

---

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 LADiR: LATENT DIFFUSION ENHANCES LLMs FOR TEXT REASONING

005      **Anonymous authors**

006      Paper under double-blind review

## 009      ABSTRACT

011      Large Language Models (LLMs) demonstrate their reasoning ability through  
012      chain-of-thought (CoT) generation. However, LLM's autoregressive decoding  
013      may limit the ability to revisit and refine earlier tokens in a holistic manner, which  
014      can also lead to inefficient exploration for diverse solutions. In this paper, we  
015      propose *LaDiR* (**L**atent **D**iffusion **R**easoner), a novel reasoning framework that  
016      unifies the expressiveness of continuous latent representation with the iterative re-  
017      finement capabilities of latent diffusion models for an existing LLM. We first con-  
018      struct a structured latent reasoning space using a Variational Autoencoder (VAE)  
019      that encodes text reasoning steps into blocks of thought tokens, preserving seman-  
020      tic information and interpretability while offering compact but expressive repre-  
021      sentations. Subsequently, we utilize a latent diffusion model that learns to denoise  
022      a block of latent *thought tokens* with a blockwise bidirectional attention mask,  
023      enabling longer horizon and iterative refinement with adaptive test-time compute.  
024      This design, combined with explicit diversity guidance during diffusion inference,  
025      enables the generation of multiple diverse reasoning trajectories that explore dis-  
026      tinct regions of the latent space, rather than producing repetitive solutions as often  
027      occurs in standard autoregressive sampling. We conduct evaluations on a suite  
028      of mathematical reasoning and planning benchmarks. Empirical results show that  
029      LaDiR consistently improves accuracy, diversity, and interpretability over exist-  
030      ing autoregressive, diffusion-based, and latent reasoning methods, revealing a new  
031      paradigm for text reasoning with latent diffusion.

## 1 INTRODUCTION

032      Large language models (LLMs) have demonstrated remarkable reasoning abilities through extensive  
033      pretraining on human languages, yet the inherent limitations of the autoregressive (AR) paradigm  
034      are becoming increasingly difficult to overlook (Zhou et al., 2024b; Bachmann & Nagarajan, 2025).  
035      As shown in Fig. 1 (top left), their sequential nature prevents revising earlier tokens, making self-  
036      refinement inefficient and difficult (Chen et al., 2024; Huang et al., 2024). Moreover, AR models  
037      with discrete CoT generate a linear chain of thought (CoT) (Dziri et al., 2023; Wei et al., 2023),  
038      which limits reasoning diversity and restricts exploration of multiple valid solutions (Naik et al.,  
039      2024; Yu et al., 2024).

040      Diffusion models (Ho et al., 2020), originally introduced for generation in continuous domains like  
041      images, have recently gained attention in text generation for their ability to maintain global coher-  
042      ence and enable iterative refinement (Ye et al., 2024b; Nie et al., 2025; Lou et al., 2023; Yu et al.,  
043      2025c; Weligalle, 2025; Sahoo et al., 2024; Gulrajani & Hashimoto, 2023). Moreover, prior works  
044      have explored continuous or latent diffusion for language generation (Li et al., 2022; Lovelace et al.,  
045      2024; Zhang et al., 2023; Lovelace et al.; Cetin et al., 2025), operating diffusion in latent spaces  
046      obtained from text autoencoders or token-embedding spaces. Existing works largely emphasize the  
047      parallelization properties of diffusion models (Israel et al., 2025; Nie et al., 2025; Weligalle, 2025)  
048      or evaluate fluency in text generation (Li et al., 2022; Lovelace et al., 2024; Zhang et al., 2023;  
049      Lovelace et al.). Arguably, a more important direction is to ask: *How can these approaches enhance*  
050      *the reasoning capabilities of LLMs?* We focus on one particularly promising capability: *the ability*  
051      *to self-correct and refine reasoning chains at semantic levels in latent space.* As shown in Fig. 1 (top  
052      right), such self-refinement cannot be achieved by discrete diffusion language models that merely  
053      transit into masked tokens.

