as dLLMs. Such reparameterization has been widely adopted in various RL and preference optimization literature. Intuitively, the log-likelihood ratio $\log \frac{p_\theta(\mathbf{x})}{p_{\text{ref}}(\mathbf{x})}$ provides a natural and principled reward signal, as it inherently measures how much a generation \mathbf{x} is preferred by the policy model (p_θ) over the reference model (p_{ref}). Then the PRMs for each denoising step t follow the same log-likelihood ratio form, but are calculated on partially observed responses $r_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \beta \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)}{p_{\text{ref}}(\mathbf{x}_{t-1}|\mathbf{x}_t)}$. Interestingly, our RFG framework also resembles the idea of the well-studied classifier-free guidance (CFG) method in diffusion models, where CFG uses conditional and unconditional models, while we use policy and reference models to guide the sampling process. We theoretically show that our step-wise RFG can yield a reward-reweighted target distribution for complete generations. Furthermore, we highlight that our proposed parameterization of the diffusion PRM is general and can be learned by any off-the-shelf dLLM instruction tuning or RL methods, without specific reliance on any training objectives.

To validate our approach, we conduct extensive experiments on challenging mathematical reasoning and code generation benchmarks. Using publicly available dLLMs like LLaDA (Nie et al., 2025b) and Dream (Ye et al., 2025), along with their instruction-tuned and RL-enhanced variants, RFG demonstrates consistent and significant performance gains across different benchmarks and various model types. Our method not only surpasses the original post-trained models but also outperforms a naive ensemble baseline with an identical compute budget, confirming that the improvements stem from our principled guidance framework rather than mere computational scaling. These results highlight the effectiveness and robustness of RFG as a powerful test-time scaling for enhancing dLLM reasoning capabilities.

2 PRELIMINARIES

Discrete Diffusion Models. Diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2020) are a class of generative models that learn a data distribution by reversing a progressive noising process. Originally developed for continuous data like images, the forward process gradually perturbs a clean sample into increasingly noisy states, eventually transforming it into pure noise. A neural network is then trained to reverse this process. Adapting this paradigm to discrete data such as text is non-trivial, as adding small amounts of Gaussian noise is not well-defined for discrete tokens. This led to the development of discrete diffusion models (Austin et al., 2021a), which define the forward process using a Markov chain that gradually randomizes discrete tokens, often converging to a uniform distribution. An effective and intuitive case of discrete diffusion is the masked diffusion model (Campbell et al., 2022; Lou et al., 2024; Shi et al., 2024; Sahoo et al., 2024). In this framework, the "noise" is a special [MASK] token. The forward process progressively masks tokens in the input sequence until the entire sequence is masked. The reverse process then learns to predict the original tokens given a partially masked sequence. This formulation naturally leverages architectures and training objectives similar to masked language modeling (MLM) (Devlin et al., 2019), while inheriting the iterative refinement of diffusion.

Diffusion Large Language Models. Diffusion large language models (dLLMs) (Nie et al., 2025b; Ye et al., 2025; Inception Labs et al., 2025; DeepMind, 2025) scale masked diffusion models to large corpora and long reasoning sequences, serving as an alternative to autoregressive (AR) language models. The forward process of a dLLM takes the original sequence \mathbf{x}_0 as input and progressively masks it following the distribution

$$q(\mathbf{x}_t|\mathbf{x}_0) = \prod_{i=0}^L q(\mathbf{x}_t^{(i)}|\mathbf{x}_0^{(i)}), \quad q(\mathbf{x}_t^{(i)}|\mathbf{x}_0^{(i)}) = \begin{cases} 1 - \alpha_t, & \mathbf{x}_t^{(i)} = [\text{MASK}] \\ \alpha_t, & \mathbf{x}_t^{(i)} = \mathbf{x}_0^{(i)} \end{cases}$$

where $t \in [0, 1]$ is the timestep, L is the sequence length, and α_t denotes a noise schedule that decreases monotonically with t , satisfying $\alpha_0 = 1$ and $\alpha_1 = 0$. At timestep t , the ratio of masked token is $1 - \alpha_t$. A key property of dLLMs is that token prediction is not tied to a fixed left-to-right ordering. At each reverse step, the model can unmask any subset of masked positions conditioned on the visible context, enabling *any-order generation* in contrast to the strictly sequential nature of AR LLMs. Any-order generation reduces exposure bias, allows parallelized inference, and supports iterative refinement, making dLLMs well-suited for reasoning tasks that benefit from revisiting or correcting intermediate states. The detailed formulation of dLLMs is provided in Appendix B.

3 REWARD FREE GUIDANCE (RFG)

In this section, we elaborate on the details of our proposed RFG framework. We show how the reward model for the denoising trajectory can be parameterized as the log-density ratio of two dLLMs, and then how PRM for each denoising step can be freely derived from the reward without additional training. Existing PRMs in AR LLM literature typically require fine-grained step labels, and then the PRM is trained to predict the quality of partially generated responses (Lightman et al., 2023; Wang et al., 2023; Lu et al., 2024). However, such labels are expensive to collect, and even impossible in the context of dLLMs since the intermediate generations are typically partial sentences with tokens masked in random positions.

To this end, in this paper, we propose to *obtain a trajectory reward with reparameterization, and then freely derive PRM for each denoising step by decomposing the trajectory level reward*. Formally, we have the following theoretical justification:

Proposition 3.1. *Given a diffusion trajectory-level reward that is parameterized as the log-likelihood ratio of two dLLMs, i.e., $r_\theta(\mathbf{x}_{0:T}) := \beta \log \frac{p_\theta(\mathbf{x}_{0:T})}{p_{\text{ref}}(\mathbf{x}_{0:T})}$. Define $Q_\theta^t(\mathbf{x}_{t-1}, \mathbf{x}_{t:T}) := \sum_{i=t}^T \beta \log \frac{p_\theta(\mathbf{x}_{i-1}|\mathbf{x}_{i:T})}{p_{\text{ref}}(\mathbf{x}_{i-1}|\mathbf{x}_{i:T})}$, then we have that q_θ^t is the expectation of exponential r_θ at step t :*

$$Q_\theta^t(\mathbf{x}_{t-1}, \mathbf{x}_{t:T}) = \beta \log \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{0:T}|\mathbf{x}_{t-1:T})} e^{\frac{1}{\beta} r_\theta(\mathbf{x}_{0:T})}, \quad (1)$$

where β is a hyperparameter for weighting the reward function.

The full proof of the proposition is provided in Appendix A. The theoretical justification indicates that when parameterizing $r_\theta(\mathbf{x}_{0:T})$ as the log-likelihood ratio, Q_θ^t can be viewed as Q function representing the expectation of the overall trajectory reward $r_\theta(\mathbf{x}_{0:T})$ at step t . Following the classic setup in RL literature, the step reward r_θ^t can be written as:

$$r_\theta^t(\mathbf{x}_{t-1}|\mathbf{x}_t) = Q_\theta^t - Q_\theta^{t+1} = \beta \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_{t:T})}{p_{\text{ref}}(\mathbf{x}_{t-1}|\mathbf{x}_{t:T})} = \beta \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)}{p_{\text{ref}}(\mathbf{x}_{t-1}|\mathbf{x}_t)}. \quad (2)$$

Notably, with such PRMs, we can then conduct reward-guided sampling from dLLMs by reweighted denoising transitions, where the log probability can be written as:

$$\begin{aligned} \log p^*(\mathbf{x}_{t-1}|\mathbf{x}_t) &= \log p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) e^{\frac{1}{\gamma} r_\theta^t(\mathbf{x}_{t-1}|\mathbf{x}_t)} + C \\ &= (1+w) \log p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) - w \log p_{\text{ref}}(\mathbf{x}_{t-1}|\mathbf{x}_t) + C, \end{aligned} \quad (3)$$

where $w = \frac{\beta}{\gamma}$ is a new hyperparameter to control the guidance strength and C is a normalizing constant. Trivially, setting $w = -1$ or $w = 0$ can recover the exact sampling of original p_{ref} or p_θ , respectively, with $w > 0$ acting an over-emphasizing behavior of p_θ . Such formulation provides a vital insight: by simply parameterizing a reward model on the whole diffusion sampling trajectory, we can freely obtain PRMs for each denoising step without any additional training. In the later paragraphs, we will discuss the connections between our RFG framework and other diffusion sampling guidance methods, and provide explanations on how we can obtain the trajectory reward model using any off-the-shelf RL or instruction-tuned checkpoints.

Algorithm 1 Sampling with Reward-Free Guidance (RFG)

Require: Reference model p_{ref} , policy model p_{θ} , guidance strength w , query q , answer length L , sampling steps N , denoising strategy \mathcal{S}

- 1: Initialize $\mathbf{x}_N \leftarrow [\text{MASK}]^L$ ▷ Start with a fully masked sequence of length L
- 2: **for** $t \leftarrow N$ **down to** 1 **do**
- 3: $\log \pi_{\text{ref}} \leftarrow \text{Logits}(p_{\text{ref}}(\cdot | \mathbf{x}_t, q))$ ▷ Get logits from reference model
- 4: $\log \pi_{\theta} \leftarrow \text{Logits}(p_{\theta}(\cdot | \mathbf{x}_t, q))$ ▷ Get logits from enhanced model
- 5: $\log \pi_{\text{RFG}} \leftarrow (1 + w) \log \pi_{\theta} - w \log \pi_{\text{ref}}$ ▷ Combine logits via RFG
- 6: $\mathbf{x}_{t-1} \leftarrow \mathcal{S}(\mathbf{x}_t, \log \pi_{\text{RFG}}, t/N)$ ▷ Generate the next state using the guided logits
- 7: **end for**
- 8: **return** \mathbf{x}_0

Connections to diffusion guidance method. A widely adopted technique for guiding diffusion sampling is called *classifier-free guidance* (CFG) (Ho & Salimans, 2022), which pushes samples towards high class-confidence regions by reweighted denoising steps. In this section, we revisit our proposed RFG from the CFG perspective, and provide a holistic view of their connections.