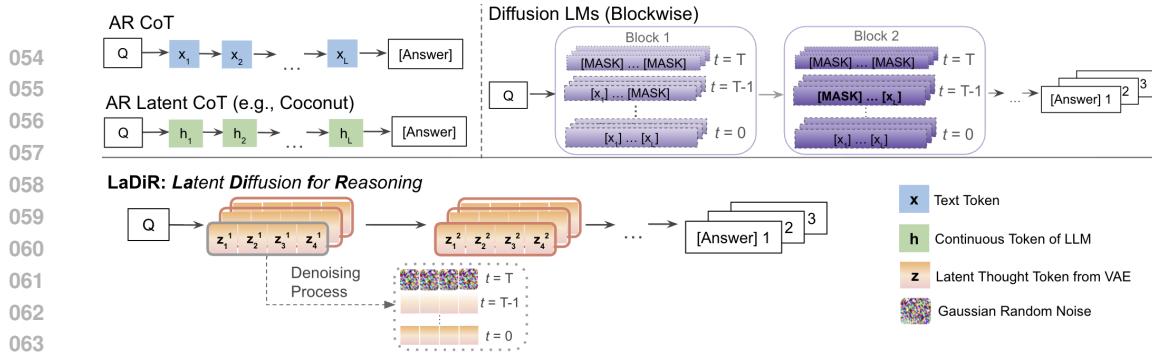


Figure 1: Comparison of reasoning paradigms: autoregressive CoT and latent CoT generate discrete or continuous tokens sequentially; diffusion LMs iteratively refine masked tokens into text in parallel; and our proposed method LaDiR reasons via latent diffusion over thought tokens, enabling iterative refinement at semantic level and diverse solution exploration.

To address this limitation, we introduce **LaDiR** (**L**atent **D**iffusion **R**easoner), a flexible reasoning framework that encodes high-level semantic representations of reasoning steps into continuous latent tokens via a Variational Autoencoder (VAE) as *latent thought tokens*, and trains a latent diffusion model over them to perform reasoning. This bridges the gap between surface-level token refinement and deeper semantic reasoning. After the reasoning process, the model generates final answer tokens conditioned on the generated latent thought tokens. **Unlike prior latent diffusion works for text generation (Li et al., 2022; Lovelace et al., 2024; Zhang et al., 2023), which focus on fluent text generation, our framework is explicitly designed for latent reasoning: it learns causal dependencies across reasoning steps through blockwise diffusion, propagates answer correctness signals back to latent tokens.**

Our proposed paradigm establishes a new reasoning framework as a post-training method, bringing several distinctive advantages. First, the iterative refinement ability of diffusion enables a better trade-off between accuracy and test-time compute, as additional denoising steps can be flexibly allocated to improve performance. **Second, our framework introduces a diversity-guidance mechanism that applies repulsive forces during diffusion inference, pushing latent trajectories apart within a batch to explore multiple diverse reasoning paths, whereas AR models tend to collapse to similar trajectories.** Finally, leveraging a VAE-based latent space enhances interpretability over continuous diffusion models, making the reasoning process more transparent and readable.

Experimentally, we demonstrate that diffusion-based latent reasoning is not only more accurate but also qualitatively different from prior approaches. On math reasoning benchmarks, including GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021), where Coconut (Hao et al., 2024) fails to surpass AR CoT supervised finetuning (SFT), LaDiR consistently outperforms it on average across 7 benchmarks with the LLaMA 3.1 8B model (Dubey et al., 2024). This suggests that modeling reasoning at the *semantic level*, rather than at the token level, may lead to more faithful intermediate steps that accumulate into stronger final answers. Moreover, on the Countdown planning task, LaDiR shows over 30% absolute improvement in both Pass@1 and Pass@100, indicating that latent thought tokens potentially enhance global planning ability, while *parallel diversity exploration* enables the model to generate diverse reasoning paths. Together, these findings suggest that diffusion-based latent reasoning provides a principled way to balance accuracy and diversity—key ingredients for advancing beyond sequential autoregressive reasoning.

## 2 PRELIMINARIES

This section introduces key concepts and notations in VAE and latent diffusion models (Rombach et al., 2022). Detailed formulations and background information are provided in Appendix B.

### 2.1 VARIATIONAL AUTOENCODER

A Variational Autoencoder (VAE) (Kingma & Welling, 2013) learns a latent representation of data by balancing reconstruction accuracy and prior regularization. Let  $x \in \mathbb{R}^{L \times d_x}$  denote a sequence of token embeddings with length  $L$  and embedding dimension  $d_x$ , and  $z \in \mathbb{R}^{M \times d_z}$  denote latent representations with  $M$  latent tokens of dimension  $d_z$ . We adopt the  $\beta$ -VAE (Higgins et al., 2017), where a scaling factor  $\beta$  controls this trade-off:

$$\mathcal{L}_{\beta\text{-VAE}} = \mathbb{E}_{q_{\phi}(z|x)}[-\log p_{\theta}(x|z)] + \beta \text{KL}(q_{\phi}(z|x) \parallel p(z)). \quad (1)$$

108 Here,  $q_\phi(z|x)$  is the encoder distribution parameterized by  $\phi$ ,  $p_\theta(x|z)$  is the decoder likelihood  
109 parameterized by  $\theta$ , and  $p(z) = \mathcal{N}(0, I)$  is the prior distribution over latents.  
110

111 Larger  $\beta$  values encourage disentangled and structured latent spaces, at the cost of reconstruction  
112 fidelity. During inference, the encoder of VAE produces a mean/variance pair  $\{(\mu, \sigma)\}$ , and a latent  
113 token  $z$  is sampled as  $z = \mu + \sigma \odot \epsilon$ ,  $\epsilon \sim \mathcal{N}(0, I)$ .  
114

## 2.2 LATENT DIFFUSION AND FLOW MATCHING

115 Latent diffusion models (Ho et al., 2020; Rombach et al., 2022) generate data by denoising latent  
116 variables from Gaussian noise in the latent space of a VAE, which preserves high-level semantic  
117 structure. Diffusion can also be viewed as a continuous-time generative flow trained via *flow matching*  
118 (Lipman et al., 2022), which we adopt as our primary framework for its superior performance  
119 (see Appendix C). The training and inference processes are as follows:  
120

121 **Training** Let  $\{z_t\}_{t \in [0,1]}$  denote a path interpolating between clean data  $z_0 \sim p_{\text{data}}$  and noise  
122  $\epsilon \sim \mathcal{N}(0, I) : z_t = (1-t)z_0 + t\epsilon$ . This path is controlled by an ordinary differential equation  
123 (ODE)  $u^*(z_t, t) = \frac{dz_t}{dt} = \epsilon - z_0$ , where  $u^*$  is the target velocity field. A neural network  $u_\theta(z_t, t)$  is  
124 trained to approximate  $u^*$  by minimizing the flow matching loss:  
125

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{t \sim \mathcal{U}(0,1), z_0 \sim p_{\text{data}}, z_1 \sim \mathcal{N}(0,I)} \left[ \|u_\theta(z_t, t) - u^*(z_t, t)\|^2 \right]. \quad (2)$$

127 **Inference.** At generation time, the process begins from Gaussian noise  $z_1 \sim \mathcal{N}(0, I)$ . The learned  
128 velocity field  $u_\theta(z_t, t)$  is then integrated backward in time using an ODE solver as follows:  $z_{t-\Delta t} =$   
129  $z_t - \Delta t u_\theta(z_t, t)$ , with steps from  $t = 1$  to  $t = 0$ . The final state  $z_0$  corresponds to a clean  
130 latent representation. This procedure naturally supports *iterative refinement*, as each integration step  
131 progressively transforms noise into a coherent latent  $z$ .  
132

## 2.3 BLOCK DIFFUSION

133 To support flexible and variable-length sequence generation, we employ a *block diffusion*  
134 scheme (Arriola et al., 2025) that integrates autoregressive modeling with diffusion. Instead of  
135 applying diffusion to individual latent tokens or full sequence, the sequence is divided into contiguous  
136 blocks, and diffusion is performed at the block level. This hybrid design retains the open-ended  
137 generation of autoregressive models while introducing global coherence within each block. See  
138 Appendix B.4 for details.  
139