Generally, CFG guidance in diffusion models involves two diffusion models $p_{\text{unconditional}}(\mathbf{x}; \mathbf{c})$ and $p_{\text{conditional}}(\mathbf{x})$. In the original CFG in Gaussian diffusion, the guidance is achieved by *extrapolating* between the scores of two denoising models with a weight w :

$$\nabla_{\mathbf{x}} \log p^*(\mathbf{x}|\mathbf{c}) = \nabla_{\mathbf{x}} \log p_{\text{conditional}}(\mathbf{x}|\mathbf{c}) + w \nabla_{\mathbf{x}} \log \frac{p_{\text{conditional}}(\mathbf{x}|\mathbf{c})}{p_{\text{unconditional}}(\mathbf{x})}. \quad (4)$$

The effect of such guidance is standard sampling from $p_{\text{conditional}}$ plus a drifting term to shift the sampling direction (for $w > 0$) towards the ratio of $p_{\text{conditional}}$ and $p_{\text{unconditional}}$. This ratio comes from the Bayesian rule that:

$$p^*(\mathbf{x}|\mathbf{c}) \propto p_{\text{conditional}}(\mathbf{x}|\mathbf{c}) \cdot p(\mathbf{c}|\mathbf{x})^w, \quad \text{where } p(\mathbf{c}|\mathbf{x}) \propto \frac{p_{\text{conditional}}(\mathbf{x}|\mathbf{c})}{p_{\text{unconditional}}(\mathbf{x})}. \quad (5)$$

The key idea of CFG is to reparameterize a hypothetical classifier as the likelihood ratio of conditional and unconditional diffusion models, which encourages the model to draw samples from density $p_{\text{conditional}}$ over $p_{\text{unconditional}}$. From this perspective, there comes a clear connection between RFG and CFG: CFG guides the sampling with classifier $p(\mathbf{c}|\mathbf{x})$ reparameterized as $\frac{p_{\text{conditional}}(\mathbf{x}|\mathbf{c})}{p_{\text{unconditional}}(\mathbf{x})}$, while RFG reweights the sampling with the reward $r(\mathbf{x})$ that is reparameterized as $\frac{p_{\theta}(\mathbf{x})}{p_{\text{ref}}(\mathbf{x})}$.

Implementation of r_{θ} . We emphasize that Proposition 3.1 is general and agnostic to any training method of the reward model. Specifically, by off-the-shelf dLLMs fine-tuned with RL and preference optimization methods, we have the conclusion that the optimal policy will converge to $p_{\theta}(\mathbf{x}) = p_{\text{ref}}(\mathbf{x}) e^{\frac{1}{\beta} r(\mathbf{x})}$ after the training (Peters & Schaal, 2007). Directly rearranging the converged policy, we have that $r_{\theta}(\mathbf{x}) = \beta \log \frac{p_{\theta}(\mathbf{x})}{p_{\text{ref}}(\mathbf{x})}$, which has already met our needs in guided sampling. Therefore, in practice, we can take any pair of a pretrained dLLM as the reference p_{ref} and an enhanced one via RL on certain tasks as the policy p_{θ} , and conduct guided sampling via RFG framework. Interestingly, from an empirical perspective, we also observe that even without explicit RL, taking an instruction fine-tuned model as p_{θ} can also offer significant performance gain, showing the generalization of RFG beyond our theoretical form.

Sampling with RFG. The practical implementation of RFG sampling involves a straightforward modification of the standard dLLM sampling loop at test time. The core idea is to compute the logits from both the policy model p_{θ} and the reference model p_{ref} , and then combine them using the guidance strength w towards the guided logit distribution. This guided distribution is then used to determine the next denoised state \mathbf{x}_{t-1} . A complete process is detailed in Algorithm 1. The denoising strategy \mathcal{S} is a generic function representing two key schedules in the dLLM denoising process. First, it determines the unmask schedule, deciding *which* and *how many* [MASK] tokens in \mathbf{x}_t to recover based on the current timestep t and the logits for all masked positions. Second, for the chosen positions, it selects the new tokens based on a chosen decoding strategy (e.g., greedy sampling, nucleus sampling) applied to the logits at the chosen position. This modularity allows RFG to be

combined with any underlying dLLM sampling implementation. We visualize the sampling process in Figure 1.

A unique advantage of our framework is that it is fully model-agnostic and training-free. In practice, p_θ can be any off-the-shelf RL-enhanced or instruction-tuned model. By employing RFG at test time, we can significantly improve performance over the policy model itself without any additional training. Furthermore, researchers can train their own specialized enhanced models and still apply RFG as a final, test-time step to further boost performance, demonstrating the flexibility and broad applicability of our approach.

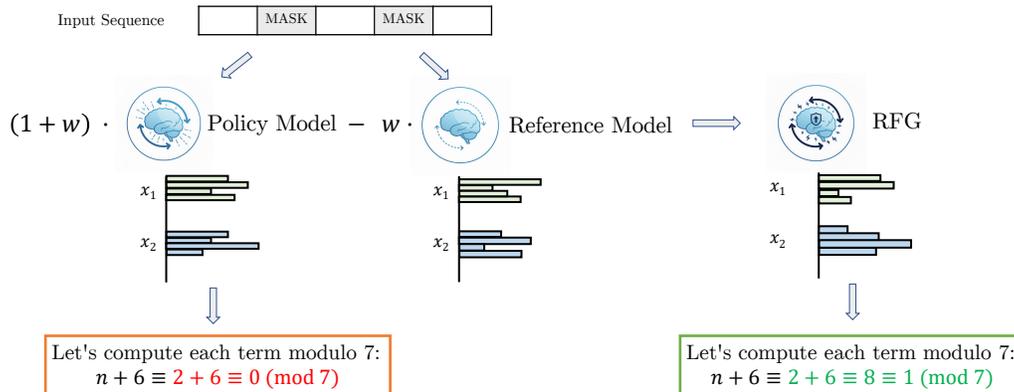


Figure 1: Sampling illustration for original policy model and RFG.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Tasks. We evaluate our method on four challenging reasoning benchmarks spanning two categories: **mathematical reasoning** and **code generation**. For mathematical reasoning, we use GSM8K (Cobbe et al., 2021), a benchmark of grade-school arithmetic word problems that requires multi-step symbolic reasoning, and MATH-500 (Hendrycks et al.; Lightman et al., 2023), a curated set of challenging competition-level mathematics problems. For code generation, we use HumanEval (Chen et al., 2021), which contains handwritten Python programming problems described in docstrings, and the MBPP (Austin et al., 2021b) (sanitized), consisting of everyday Python tasks with natural language prompts and associated unit tests.

Models. We leverage two state-of-the-art families of dLLMs: LLaDA (Nie et al., 2025b) and Dream (Ye et al., 2025). We use **LLaDA-Base** and **Dream-Base** as the reference model, respectively. For each family, we use two categories of post-trained models as the policy model: instruction fine-tuned and RL-enhanced. For instruction fine-tuned models, we employ **LLaDA-Instruct** and **Dream-Instruct**. For RL-enhanced models, we include **d1-LLaDA** (Zhao et al., 2025), which applies GRPO (Shao et al., 2024) to enhance LLaDA on mathematical and logical tasks; **LLaDA-1.5** (Zhu et al., 2025), which introduces VRPO techniques to reduce variance when applying DPO (Rafailov et al., 2023) to LLaDA; and **DiffuCoder** (Gong et al., 2025b), which builds on the Dream backbone and leverages coupled GRPO for code generation tasks.

Baselines. We compare our RFG, against two sets of baselines. The first is the original post-trained model applied directly for inference. The second baseline is a naive ensemble, which computes the final logits by taking the average of the logits from the post-trained model and the reference base model. This naive ensemble baseline is designed to match the exact compute budget of our RFG and serves as a controlled ablation.

Evaluation. For all benchmarks, we evaluate models under a zero-shot setting to assess their intrinsic reasoning capabilities without task-specific examples. We report accuracy for math reasoning

tasks and pass@1 for code generation tasks. All results are reported on the official test sets of each benchmark. We use official checkpoints for all models whenever publicly available; for d1-LLaDA, whose checkpoint has not yet been released, we reproduce the model with the official source code. To ensure a fair comparison, all baselines are implemented and evaluated under the identical inference setting with the same hyperparameters. Additional implementation details and hyperparameters are provided in Appendix C.

4.2 MAIN RESULTS

The results for mathematical reasoning and code generation are presented in Table 1 and Table 2, respectively. Across all benchmarks and for every post-trained model, our RFG consistently and significantly outperforms both the original post-trained model and the naive ensemble baseline. The improvement over the original post-trained models demonstrates that RFG is highly effective at guiding the generation process towards more accurate and logically sound reasoning traces. Furthermore, the substantial gains over the naive ensemble baseline with identical computational budget confirms that the performance gain is not merely from incorporating the reference model and additional compute, but from our principled guidance formulation.

Table 1: Performance on GSM8K and MATH-500. We report the accuracy of the original model, the naïve ensemble baseline, and RFG, along with the relative gain over the original model.

Model	GSM8K				MATH-500			
	Original	Ensemble	RFG	Rel. Gain	Original	Ensemble	RFG	Rel. Gain
Instruction Fine-tuning								
LLaDA 8B Instruct	79.8	68.0	81.3	+1.5	38.4	33.6	39.6	+1.2
Dream 7B Instruct	80.9	68.3	82.1	+1.2	43.6	30.6	46.4	+2.8
Reinforcement Learning								
d1-LLaDA	82.5	68.5	84.7	+2.2	38.8	34.8	41.6	+2.8
LLaDA 1.5	81.6	70.9	82.1	+0.5	42.4	40.6	44.4	+2.0

Table 2: Performance on HumanEval and MBPP. We report the pass@1 accuracy of the original model, the naïve ensemble baseline, and RFG, along with the relative gain over the original model.

Method	HumanEval				MBPP			
	Original	Ensemble	RFG	Rel. Gain	Original	Ensemble	RFG	Rel. Gain
Instruction Fine-tuning								
LLaDA 8B Instruct	45.1	40.9	47.6	+2.5	49.8	49.8	52.9	+3.1
Dream 7B Instruct	62.2	58.5	65.2	+3.0	64.6	69.6	70.4	+5.8
Reinforcement Learning								
diffuCoder	69.5	73.2	78.7	+9.2	72.8	70.8	74.3	+1.5

We also conduct experiments to demonstrate RFG’s effectiveness and generality across different generation settings. We use different numbers of denoising steps (128, 256, 512), with results shown in Table 3, 4, 5, and 6. The results demonstrate that RFG achieves consistent improvement over baselines across different number of denoising steps on all benchmarks and for all models. We also double the number of token per denoising step in Table 1 and 2, with results shown in Table 7 and 8. The results demonstrate that RFG achieves consistent improvement on multi-token prediction setup, strengthening its effectiveness and applicability.

4.3 ANALYSIS AND DISCUSSION

Qualitative Analysis. We present qualitative examples from all models on every benchmark to illustrate how RFG improves reasoning quality in the main paper as well as in Appendix D. On mathematical reasoning tasks, RFG produces more coherent multi-step derivations and avoids hallucination or inconsistent conclusions. On code generation tasks, RFG generates code that is not only syntactically correct but also more robust, reducing common errors such as missing edge conditions

Table 3: Performance on GSM8K at different denoising steps. Best in **bold**.

Model	# Steps = 128			# Steps = 256			# Steps = 512		
	Original	Ensemble	RFG	Original	Ensemble	RFG	Original	Ensemble	RFG
Instruction Fine tuning									
LLaDA 8B Instruct	76.3	39.3	77.2	79.8	68.0	81.3	84.1	80.1	84.7
Dream 7B Instruct	67.9	27.7	74.2	80.9	68.3	82.1	84.8	84.0	85.1
Reinforcement Learning									
d1 LLaDA	79.3	42.4	79.8	82.5	68.5	84.7	82.8	80.1	84.7
LLaDA 1.5	76.3	45.3	78.0	81.6	70.9	82.1	84.7	80.5	85.4

Table 4: Performance on MATH-500 at different denoising steps. Best in **bold**.

Model	# Steps = 128			# Steps = 256			# Steps = 512		
	Original	Ensemble	RFG	Original	Ensemble	RFG	Original	Ensemble	RFG
Instruction Fine tuning									
LLaDA 8B Instruct	32.4	22.8	34.2	38.4	33.6	39.6	42.6	37.6	43.6
Dream 7B Instruct	32.2	11.6	33.2	43.6	30.6	46.4	50.4	51.2	52.6
Reinforcement Learning									
d1 LLaDA	34.0	25.2	37.4	38.8	34.8	41.6	41.6	38.8	45.4
LLaDA 1.5	32.0	23.0	34.0	38.6	34.2	39.6	42.4	40.6	44.4

or incomplete logic. This qualitative case study corroborates our quantitative findings, illustrating how RFG effectively corrects and refines the reasoning trajectory during inference.

Query (GSM8K)

Christina is planning a birthday party and needs .75 gift bags per invited guest, because 1/4 of attendees don't show up. She invited 16 friends. Gift bags are \$2 each. How much will she spend?

Original (LLaDA 1.5)

To determine how much Christina will spend on gift bags, we need to follow these steps:

1. Calculate the number of friends who don't show up. 2. Determine the number of gift bags needed. 3. Calculate the total cost of the gift bags.

First, let's find out how many friends don't show up. Christina invited 16 friends, and 1/4 of them don't show up:

$$\frac{1}{4} \times 16 = 4$$

So, 4 friends don't show up.

Next, we calculate the number of friends who actually show up:

$$16 - 4 = 12$$

Since Christina needs 0.75 gift bags per invited guest, the number of gift bags needed is:

$$0.75 \times 12 = 9$$

Finally, we calculate the total cost of the gift bags. Each gift bag costs \$2, so the total cost is:

$$9 \times 2 = 18$$

Therefore, the amount Christina will spend is:

18

RFG (LLaDA 1.5)

To determine how much Christina will spend on gift bags, we need to follow these steps:

1. Calculate the number of friends who don't show up. 2. Determine the number of gift bags needed based on the actual attendees. 3. Calculate the total cost of the gift bags.

First, we calculate the number of friends who don't show up. Since 1/4 of the attendees don't show up and she invited 16 friends:

$$\frac{1}{4} \times 16 = 4$$

So, 4 friends don't show up.

Next, we determine the number of gift bags needed based on the actual attendees. Christina needs 0.75 gift bags per invited guest, so for 16 invited guests:

$$0.75 \times 16 = 12$$

So, she needs 12 gift bags.

Finally, we calculate the total cost of the gift bags. Each gift bag costs \$2, so:

$$12 \times 2 = 24$$

Therefore, the total amount Christina spends is:

24

378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431

Table 5: Performance on HumanEval at different denoising steps. Best in **bold**.

Model	# Steps = 128			# Steps = 256			# Steps = 512		
	Original	Ensemble	RFG	Original	Ensemble	RFG	Original	Ensemble	RFG
Instruction Fine tuning									
LLaDA 8B Instruct	32.3	18.3	33.5	39.6	37.2	40.9	45.1	40.9	47.6
Dream 7B Instruct	52.4	48.2	60.4	61.0	59.1	66.5	62.2	58.5	65.2
Reinforcement Learning									
diffuCoder	70.1	64.6	73.8	69.5	72.6	76.2	69.5	73.2	78.7

Table 6: Performance on MBPP at different denoising steps. Best in **bold**.

Model	# Steps = 128			# Steps = 256			# Steps = 512		
	Original	Ensemble	RFG	Original	Ensemble	RFG	Original	Ensemble	RFG
Instruction Fine tuning									
LLaDA 8B Instruct	51.0	52.5	54.1	54.9	51.4	55.6	49.8	49.8	52.9
Dream 7B Instruct	65.0	65.8	68.1	64.2	67.3	68.5	64.6	69.6	70.4
Reinforcement Learning									
diffuCoder	72.4	70.0	73.9	72.4	72.8	73.5	72.8	70.8	74.3

```
Query (MBPP)

Write a function to return two words from a list of words starting with letter 'p'. Your code should pass the test:

assert start_withp(["Python PHP", "Java JavaScript", "c c++"])==('Python', 'PHP')
```

```
Original (DiffuCoder)

def start_withp(lst):
    result=[]
    for i in lst:
        x=i.split()
        for j in x:
            if j[0]=='p':
                result.append(j)
    return tuple(result[:2])
```

```
RFG (DiffuCoder)

def start_withp(lst):
    result= []
    for i in lst:
        for j in i.split() :
            if j[0].lower()=='p' :
                result.append(j)
    return tuple(result[:2])
```

Sensitivity to Guidance Strength. To elucidate the impact of the guidance strength on our model’s performance, we conducted a sensitivity analysis by varying the hyperparameter w . The results of this investigation are presented in Figure 2, which plots model accuracy as a function of the guidance strength across all four evaluation benchmarks. The findings clearly demonstrate the robustness of our RFG framework. Crucially, this improvement is not confined to a narrow, fine-tuned peak; instead, the model exhibits a wide plateau of strong performance across a broad range of w values. This indicates that substantial gains can be achieved without intensive hyperparameter tuning. This characteristic is a significant practical advantage, making RFG a reliable and easy-to-implement enhancement for dLLM reasoning.

Interpreting the Guidance Direction and Strength. The RFG formulation, $\log \pi_{\text{RFG}} = \log \pi_{\theta} + w(\log \pi_{\theta} - \log \pi_{\text{ref}})$, can be interpreted as steering the enhanced model’s distribution, $\log \pi_{\theta}$, along an optimization direction defined by the difference $(\log \pi_{\theta} - \log \pi_{\text{ref}})$ with strength w . While RFG is motivated by theoretical derivation on RL models, our empirical results show RFG works well for SFT models too, and this interpretation helps explain why. We also note that in most cases we use $w > 0$ to push generation further along this beneficial direction, though a negative strength may also be optimal. The post-trained checkpoints are optimized jointly for multiple tasks, and they may perform sub-optimally on a specific benchmark due to over-optimization. In such cases, the optimization direction is still valuable, but the enhanced model has moved too far along it. By

Table 7: Performance on GSM8K and MATH-500, doubling the number of token per denoising step in Table 1.

Model	GSM8K				MATH-500			
	Original	Ensemble	RFG	Rel. Gain	Original	Ensemble	RFG	Rel. Gain
Instruction Fine-tuning								
LLaDA 8B Instruct	78.6	64.3	80.7	+2.1	34.2	29.4	38.0	+3.8
Dream 7B Instruct	68.5	58.6	73.6	+5.1	35.4	27.4	37.6	+2.2
Reinforcement Learning								
d1 LLaDA	79.8	67.3	82.0	+2.2	36.6	29.4	39.2	+2.6
LLaDA 1.5	81.5	70.2	82.1	+0.6	38.6	37.2	41.4	+2.8

Table 8: Performance on HumanEval and MBPP, doubling the number of token per denoising step in Table 2.

Method	HumanEval				MBPP			
	Original	Ensemble	RFG	Rel. Gain	Original	Ensemble	RFG	Rel. Gain
Instruction Fine-tuning								
LLaDA 8B Instruct	37.2	32.9	38.4	+1.2	43.6	39.7	44.4	+0.8
Dream 7B Instruct	37.8	39.0	42.7	+4.9	48.2	52.1	55.3	+7.1
Reinforcement Learning								
diffuCoder	57.3	51.2	60.4	+3.1	61.1	36.6	63.0	+1.9

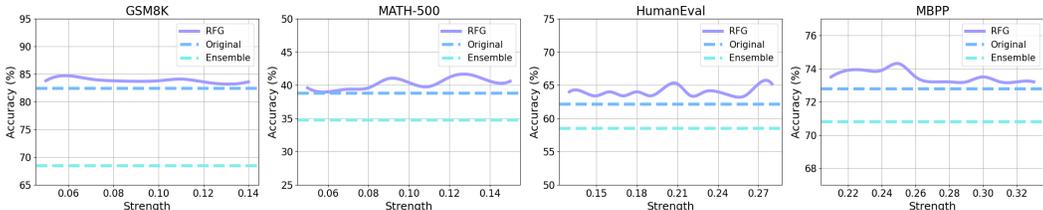


Figure 2: Accuracy of RFG under varying guidance strength w across four benchmarks: GSM8K and MATH-500 for mathematical reasoning using d1-LLaDA, and HumanEval and MBPP for code generation using Dream-Instruct and DiffuCoder, respectively. We observe that RFG consistently improves performance over a broad range of guidance strength.

selecting $-1 < w < 0$, RFG effectively takes a step back from the over-optimized model and finds a "sweet spot" for the specific task.

Future Directions. We discuss several promising future directions. (1) RFG is compatible with advanced dLLM paradigms from concurrent works such as guiding the block diffusion sampling in SDAR (Cheng et al., 2025) or enhancing models trained with trajectory-aware RL like TraDo (Wang et al., 2025b). (2) Another direction is to adapt RFG’s formulation to use a single dLLM, instantiating the policy and reference models as the same model with and without the prompt, though we note the theoretical decomposition for this setup does not hold so similar performance is not guaranteed. (3) RFG’s framework naturally provides a step-wise reward signal that could be leveraged for the RL-based training of dLLMs, offering a new path for alignment beyond our test-time focus.

5 RELATED WORK

Diffusion Language Models. Early attempts at text diffusion relied on operating in continuous spaces, either by relaxing text tokens in a continuous form (Li et al., 2022; Dieleman et al., 2022; Gong et al., 2023; Chen et al., 2023; Wu et al., 2023) or by diffusing the continuous parameters of discrete distributions (Lou & Ermon, 2023; Lin et al., 2023; Graves et al., 2023; Xue et al.,

2024). Despite conceptual simplicity, these approaches faced challenges in scalability (Gulrajani & Hashimoto, 2023). Alternatively, Austin et al. (2021a) introduced discrete diffusion that operates directly on discrete text tokens, leading to a proliferation of variants (Hoogeboom et al., 2021; He et al., 2022; Campbell et al., 2022; Meng et al., 2022; Sun et al., 2023; Gat et al., 2024; Sahoo et al., 2024; Shi et al., 2024; Zheng et al., 2025; Ou et al., 2025; Nie et al., 2025a; Gong et al., 2025a). In terms of generation order, Arriola et al. (2025) proposed Block Diffusion which generates block-by-block autoregressively while applying parallel diffusion within each block. Building upon these foundations, a significant breakthrough was the successful scaling of discrete diffusion language models, with LLaDA (Nie et al., 2025b) and Dream (Ye et al., 2025) demonstrating performance comparable to their autoregressive counterparts. To further boost reasoning and alignment, reinforcement learning has been applied. For instance, Zhao et al. (2025), Yang et al. (2025), and Tang et al. (2025) adapted the GRPO (Shao et al., 2024) objectives for dLLMs, while Zhu et al. (2025) introduced unbiased variance reduction techniques when applying DPO (Rafailov et al., 2023) to dLLMs. The scope has also expanded beyond text to multimodal domains, with models such as LaViDa (Li et al., 2025), MMaDA (Yang et al., 2025), and Dimple (Yu et al., 2025) integrating text diffusion with vision capabilities. Concurrent work SDAR (Cheng et al., 2025) introduces a hybrid paradigm that turns a pretrained autoregressive model into a blockwise diffusion model through a lightweight adaptation phase, and concurrent work TraDo (Wang et al., 2025b) presents TraceRL, a trajectory aware reinforcement learning framework that aligns the post training objective with the model’s inference trajectory.