## 3 METHODOLOGY

140 Our approach separates reasoning from answering. A variational autoencoder (VAE) constructs a  
141 latent space of intermediate reasoning steps, encoding each step as a block of *thought tokens*. We  
142 further utilize a reasoning model that predicts and refines *thought tokens* via latent diffusion, and  
143 then generates the final answer tokens conditioned on the denoised latent tokens.  
144

### 3.1 ARCHITECTURE

145 We employ a VAE to construct the latent space of intermediate reasoning steps, and a reasoning  
146 model that predicts latent tokens via diffusion and generates the final text answer.  
147

148 **Blockization.** We separate the chain-of-thought (CoT) reasoning and the final answer in the  
149 dataset using the prefix ‘‘The answer is’’. The text preceding the prefix is treated as CoT  $c$ ,  
150 while the text following is treated as the final answer  $y$ . We then split  $c$  into individual sentences,  
151 each treated as a *block* of latent tokens with block size  $L_b$ :  
152

$$\mathbf{Z}^{(b)} = \{z_1^{(b)}, \dots, z_{L_b}^{(b)}\}, \quad b = 1, \dots, N.$$

153 This one-sentence-per-block design ensures that each reasoning step is localized in latent space.  
154

155 **VAE architecture.** As shown in Figure 2 (left), our VAE encoder is initialized from a pretrained  
156 LLM and fine-tuned with all parameters, along with  $L_b$  learnable embeddings. The encoder’s last  
157 hidden state is passed through two linear projections to obtain the mean  $\mu$  and variance  $\sigma^2$ , from  
158 which we sample  $\mathbf{Z}^{(b)} \sim \mathcal{N}(\mu, \sigma^2)$ . The decoder is a *frozen* pretrained LLM that conditions on  
159 the sampled  $\mathbf{Z}^{(b)}$  to reconstruct the corresponding block of text. This design enables the encoder to  
160 compress each reasoning step into a structured latent representation aligned with the semantic space  
161 of the language model.  
162

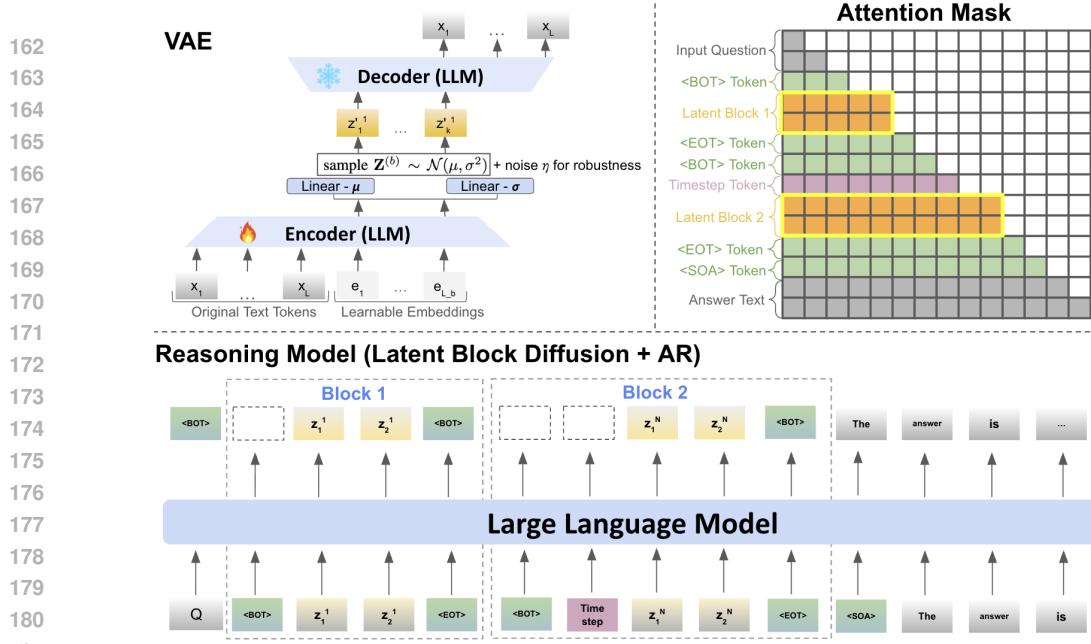


Figure 2: Illustration of our block-wise latent reasoning framework. A question  $Q$  is first input as condition to generate latent blocks, each delimited by  $<\text{BOT}>$  and  $<\text{EOT}>$ . For each block, the model iteratively denoises latent tokens  $\hat{\mathbf{Z}}^{(b)}$  across timesteps, with bidirectional attention inside a block and causal attention across blocks. The reasoning process terminates when the model emits the  $<\text{SOA}>$  token, after which the model generates the answer text autoregressively.