Steering Generative Models Reasoning. Steering large generative models toward desired behaviors is a fundamental problem. In diffusion, guidance methods steer the generative process by adjusting the score or noise prediction with auxiliary signals such as class labels, weaker models, or reward gradients to bias sampling toward desired outputs without retraining. Classifier guidance (Dhariwal & Nichol, 2021) uses gradients from an external classifier to steer the diffusion process toward a target label, while classifier-free guidance (CFG) (Ho & Salimans, 2022) combines conditional and unconditional diffusion predictions to achieve the same effect without requiring a separate classifier. Karras et al. (2024) extended CFG and proposed autoguidance that guides the generation process with a deliberately less-trained and smaller version of itself rather than an unconditional model, decoupling quality from diversity and yielding state-of-the-art image generation. In contrast, RFG is entirely training-free: it does not deliberately train a weaker model and instead reuses off-the-shelf dLLMs that directly steer the reverse process. Similar ideas to autoguidance have also been explored in autoregressive LLMs, such as contrastive decoding (Li et al., 2023), which mitigates repetitiveness in generation. Nisonoff et al. (2025) applied guidance to discrete diffusion models, focusing on applications to scientific data such as molecules, DNA, and proteins. Another line of work uses reward-based guidance that typically trains an external reward model to score outputs, which has been successful in aligning autoregressive LLMs (Ouyang et al., 2022; Rafailov et al., 2023; Shao et al., 2024). Outcome reward models (ORMs) evaluate complete responses or final answers without shaping intermediate reasoning, whereas process reward models (PRMs) assign rewards to step-level traces (Uesato et al., 2022; Lightman et al., 2023; Wang et al., 2023; Lu et al., 2024)—a strategy effective for autoregressive LLMs but dependent on dense annotations and a fixed generation order, making it ill-suited to dLLMs with masked, any-order intermediates.

6 CONCLUSION

We introduced reward-free guidance (RFG), a novel and principled framework for enhancing the reasoning capabilities of diffusion large language models (dLLMs) at test time. Our core contribution is a method that guides the denoising process without requiring an explicitly trained process reward model. By parameterizing the reward as the log-likelihood ratio of a policy and reference dLLM, RFG elegantly connects a trajectory-level reward with step-wise guidance. Our theoretical analysis shows that RFG’s sampling process is equivalent to reward-guided sampling. Comprehensive experiments on challenging benchmarks empirically validate the effectiveness of RFG. We envision this framework as a foundation for broader alignment and reasoning improvements in generative models, including multimodal diffusion and agentic reasoning systems, where test-time guidance offers a scalable and general alternative to costly retraining.

REFERENCES

- 540
541
542 Marianne Arriola, Aaron Gokaslan, Justin T Chiu, Zhihan Yang, Zhixuan Qi, Jiaqi Han, Sub-
543 ham Sekhar Sahoo, and Volodymyr Kuleshov. Block diffusion: Interpolating between autore-
544 gressive and diffusion language models. In *The Thirteenth International Conference on Learning*
545 *Representations*, 2025. URL <https://arxiv.org/abs/2503.09573>.
- 546 Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured
547 denoising diffusion models in discrete state-spaces. *Advances in neural information processing*
548 *systems*, 34:17981–17993, 2021a.
- 549 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan,
550 Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language
551 models. *arXiv preprint arXiv:2108.07732*, 2021b.
- 552
553 Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V Le, Christopher Ré, and
554 Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling.
555 *arXiv preprint arXiv:2407.21787*, 2024.
- 556
557 Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and
558 Arnaud Doucet. A continuous time framework for discrete denoising models. *Advances in Neural*
559 *Information Processing Systems*, 35:28266–28279, 2022.
- 560 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared
561 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large
562 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- 563
564 Ting Chen, Ruixiang ZHANG, and Geoffrey Hinton. Analog bits: Generating discrete data using
565 diffusion models with self-conditioning. In *The Eleventh International Conference on Learning*
566 *Representations*, 2023. URL <https://openreview.net/forum?id=3itjR9QxFw>.
- 567 Shuang Cheng, Yihan Bian, Dawei Liu, Yuhua Jiang, Yihao Liu, Linfeng Zhang, Wenhai Wang,
568 Qipeng Guo, Kai Chen, Biqing Qi, et al. Sdar: A synergistic diffusion-autoregression paradigm
569 for scalable sequence generation. *arXiv preprint arXiv:2510.06303*, 2025.
- 570
571 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
572 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John
573 Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*,
574 2021.
- 575 Meihua Dang, Jiaqi Han, Minkai Xu, Kai Xu, Akash Srivastava, and Stefano Ermon. Inference-
576 time scaling of diffusion language models with particle gibbs sampling. *arXiv preprint*
577 *arXiv:2507.08390*, 2025.
- 578
579 DeepMind. Gemini diffusion. 2025. URL [https://deepmind.google/models/](https://deepmind.google/models/gemini-diffusion/)
580 [gemini-diffusion/](https://deepmind.google/models/gemini-diffusion/).
- 581 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep
582 bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of*
583 *the North American Chapter of the Association for Computational Linguistics: Human Language*
584 *Technologies, Volume 1 (Long and Short Papers)*, June 2019.
- 585
586 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances*
587 *in neural information processing systems*, 34:8780–8794, 2021.
- 588 Sander Dieleman, Laurent Sartran, Arman Roshannai, Nikolay Savinov, Yaroslav Ganin, Pierre H
589 Richemond, Arnaud Doucet, Robin Strudel, Chris Dyer, Conor Durkan, et al. Continuous diffu-
590 sion for categorical data. *arXiv preprint arXiv:2211.15089*, 2022.
- 591
592 Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and
593 Yaron Lipman. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37:
133345–133385, 2024.

- 594 Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and Lingpeng Kong. Diffuseq: Sequence
595 to sequence text generation with diffusion models. In *The Eleventh International Conference on*
596 *Learning Representations*, 2023. URL [https://openreview.net/forum?id=jQj-_](https://openreview.net/forum?id=jQj-_rLVXsj)
597 [rLVXsj](https://openreview.net/forum?id=jQj-_rLVXsj).
- 598 Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An,
599 Peilin Zhao, Wei Bi, Jiawei Han, Hao Peng, and Lingpeng Kong. Scaling diffusion language
600 models via adaptation from autoregressive models. In *The Thirteenth International Confer-*
601 *ence on Learning Representations*, 2025a. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=j1tSLYKwg8)
602 [j1tSLYKwg8](https://openreview.net/forum?id=j1tSLYKwg8).
- 603 Shansan Gong, Ruixiang Zhang, Huangjie Zheng, Jiatao Gu, Navdeep Jaitly, Lingpeng Kong, and
604 Yizhe Zhang. Diffucoder: Understanding and improving masked diffusion models for code gen-
605 eration. *arXiv preprint arXiv:2506.20639*, 2025b.
- 606 Alex Graves, Rupesh Kumar Srivastava, Timothy Atkinson, and Faustino Gomez. Bayesian flow
607 networks. *arXiv preprint arXiv:2308.07037*, 2023.
- 608 Ishaan Gulrajani and Tatsunori B Hashimoto. Likelihood-based diffusion language models. *Ad-*
609 *vances in Neural Information Processing Systems*, 36:16693–16715, 2023.
- 610 Zhengfu He, Tianxiang Sun, Kuanning Wang, Xuanjing Huang, and Xipeng Qiu. Diffusion-
611 bert: Improving generative masked language models with diffusion models. *arXiv preprint*
612 *arXiv:2211.15029*, 2022.
- 613 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
614 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *Sort*, 2
615 (4):0–6.
- 616 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint*
617 *arXiv:2207.12598*, 2022.
- 618 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*
619 *neural information processing systems*, 33:6840–6851, 2020.
- 620 Emiel Hooeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows
621 and multinomial diffusion: Learning categorical distributions. *Advances in neural information*
622 *processing systems*, 34:12454–12465, 2021.
- 623 Inception Labs, Samar Khanna, Siddhant Kharbanda, Shufan Li, Harshit Varma, Eric Wang, Sawyer
624 Birnbaum, Ziyang Luo, Yanis Miraoui, Akash Palrecha, et al. Mercury: Ultra-fast language
625 models based on diffusion. *arXiv preprint arXiv:2506.17298*, 2025.
- 626 Tero Karras, Miika Aittala, Tuomas Kynkäänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine.
627 Guiding a diffusion model with a bad version of itself. *Advances in Neural Information Processing*
628 *Systems*, 37:52996–53021, 2024.
- 629 Shufan Li, Konstantinos Kallidromitis, Hritik Bansal, Akash Gokul, Yusuke Kato, Kazuki Kozuka,
630 Jason Kuen, Zhe Lin, Kai-Wei Chang, and Aditya Grover. Lavidia: A large diffusion language
631 model for multimodal understanding. *ArXiv preprint*, abs/2505.16839, 2025. URL [https:](https://arxiv.org/abs/2505.16839)
632 [//arxiv.org/abs/2505.16839](https://arxiv.org/abs/2505.16839).
- 633 Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-
634 lm improves controllable text generation. *Advances in neural information processing systems*, 35:
635 4328–4343, 2022.
- 636 Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke
637 Zettlemoyer, and Mike Lewis. Contrastive decoding: Open-ended text generation as optimization.
638 In *The 61st Annual Meeting Of The Association For Computational Linguistics*, 2023.
- 639 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
640 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth*
641 *International Conference on Learning Representations*, 2023.

- 648 Zhenghao Lin, Yeyun Gong, Yelong Shen, Tong Wu, Zhihao Fan, Chen Lin, Nan Duan, and Weizhu
649 Chen. Text generation with diffusion language models: A pre-training approach with continuous
650 paragraph denoise. In *International Conference on Machine Learning*, pp. 21051–21064. PMLR,
651 2023.
- 652 Aaron Lou and Stefano Ermon. Reflected diffusion models. In *International Conference on Machine*
653 *Learning*, pp. 22675–22701. PMLR, 2023.
- 654 Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios
655 of the data distribution. In *International Conference on Machine Learning*, pp. 32819–32848.
656 PMLR, 2024.
- 657 Jianqiao Lu, Zhiyang Dou, Hongru Wang, Zeyu Cao, Jianbo Dai, Yunlong Feng, and Zhijiang Guo.
658 Autopsv: Automated process-supervised verifier. *Advances in Neural Information Processing*
659 *Systems*, 37:79935–79962, 2024.
- 660 Chenlin Meng, Kristy Choi, Jiaming Song, and Stefano Ermon. Concrete score matching: General-
661 ized score matching for discrete data. *Advances in Neural Information Processing Systems*, 35:
662 34532–34545, 2022.
- 663 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke
664 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time
665 scaling, 2025. URL <https://arxiv.org/abs/2501.19393>.
- 666 Shen Nie, Fengqi Zhu, Chao Du, Tianyu Pang, Qian Liu, Guangtao Zeng, Min Lin, and Chongx-
667 uan Li. Scaling up masked diffusion models on text. In *The Thirteenth International Confer-*
668 *ence on Learning Representations*, 2025a. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=WNvwwK0tut)
669 [WNvwwK0tut](https://openreview.net/forum?id=WNvwwK0tut).
- 670 Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin,
671 Ji-Rong Wen, and Chongxuan Li. Large language diffusion models, 2025b. URL <https://arxiv.org/abs/2502.09992>.
- 672 Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking guid-
673 ance for discrete state-space diffusion and flow models. In *The Thirteenth International Confer-*
674 *ence on Learning Representations*, 2025. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=XsgH154yO7)
675 [XsgH154yO7](https://openreview.net/forum?id=XsgH154yO7).
- 676 Jingyang Ou, Shen Nie, Kaiwen Xue, Fengqi Zhu, Jiacheng Sun, Zhenguo Li, and Chongxuan
677 Li. Your absorbing discrete diffusion secretly models the conditional distributions of clean data.
678 In *The Thirteenth International Conference on Learning Representations*, 2025. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=sMyXP8Tanm)
679 [sMyXP8Tanm](https://openreview.net/forum?id=sMyXP8Tanm).
- 680 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
681 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-
682 low instructions with human feedback. *Advances in neural information processing systems*, 35:
683 27730–27744, 2022.
- 684 Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational
685 space control. In *Proceedings of the 24th international conference on Machine learning*, pp. 745–
686 750, 2007.
- 687 Mihir Prabhudesai, Mengning Wu, Amir Zadeh, Katerina Fragkiadaki, and Deepak Pathak. Diffu-
688 sion beats autoregressive in data-constrained settings. *arXiv preprint arXiv:2507.15857*, 2025.
- 689 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
690 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances*
691 *in neural information processing systems*, 36:53728–53741, 2023.
- 692 Subham Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin Chiu,
693 Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language
694 models. *Advances in Neural Information Processing Systems*, 37:130136–130184, 2024.

- 702 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
703 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical
704 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
705
- 706 Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis Titsias. Simplified and general-
707 ized masked diffusion for discrete data. *Advances in neural information processing systems*, 37:
708 103131–103167, 2024.
- 709 Raghav Singhal, Zachary Horvitz, Ryan Teehan, Mengye Ren, Zhou Yu, Kathleen McKeown, and
710 Rajesh Ranganath. A general framework for inference-time scaling and steering of diffusion
711 models. *arXiv preprint arXiv:2501.06848*, 2025.
- 712 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally
713 can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.
714
- 715 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
716 learning using nonequilibrium thermodynamics. In *International conference on machine learn-*
717 *ing*, pp. 2256–2265. pmlr, 2015.
- 718 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
719 Poole. Score-based generative modeling through stochastic differential equations. In *Interna-*
720 *tional Conference on Learning Representations*, 2020.
721
- 722 Haoran Sun, Lijun Yu, Bo Dai, Dale Schuurmans, and Hanjun Dai. Score-based continuous-time
723 discrete diffusion models. In *The Eleventh International Conference on Learning Representa-*
724 *tions*, 2023. URL <https://openreview.net/forum?id=BYWWwSY2G5s>.
- 725 Xiaohang Tang, Rares Dolga, Sangwoong Yoon, and Ilija Bogunovic. wd1: Weighted policy opti-
726 mization for reasoning in diffusion language models. *arXiv preprint arXiv:2507.08838*, 2025.
727
- 728 Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia
729 Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process-and
730 outcome-based feedback. *arXiv preprint arXiv:2211.14275*, 2022.
- 731 Guanghan Wang, Yair Schiff, Subham Sekhar Sahoo, and Volodymyr Kuleshov. Remasking discrete
732 diffusion models with inference-time scaling. *arXiv preprint arXiv:2503.00307*, 2025a.
733
- 734 Peiyi Wang, Lei Li, Zhihong Shao, RX Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang
735 Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. *arXiv*
736 *preprint arXiv:2312.08935*, 2023.
- 737 Yinjie Wang, Ling Yang, Bowen Li, Ye Tian, Ke Shen, and Mengdi Wang. Revolutioniz-
738 ing reinforcement learning framework for diffusion large language models. *arXiv preprint*
739 *arXiv:2509.06949*, 2025b.
- 740 Tong Wu, Zhihao Fan, Xiao Liu, Hai-Tao Zheng, Yeyun Gong, Jian Jiao, Juntao Li, Jian Guo, Nan
741 Duan, Weizhu Chen, et al. Ar-diffusion: Auto-regressive diffusion model for text generation.
742 *Advances in Neural Information Processing Systems*, 36:39957–39974, 2023.
743
- 744 Minkai Xu, Tomas Geffner, Karsten Kreis, Weili Nie, Yilun Xu, Jure Leskovec, Stefano Ermon, and
745 Arash Vahdat. Energy-based diffusion language models for text generation. In *The Thirteenth*
746 *International Conference on Learning Representations*, 2025. URL [https://openreview.](https://openreview.net/forum?id=sL2F9YCMXf)
747 [net/forum?id=sL2F9YCMXf](https://openreview.net/forum?id=sL2F9YCMXf).
- 748 Kaiwen Xue, Yuhao Zhou, Shen Nie, Xu Min, Xiaolu Zhang, Jun Zhou, and Chongxuan Li. Uni-
749 fying bayesian flow networks and diffusion models through stochastic differential equations. In
750 *International Conference on Machine Learning*, pp. 55656–55681. PMLR, 2024.
- 751 Ling Yang, Ye Tian, Bowen Li, Xinchun Zhang, Ke Shen, Yunhai Tong, and Mengdi Wang. Mmada:
752 Multimodal large diffusion language models. *ArXiv preprint*, abs/2505.15809, 2025. URL
753 <https://arxiv.org/abs/2505.15809>.
754
- 755 Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng
Kong. Dream 7b: Diffusion large language models, 2025.

756 Runpeng Yu, Xinyin Ma, and Xinchao Wang. Dimple: Discrete diffusion multimodal large language
757 model with parallel decoding. *ArXiv preprint*, abs/2505.16990, 2025. URL <https://arxiv.org/abs/2505.16990>.
758
759
760 Lifan Yuan, Wendi Li, Huayu Chen, Ganqu Cui, Ning Ding, Kaiyan Zhang, Bowen Zhou, Zhiyuan
761 Liu, and Hao Peng. Free process rewards without process labels. In *Forty-second International
762 Conference on Machine Learning*, 2025. URL [https://openreview.net/forum?id=
763 8ThnPFhGm8](https://openreview.net/forum?id=8ThnPFhGm8).
764
765 Siyan Zhao, Devaansh Gupta, Qinqing Zheng, and Aditya Grover. d1: Scaling reasoning in diffusion
766 large language models via reinforcement learning. *ArXiv preprint*, abs/2504.12216, 2025. URL
767 <https://arxiv.org/abs/2504.12216>.
768
769 Kaiwen Zheng, Yongxin Chen, Hanzi Mao, Ming-Yu Liu, Jun Zhu, and Qinsheng Zhang. Masked
770 diffusion models are secretly time-agnostic masked models and exploit inaccurate categorical
771 sampling. In *The Thirteenth International Conference on Learning Representations*, 2025. URL
772 <https://openreview.net/forum?id=CTC7CmirNr>.
773
774 Fengqi Zhu, Rongzhen Wang, Shen Nie, Xiaolu Zhang, Chunwei Wu, Jun Hu, Jun Zhou, Jianfei
775 Chen, Yankai Lin, Ji-Rong Wen, et al. Llada 1.5: Variance-reduced preference optimization for
776 large language diffusion models. *arXiv preprint arXiv:2505.19223*, 2025.
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

A PROOF

Proposition 3.1. Given a diffusion trajectory-level reward that is parameterized as the log-likelihood ratio of two dLLMs, i.e., $r_\theta(\mathbf{x}_{0:T}) := \beta \log \frac{p_\theta(\mathbf{x}_{0:T})}{p_{\text{ref}}(\mathbf{x}_{0:T})}$. Define $Q_\theta^t(\mathbf{x}_{t-1}, \mathbf{x}_{t:T}) := \sum_{i=t}^T \beta \log \frac{p_\theta(\mathbf{x}_{i-1}|\mathbf{x}_{i:T})}{p_{\text{ref}}(\mathbf{x}_{i-1}|\mathbf{x}_{i:T})}$, then we have that q_θ^t is the expectation of exponential r_θ at step t :

$$Q_\theta^t(\mathbf{x}_{t-1}, \mathbf{x}_{t:T}) = \beta \log \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{0:T}|\mathbf{x}_{t-1:T})} e^{\frac{1}{\beta} r_\theta(\mathbf{x}_{0:T})}, \quad (6)$$

where β is a hyperparameter for weighting the reward function.

Proof. The proposition is mainly proven by induction. The proposition and proof are largely borrowed from related literature in autoregressive LLMs (Yuan et al., 2025). The key difference is that in previous work, the analyses are mainly for autoregressive generation, while in this paper, we focus on the sampling trajectory of dLLMs.

Suppose we are given dLLM that discrete the sampling trajectory into T steps. The proof of Theorem 3.1 can be decomposed into the two following arguments:

1. At $t = 1$, $Q_\theta^1(\mathbf{x}_0, \mathbf{x}_{1:T}) = r_\theta(\mathbf{x}_{0:T}) = \beta \log \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{0:T}|\mathbf{x}_{0:T})} e^{\frac{1}{\beta} r_\theta(\mathbf{x}_{0:T})}$;
2. For $\forall t \in \{1, \dots, T-1\}$, if $Q_\theta^t(\mathbf{x}_{t-1}, \mathbf{x}_{t:T}) = \beta \log \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{0:T}|\mathbf{x}_{t-1:T})} e^{\frac{1}{\beta} r_\theta(\mathbf{x}_{0:T})}$, then we would also have $Q_\theta^{t+1}(\mathbf{x}_t, \mathbf{x}_{t+1:T}) = \beta \log \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{0:T}|\mathbf{x}_t:T)} e^{\frac{1}{\beta} r_\theta(\mathbf{x}_{0:T})}$.

Proof of 1. In dLLM, we have that $p(\mathbf{x}_{0:T}) = p(x_T) \prod_{t=1}^T p(\mathbf{x}_{t-1}|\mathbf{x}_t)$. Then we have

$$r_\theta(\mathbf{x}) := \beta \log \frac{p_\theta(\mathbf{x}_{0:T})}{p_{\text{ref}}(\mathbf{x}_{0:T})} = \beta \log \prod_{i=1}^T \frac{p_\theta(\mathbf{x}_{i-1}|\mathbf{x}_i)}{p_{\text{ref}}(\mathbf{x}_{i-1}|\mathbf{x}_i)} = \sum_{i=1}^T \beta \log \frac{p_\theta(\mathbf{x}_{i-1}|\mathbf{x}_i)}{p_{\text{ref}}(\mathbf{x}_{i-1}|\mathbf{x}_i)}.$$

Then we trivially have that:

$$Q_\theta^1(\mathbf{x}_0, \mathbf{x}_{1:T}) = \sum_{i=1}^T \beta \log \frac{p_\theta(\mathbf{x}_{i-1}|\mathbf{x}_i)}{p_{\text{ref}}(\mathbf{x}_{i-1}|\mathbf{x}_i)} = r_\theta(\mathbf{x}_{0:T}) = \beta \log e^{\frac{1}{\beta} r_\theta(\mathbf{x}_{0:T})} = \beta \log \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{0:T}|\mathbf{x}_{0:T})} e^{\frac{1}{\beta} r_\theta(\mathbf{x}_{0:T})}.$$

Proof of 2. For $\forall t \in \{1, \dots, T-1\}$, given $Q_\theta^t(\mathbf{x}_{t-1}, \mathbf{x}_{t:T}) = \beta \log \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{0:T}|\mathbf{x}_{t-1:T})} e^{\frac{1}{\beta} r_\theta(\mathbf{x}_{0:T})}$, we have:

$$\begin{aligned} Q_\theta^{t+1}(\mathbf{x}_t, \mathbf{x}_{t+1:T}) &= \beta \sum_{i=t}^T \log \frac{p_\theta(\mathbf{x}_i|\mathbf{x}_{i+1:T})}{p_{\text{ref}}(\mathbf{x}_i|\mathbf{x}_{i+1:T})} \\ &= \beta \log \prod_{i=t}^T \frac{p_\theta(\mathbf{x}_i|\mathbf{x}_{i+1:T})}{p_{\text{ref}}(\mathbf{x}_i|\mathbf{x}_{i+1:T})} \\ &= \beta \log \prod_{i=t}^T \frac{p_\theta(\mathbf{x}_i|\mathbf{x}_{i+1:T})}{p_{\text{ref}}(\mathbf{x}_i|\mathbf{x}_{i+1:T})} \sum_{\mathbf{y}_{t-1}} p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_{t:T}) \\ &= \beta \log \prod_{i=t}^T \frac{p_\theta(\mathbf{x}_i|\mathbf{x}_{i+1:T})}{p_{\text{ref}}(\mathbf{x}_i|\mathbf{x}_{i+1:T})} \sum_{\mathbf{y}_{t-1}} p_{\text{ref}}(\mathbf{x}_{t-1}|\mathbf{x}_{t:T}) \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_{t:T})}{p_{\text{ref}}(\mathbf{x}_{t-1}|\mathbf{x}_{t:T})} \\ &= \beta \log \prod_{i=t}^T \frac{p_\theta(\mathbf{x}_i|\mathbf{x}_{i+1:T})}{p_{\text{ref}}(\mathbf{x}_i|\mathbf{x}_{i+1:T})} \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{t-1}|\mathbf{x}_{t:T})} \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_{t:T})}{p_{\text{ref}}(\mathbf{x}_{t-1}|\mathbf{x}_{t:T})} \\ &= \beta \log \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{t-1}|\mathbf{x}_{t:T})} \prod_{i=t-1}^T \frac{p_\theta(\mathbf{x}_i|\mathbf{x}_{i+1:T})}{p_{\text{ref}}(\mathbf{x}_i|\mathbf{x}_{i+1:T})} \\ &= \beta \log \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{t-1}|\mathbf{x}_{t:T})} e^{\frac{1}{\beta} Q_\theta^t(\mathbf{x}_{t-1}, \mathbf{x}_{t:T})} \end{aligned}$$

$$\begin{aligned}
&= \beta \log \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{t-1}|\mathbf{x}_{t:T})} \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{0:T}|\mathbf{x}_{t-1:T})} e^{\frac{1}{\beta} r_{\theta}(\mathbf{x}_{0:T})} \\
&= \beta \log \mathbb{E}_{p_{\text{ref}}(\mathbf{x}_{0:T}|\mathbf{x}_{t:T})} e^{\frac{1}{\beta} r_{\theta}(\mathbf{x}_{0:T})}
\end{aligned}$$

which completes the proof. \square

B MASKED DIFFUSION LARGE LANGUAGE MODEL FORMULATION

We can frame a diffusion large language model (dLLM) as a scaled-up implementation of a masked diffusion model tailored for text generation. It operates on a sequence of tokens \mathbf{x}_0 and learns to reverse a masking process.

Forward Process. The forward process $q(\mathbf{x}_t|\mathbf{x}_{t-1})$ is a Markov process that replaces a subset of non-masked tokens in \mathbf{x}_{t-1} with a [MASK] token to produce \mathbf{x}_t . The number of tokens to mask at each step is determined by a predefined noise schedule α_t , where $t \in [0, 1]$ is the timestep. At timestep t , the ratio of masked token is $1 - \alpha_t$. α_t strictly decreases with t , satisfying $\alpha_0 = 1$ and $\alpha_1 = 0$. As an example, LLaDA (Nie et al., 2025b) adopts a linear schedule defined as $\alpha_t = 1 - t$. We then have the following the distribution

$$q_{t|0}(\mathbf{x}_t|\mathbf{x}_0) = \prod_{i=0}^L q_{t|0}(\mathbf{x}_t^{(i)}|\mathbf{x}_0^{(i)}), \quad q_{t|0}(\mathbf{x}_t^{(i)}|\mathbf{x}_0^{(i)}) = \begin{cases} 1 - \alpha_t, & \mathbf{x}_t^{(i)} = [\text{MASK}] \\ \alpha_t, & \mathbf{x}_t^{(i)} = \mathbf{x}_0^{(i)} \end{cases}$$

Reverse Process. The reverse process aims to predict the original tokens that were masked in \mathbf{x}_t . It is worth noting that during the forward process once a token is masked, it remains in the masked state and cannot transition to other states. Given this, the conditional distribution for the reverse process moving from a time step t to an earlier step s , where $0 \leq s < t \leq 1$, is expressed as

$$q_{s|t}(\mathbf{x}_s^{(i)}|\mathbf{x}_t) = \begin{cases} 1, & \mathbf{x}_t^{(i)} \neq [\text{MASK}], \mathbf{x}_s^{(i)} = \mathbf{x}_t^{(i)} \\ \frac{1-\alpha_s}{1-\alpha_t}, & \mathbf{x}_t^{(i)} = [\text{MASK}], \mathbf{x}_s^{(i)} = [\text{MASK}] \\ \frac{\alpha_s - \alpha_t}{1-\alpha_t} q_{0|t}(\mathbf{x}_s^{(i)}|\mathbf{x}_t), & \mathbf{x}_t^{(i)} = [\text{MASK}], \mathbf{x}_s^{(i)} \neq [\text{MASK}] \\ 0, & \text{otherwise} \end{cases}$$

In practice, $q_{0|t}(\mathbf{x}_s^{(i)}|\mathbf{x}_t)$ is approximated by the dLLM that recover the original token in sequence \mathbf{x}_0 given the partially masked sequence \mathbf{x}_t .

Training. The dLLM, denoted as $p_{\theta}(\cdot|\mathbf{x}_t)$, is trained to reconstruct the original sequence \mathbf{x}_0 by predicting all masked tokens in the sequence \mathbf{x}_t , analogous to the masked language modeling objective. The loss is an upper bound of the negative log-likelihood of the model distribution:

$$\mathcal{L}(\theta) := \mathbb{E}_{\mathbf{x}_0 \sim p_{\text{data}}(\mathbf{x}_0), \mathbf{x}_t \sim q_{t|0}(\cdot|\mathbf{x}_0)} \left[\int_{t=0}^{t=1} \frac{\alpha'_t}{1-\alpha_t} \sum_{i=1}^L \mathbb{I}[\mathbf{x}_t^{(i)} = [\text{MASK}]] \log p_{\theta}(\mathbf{x}_0^{(i)}|\mathbf{x}_t) \right],$$

where p_{data} is the distribution of training data. For models that employ linear noise schedule $\alpha_t = 1 - t$ (e.g., LLaDA), the loss function simplifies to

$$\mathcal{L}(\theta) = -\mathbb{E}_{t \sim \mathcal{U}[0,1], \mathbf{x}_0 \sim p_{\text{data}}(\mathbf{x}_0), \mathbf{x}_t \sim q_{t|0}(\cdot|\mathbf{x}_0)} \left[\frac{1}{t} \sum_{i=1}^L \mathbb{I}[\mathbf{x}_t^{(i)} = [\text{MASK}]] \log p_{\theta}(\mathbf{x}_0^i|\mathbf{x}_t) \right].$$

Inference. Inference starts with a sequence of L [MASK] tokens. The model then iteratively un-masks the sequence over T steps. At each step t , the model $p_{\theta}(\mathbf{x}_0|\mathbf{x}_t)$ predicts a full sequence. A key feature of dLLMs is their flexibility in the unmasking order. Instead of a fixed left-to-right generation, dLLMs can employ an *any-order* generation scheme. For instance, at each step, the model might only replace the [MASK] tokens for which it has the highest prediction confidence (i.e., the lowest entropy). This new, partially filled sequence becomes \mathbf{x}_{t-1} , and the process repeats. This allows the model to build the sequence from the easiest parts to the hardest, leveraging bidirectional context at every step.

C IMPLEMENTATION DETAILS

Hyperparameters. For mathematical reasoning benchmarks (GSM8K and MATH-500), we set the generation length to 256 for all models, except LLaDA-1.5 on MATH-500, where we follow the original paper and use a generation length of 512 for consistency. For code generation benchmarks (HumanEval and MBPP), the generation length is set to 512 across all models. In all cases, the number of generation steps is set equal to the generation length. For the LLaDA family (LLaDA-Instruct, d1-LLaDA, and LLaDA-1.5), we adopt block diffusion with a block length of 32 for all tasks, except again for LLaDA-1.5 on MATH-500, where we use a block length of 64 to match the original setup. We apply low-confidence remasking for all LLaDA-based models. The temperature is set to 0 for LLaDA family, following the recommendation from the original paper. For the Dream family (Dream-Instruct and DiffuCoder), we use nucleus sampling with $p = 0.95$ and a temperature of 0.1. To ensure fairness, we use exactly the same set of hyperparameters for the original models, the naive ensemble baselines, and our RFG.

Task Prompts. For GSM8k and MATH-500, we provide the original problem statement followed by an instruction asking the model to output the final answer enclosed in a `\boxed` command. For HumanEval and MBPP, we prepend a generation prefix: “Here is the code to solve this problem `\n```\npython`”. For MBPP specifically, we include only a single test case in the prompt, striking a balance between clarifying the task details and avoiding potential test-data leakage. We use exact same setting for all model types for each task.

D QUALITATIVE EXAMPLES

D.1 GSM8K

D.1.1 LLaDA INSTRUCT

Query (GSM8K)

Marissa is hiking a 12-mile trail. She took 1 hour to walk the first 4 miles, then another hour to walk the next two miles. If she wants her average speed to be 4 miles per hour, what speed (in miles per hour) does she need to walk the remaining distance?

Original (LLaDA Instruct)

First, let's determine the total distance Marissa has already walked:

$$4 \text{ miles} + 2 \text{ miles} = 6 \text{ miles}$$

Next, let's calculate the total time she has spent walking:

$$1 \text{ hour} + 1 \text{ hour} = 2 \text{ hours}$$

Now, let's find the remaining distance she needs to walk:

$$12 \text{ miles} - 6 \text{ miles} = 6 \text{ miles}$$

Next, let's determine the remaining time she has:

$$4 \text{ hours} - 2 \text{ hours} = 2 \text{ hours}$$

Finally, let's calculate the speed she needs to walk the remaining distance:

$$\text{Speed} = \frac{6 \text{ miles}}{2 \text{ hours}} = 3 \text{ miles per hour}$$

Therefore, Marissa needs to walk the remaining distance at a speed of

3 miles per hour to achieve an average speed of 4 miles per hour.

RFQ (LLaDA Instruct)

To determine how much Christina will spend on gift bags, we need to follow these steps:

First, let's determine the total time Marissa has already spent hiking. She took 1 hour to walk the first 4 miles and another 1 hour to walk the next 2 miles. Therefore, the total time spent is:

$$1 + 1 = 2 \text{ hours}$$

Next, we know the total distance of the trail is 12 miles. She has already walked 4 miles + 2 miles = 6 miles. Therefore, the remaining distance is:

$$12 - 6 = 6 \text{ miles}$$

Marissa wants her average speed to be 4 miles per hour for the entire trail. Therefore, the total time she needs to complete the trail is:

$$\frac{12}{4} = 3 \text{ hours}$$

Since she has already spent 2 hours hiking, she has 1 hour left to complete the remaining 6 miles. Therefore, the speed she needs to walk the remaining distance is:

$$\frac{6}{1} = 6 \text{ miles per hour}$$

The final answer is:

6

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

1026 D.1.2 DREAM INSTRUCT
10271028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039

1040 Query (GSM8K)

1041
1042 Carla is downloading a 200 GB file. Normally she can download 2 GB/minute, but 40% of the way
1043 through the download, Windows forces a restart to install updates, which takes 20 minutes. Then Carla
1044 has to restart the download from the beginning. How long does it take to download the file?
10451046
1047
1048 Original (Dream Instruct)1049
1050 First, we calculate the time it takes to
1051 download 40% of the 200 GB file at
1052 a rate of 2 GB/minute. 40% of 200
1053 GB is 80 GB, so it takes $80 \text{ GB} / 2$
1054 $\text{GB/minute} = 40$ minutes to download
80 GB.1055 Next, we calculate the time it takes to
1056 download the remaining 60% of the
1057 file at the same rate. 60% of 200 GB
1058 is 120 GB, so it takes $120 \text{ GB} / 2$
1059 $\text{GB/minute} = 60$ minutes to download
120 GB.1060 Finally, we add the time it takes to
1061 download the first 40% of the file, the
1062 time it takes to download the remain-
1063 ing 60%, and the time it takes to install
1064 updates. The total time is 40 minutes
1065 + 60 minutes + 20 minutes = 120 min-
1066 utes.1067 Therefore, it takes Carla 120 min-
1068 utes to download the file.1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079 RFG (Dream Instruct)1. First, calculate the amount of data Carla has down-
loaded when the 40

$$200 \times 0.4 = 80 \text{ GB}$$

2. Since Carla downloads at a rate of 2 GB/minute,
the time it takes to download 80 GB is:

$$\frac{80}{2} = 40 \text{ minutes}$$

3. After the restart, Carla has to download the entire
200 GB again. The time it takes to download 200 GB
at a rate of 2 GB/minute is:

$$\frac{200}{2} = 100 \text{ minutes}$$

4. Add the time for the restart and the time to down-
load the entire file again to find the total time:

$$40 + 20 + 100 = 160 \text{ minutes}$$

$$\boxed{160} \text{ minutes}$$

D.1.3 D1-LLADA

Query (GSM8K)

Candice put 80 post-it notes in her purse before she headed out to her job at the coffee shop. On her way, she stopped off at the store and purchased a package of Post-it notes; At work, she placed a single Post-it note on each of 220 different cups of coffee. If she had 23 post-it notes remaining overall, how many Post-it notes were in the package that she purchased?

Original (d1 LLaDA)

To determine how many Post-it notes were in the package that Candice purchased, we need to follow these steps:

1. Calculate the total number of Post-it notes Candice used at work. 2. Add the number of Post-it notes she had remaining to the number of Post-it notes she initially had in her purse. 3. Subtract the number of Post-it notes she used at work from the total number of Post-it notes she had after purchasing the package.

First, we know that Candice had 80 Post-it notes in her purse initially and 23 Post-it notes remaining after work. Therefore, the total number of Post-it notes she had after purchasing the package is:

$$80 + 23 = 303$$

Next, we know that Candice used 220 Post-it notes at work. To find the number of Post-it notes in the package, we subtract the number of Post-it notes used from the total number of Post-it notes she had:

$$303 - 220 = 83$$

Thus, the number of Post-it notes in the package is:

83

RFG (d1 LLaDA)

To determine how many Post-it notes were in the package that Candice purchased, we need to follow these steps:

1. Calculate the total number of Post-it notes Candice used at work. 2. Add the number of Post-it notes she had remaining to the number used at work to find the total number of Post-it notes she had. 3. Subtract the number of Post-it notes she originally had in her purse from the total number of Post-it notes she had to find out how many she purchased in the package.

First, we know that Candice used 220 Post-it notes at work and had 23 Post-it notes remaining. Therefore, the total number of Post-it notes she had is:

$$220 + 23 = 243$$

Next, we know that Candice originally had 80 Post-it notes in her purse. So, the number of Post-it notes she purchased in the package is:

$$243 - 80 = 163$$

Thus, the number of Post-it notes in the package that Candice purchased is:

163

1134 D.1.4 LLADA 1.5

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

Query (GSM8K)

1146

1147 Julia was preparing for a dinner party at her house, where she intended to serve stew. She noticed that
 1148 she was out of plastic spoons, so she bought a new package of spoons. Later, her husband also bought
 1149 a package of 5 new spoons and gave them to Julia. While Julia was making the stew, she used three of
 1150 the spoons to sample her stew. Later, when she went to set the table, she had a total of 12 spoons. How
 1151 many spoons were in the package that Julia bought?

1152

1153

1154 Original (LLaDA 1.5)

1155

1156 To determine how many spoons were in the pack-
 1157 age that Julia bought, we need to follow these
 1158 steps:

1159 1. **Identify the total number of spoons Julia**
 1160 **had before setting the table:** Julia had 12 spoons
 1161 before setting the table.

1162 2. **Account for the spoons used while making**
 1163 **the stew:** Julia used 3 spoons to sample her stew.

1164 3. **Determine the number of spoons given by**
 1165 **her husband:** Her husband gave her 5 spoons.

1166 4. **Calculate the total number of spoons before**
 1167 **buying the package:** Subtract the spoons used
 1168 from the total number of spoons before setting
 1169 the table:

$$12 - 3 = 9$$

1170 So, Julia had 9 spoons before buying the package.

1171 5. **Calculate the number of spoons in the pack-**
 1172 **age Julia bought:** Subtract the spoons given by
 1173 her husband from the total number of spoons be-
 1174 fore buying the package:

$$9 - 5 = 4$$

1175 Therefore, the package that Julia bought con-
 1176 tained 4 spoons.

1177 The final answer is:

1178 4

1179

1180

1181

1182

1183

1184

1185

1186

1187

RFG (LLaDA 1.5)

To determine how many spoons were in the package that Julia bought, we need to follow these steps:

1. **Identify the total number of spoons Julia had when setting the table:** Julia had 12 spoons when setting the table.

2. **Account for the spoons used while making the stew:** Julia used 3 spoons to sample her stew.

3. **Determine the number of spoons given by her husband:** Her husband gave her 5 spoons.

4. **Calculate the total number of spoons before setting the table:** Let x be the number of spoons in the package Julia bought. After receiving 5 spoons from her husband, she had $x + 5$ spoons. After using 3 spoons, she had 12 spoons left.

$$(x + 5) - 3 = 12$$

Simplify the equation:

$$x + 2 = 12$$

Subtract 2 from both sides:

$$x = 10$$

Therefore, Julia bought 10 spoons.

1188 D.2 MATH-500

1189
1190 D.2.1 LLaDA INSTRUCT1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Query (MATH-500)

If $n \equiv 2 \pmod{7}$, then find the remainder when $(n+2)(n+4)(n+6)$ is divided by 7.

Original (LLaDA Instruct)

First, we know that $n \equiv 2 \pmod{7}$. We need to find the remainder when $(n+2)(n+4)(n+6)$ is divided by 7.

Let's compute each term modulo 7:

1. $n+2 \equiv 2+2 \equiv 4 \pmod{7}$ 2. $n+4 \equiv 2+4 \equiv 6 \pmod{7}$ 3. $n+6 \equiv 2+6 \equiv 0 \pmod{7}$

Now, we multiply these results together:

$$(n+2)(n+4)(n+6) \equiv 4 \cdot 6 \cdot 0 \pmod{7}$$

Since one of the factors is 0, the entire product is:

$$0 \pmod{7}$$

Thus, the remainder when $(n+2)(n+4)(n+6)$ is divided by 7 is

0

RFG (LLaDA Instruct)

First, we know that $n \equiv 2 \pmod{7}$. We need to find the remainder when $(n+2)(n+4)(n+6)$ is divided by 7.

Let's compute each term modulo 7:

1. $n+2 \equiv 2+2 \equiv 4 \pmod{7}$ 2. $n+4 \equiv 2+4 \equiv 6 \equiv -1 \pmod{7}$ 3. $n+6 \equiv 2+6 \equiv 8 \equiv 1 \pmod{7}$

Next, we multiply these results together:

$$(4)(-1)(1) \equiv -4 \pmod{7}$$

Finally, we convert -4 to a positive equivalent modulo 7:

$$-4 \equiv 3 \pmod{7}$$

Thus, the remainder when $(n+2)(n+4)(n+6)$ is divided by 7 is

3

1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

D.2.2 DREAM INSTRUCT

Query (MATH-500)

CTwo fair, 6-sided dice are thrown. What is the probability that the product of the two numbers is a multiple of 5? Express your answer as a common fraction.

Original (Dream Instruct)

To find the probability that the product of the two numbers is a multiple of 5, we need to consider the cases where at least one of the numbers is a multiple of 5. There are 6 possible outcomes for each die, so there are a total of $6 \times 6 = 36$ possible outcomes when two dice are thrown.

The multiples of 5 on a 6-sided die are 5 itself. Therefore, there is only 1 way for each die to show a multiple of 5. Thus, there are $1 \times 1 = 1$ outcome where both dice show a multiple of 5.

The probability that the product of the two numbers is a multiple of 5 is the number of favorable outcomes (where at least one die shows a multiple of 5) divided by the total number of outcomes. Since there is only one outcome where both dice show a multiple of 5, the probability is $\frac{1}{36}$.

Therefore, the probability that the product of the two numbers is a multiple of 5 is $\frac{1}{36}$.

RFG (Dream Instruct)

To find the probability that the product of the two numbers is a multiple of 5, we need to consider the cases where at least one of the numbers is a multiple of 5. There are 6 possible outcomes for each die, so there are a total of $6 \times 6 = 36$ possible outcomes when two dice are thrown.

The multiples of 5 on a 6-sided die are 5 itself. Therefore, there is only one way for each die to show a multiple of 5. The probability of neither die showing a multiple of 5 is $(\frac{5}{6})^2 = \frac{25}{36}$. Thus, the probability that at least one die shows a multiple of 5 is $1 - \frac{25}{36} = \frac{11}{36}$.

Therefore, the probability that the product of the two numbers is a multiple of 5 is $\frac{11}{36}$.

D.2.3 D1-LLADA

Query (MATH-500)

Compute $\arcsin(-\frac{1}{2})$. Express your answer in radians.

Original (d1 LLaDA)

To compute $\arcsin(-\frac{1}{2})$, we need to find the angle θ such that $\sin(\theta) = -\frac{1}{2}$. The angle whose sine is $-\frac{1}{2}$ is

$\frac{5\pi}{6}$ radians.
Thus, the answer is:

$$\frac{5\pi}{6}$$

RFG (d1 LLaDA)

To compute $\arcsin(-\frac{1}{2})$, we need to find the angle θ such that $\sin \theta = -\frac{1}{2}$ and θ lies within the range of the arcsine function, which is $[-\frac{\pi}{2}, \frac{\pi}{2}]$.

We know that $\sin(-\frac{\pi}{6}) = -\frac{1}{2}$. The angle $-\frac{\pi}{6}$ is within the range $[-\frac{\pi}{2}, \frac{\pi}{2}]$.

Thus, $\arcsin(-\frac{1}{2}) = -\frac{\pi}{6}$.

$$-\frac{\pi}{6}$$

1296 D.2.4 LLaDA 1.5

1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

Query (MATH-500)

What power of 4 is equal to 8? Express your answer as a common fraction.

Original (LLaDA 1.5)

To determine what power of 4 is equal to 8,
we need to solve the equation $4^x = 8$.
First, we can express 8 as a power of 4:

$$8 = 4^2$$

So, the equation becomes:

$$4^x = 4^2$$

Since the bases are the same, we can set the
exponents equal to each other:

$$x = 2$$

Therefore, the power of 4 that is equal to 8
is:

$$\boxed{2}$$

RFG (LLaDA 1.5)

To determine what power of 4 is equal to 8,
we need to solve the equation $4^x = 8$.
First, let's express 8 as a power of 2:

$$8 = 2^3$$

So, we can rewrite the equation as:

$$4^x = 2^3$$

Next, we express 4 as a power of 2:

$$4 = 2^2$$

Thus, we can rewrite the equation as:

$$(2^2)^x = 2^3$$

Using the power of a power property
(a^m)ⁿ = a^{mn} , we get:

$$2^{2x} = 2^3$$

Since the bases are the same, we can set the
exponents equal to each other:

$$2x = 3$$

Solving for x , we divide both sides by 2:

$$x = \frac{3}{2}$$

Therefore, the power of 4 that is equal to 8
is:

$$\boxed{\frac{3}{2}}$$

1350 D.3 HUMANEVAL

1351

1352 D.3.1 LLADA INSTRUCT

1353

1354

1355

Query (HumanEval)

1356

1357

```
def sort_array(array):
```

1358

```
"""
Given an array of non-negative integers, return a copy of the given array after sorting,
you will sort the given array in ascending order if the sum
( first index value, last index value) is odd,
or sort it in descending order if the sum( first index value, last index value) is even.
```

1361

```
Note:
```

1362

```
* don't change the given array.
```

1363

```
Examples:
```

1364

```
* sort_array([]) => []
```

1365

```
* sort_array([5]) => [5]
```

1366

```
* sort_array([2, 4, 3, 0, 1, 5]) => [0, 1, 2, 3, 4, 5]
```

1367

```
* sort_array([2, 4, 3, 0, 1, 5, 6]) => [6, 5, 4, 3, 2, 1, 0]
```

1368

1369

1370

Original (LLaDA Instruct)

1371

```
if len(array) == 0:
```

1372

```
    return []
```

1373

```
sum_last = array[0] + array[-1]
```

1374

```
if sum_last % 2 == 1:
```

1375

```
    return sorted(array, reverse=True)
```

1376

```
else:
```

1377

```
    return sorted(array, reverse=False)
```

1378

1379

1380

D.3.2 DREAM INSTRUCT

1381

1382

1383

Query (HumanEval)

1384

```
def sort_numbers(numbers: str) -> str:
```

1385

```
    """ Input is a space-delimited string of numerals from 'zero' to 'nine'.
```

1386

```
    Valid choices are 'zero', 'one', 'two', 'three', 'four', 'five', 'six',
```

1387

```
    'seven', 'eight' and 'nine'.
```

1388

```
    Return the string with numbers sorted from smallest to largest
```

1389

```
    >>> sort_numbers('three one five')
```

1390

```
    'one three five'
```

1391

1392

1393

Original (Dream Instruct)

1394

```
numbers = numbers.split()
```

1395

```
numbers.sort(key=lambda x: numbers.index(x))
```

1396

```
return ' '.join(numbers)
```

1397

1398

1399

RFG (Dream Instruct)

1400

```
valid_numbers = ['zero', 'one', 'two', 'three', 'four', 'five', 'six', 'seven', 'eight', 'nine']
```

1401

```
numbers = [valid_numbers.index(num) for num in numbers.split()]
```

1402

```
numbers.sort()
```

1403

```
return ' '.join(str(valid_numbers[num]) for num in numbers)
```

1404 D.3.3 DIFFUCODER

1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419

Query (HumanEval)

```
def how_many_times(string: str, substring: str) -> int:
    """ Find how many times a given substring can be found in the original string.
    Count overlapping cases.
    >>> how_many_times('', 'a')
    0
    >>> how_many_times('aaa', 'a')
    3
    >>> how_many_times('aaaa', 'aa')
    3
    """
```

1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432

Original (DiffuCoder)

```
return string.count(substring)
```

RFG (DiffuCoder)

```
count = 0
for i in range(len(string) - len(substring) + 1):
    if string[i:i+len(substring)] == substring:
        count += 1
return count
```

1433 D.4 MBPP

1434

1435 D.4.1 LLaDA INSTRUCT

1436
1437
1438
1439
1440

Query (MBPP)

Write a function to append the given list to the given tuples. Your code should pass the test:

```
assert add_lists([5, 6, 7], (9, 10)) == (9, 10, 5, 6, 7)
```

1446
1447
1448

Original (LLaDA Instruct)

```
def add_lists(lst, tuples):
    return tuple(lst + tuples)
```

RFG (LLaDA Instruct)

```
def add_lists(lst, tuples):
    # Convert the tuples to a list
    tuples_list = list(tuples)
    # Append the given list to the tuples list
    tuples_list.extend(lst)
    # Convert the list back to a tuple
    return tuple(tuples_list)
```

1454
1455
1456
1457

1458 D.4.2 DREAM INSTRUCT
1459

1460 Query (MBPP)

1461 Write a python function to find the number of numbers with an odd value when rotating a binary string the given number of times.
1462 Your code should pass the test:

```
1463 assert odd_Equivalent("011001",6) == 3
1464 """
```

1465 Original (Dream Instruct)

```
1466
1467
1468 def odd_Equivalent(binary_string, rotations):
1469     n = len(binary_string)
1470     odd_count = 0
1471     for i in range(rotations):
1472         binary_string = binary_string[1:] + binary_string[0]
1473         odd_count += sum(int(bit) for bit in binary_string) % 2
1474     return odd_count
```

1475 RFG (Dream Instruct)

```
1476
1477 def odd_Equivalent(binary, n):
1478     count = 0
1479     for i in range(n):
1480         binary = binary[1:] + binary[0]
1481         if int(binary, 2) % 2 == 1:
1482             count += 1
1483     return count
```

1484 D.4.3 DIFFUCODER
1485

1486 Query (MBPP)

1487 Write a function to find the largest sum of a contiguous array in the modified array which is formed by repeating the given array k
1488 times. Your code should pass the test:

```
1489 assert max_sub_array_sum_repeated([10, 20, -30, -1], 4, 3) == 30
1490
```

1491 Original (DiffuCoder)

```
1492
1493 def max_sub_array_sum_repeated(arr, n, k):
1494     repeated_arr = arr * k
1495     max_sum = float('-inf')
1496     for i in range(n):
1497         current_sum = 0
1498         for j in range(i, i + n):
1499             current_sum += repeated_arr[j % n]
1500             max_sum = max(max_sum, current_sum)
1501     return max_sum
```

1502 RFG (DiffuCoder)

```
1503
1504 def max_sub_array_sum_repeated(arr, n, k):
1505     repeated_arr = arr * k
1506     max_sum = float('-inf')
1507     for i in range(n * k):
1508         current_sum = 0
1509         for j in range(i, n * k):
1510             current_sum += repeated_arr[j % n]
1511             max_sum = max(max_sum, current_sum)
1512     return max_sum
```