**Reasoning model architecture.** We utilize an existing LLM as our reasoning model. As illustrated in Fig. 2, consider the prediction of the second latent block. After the input question  $Q$ , we insert a special token  $<\text{BOT}>$  to mark the start of a block, followed by the first block tokens  $\mathbf{Z}^{(1)}$ , and a token  $<\text{EOT}>$  to mark its end. For the second block, since it is being predicted, we add a timestep embedding between  $<\text{BOT}>$  and  $z_1^{(2)}$  to encode the timestep information. Once the latent reasoning process is complete, we switch to text generation mode by appending a  $<\text{SOA}>$  token to indicate the start of the answer, which is then generated autoregressively. To balance lookahead and variable-length generation, we adopt a hybrid attention mask  $\mathcal{M}$  (Fig. 2, top right). Within each block, tokens attend *bidirectionally*, enabling the model to internally *reason* over a horizon defined by the block size and capture richer local dependencies. Across blocks, attention is strictly *causal*, so later steps depend on earlier ones in an autoregressive manner.

### 3.2 TRAINING

We train the two components separately: the VAE is first trained to learn latent representations of *thought tokens*, after which the reasoning model is trained to predict these *thought tokens*. We describe each stage in turn, beginning with VAE training and followed by reasoning model training.

#### 3.2.1 VAE TRAINING

We build on the standard  $\beta$ -VAE training and inference framework described in Section 2.1, with the following adaptations tailored to our task.

**Robustness augmentations.** To improve generalization and make the latent space resilient to noise and input variability, we introduce two augmentation strategies during training:

- **Latent Gaussian noise.** For each latent token  $z_i^{(b)}$ , we inject isotropic Gaussian perturbations:

$$z_i'^{(b)} = z_i^{(b)} + \eta_i, \quad \eta_i \sim \mathcal{N}(0, k^2 I)$$

where we find  $k = 3$  achieve the best downstream performance. This enhances robustness by smoothing the latent space and mitigating sensitivity to small semantic variations.

- **Input token substitution.** For the encoder input sequence, with probability  $p = 0.3$  we replace a token with another randomly chosen token (sampled uniformly from the LLM vocabulary). This

216 forces the encoder to learn invariances to paraphrasing, typos, or minor corruptions in input text,  
217 ensuring that latent representations capture semantic content rather than exact lexical form.  
218

219 Together, these augmentations encourage the VAE to build a smoother latent space that is both  
220 robust to perturbations and expressive enough to encode thought-level reasoning steps. A more  
221 detailed diagram of the VAE can be seen in Appendix F.

### 222 3.3 REASONING MODEL TRAINING

223 After constructing a latent reasoning space with the VAE, we train a latent diffusion model  $f_\psi$  from  
224 the same pretrained LLM as in our VAE model to denoise latent blocks, gradually transforming  
225 noisy latent representations to coherent reasoning blocks. Empirically, we observe that training with  
226 the flow-matching objective yields the best performance (see Appendix C), and therefore adopt it as  
227 our default training objective in the paper.

228 **Answer Token Loss.** While  $f_\psi$  learns to predict latent reasoning trajectories, to avoid explicitly  
229 decoding these steps through the VAE decoder for efficiency during inference, we use the same  
230 [transformer backbone  \$\psi\$  with a LM head](#) to autoregressively predict answer text tokens conditioned  
231 on the latent reasoning blocks. To this end, given the question  $q$ , the reasoning blocks  $\mathbf{Z}^{(\leq B)}$ ,  
232 and the past answer tokens  $y_{<w}$ , the model predicts the next answer token  $y_w$  with distribution  
233  $p_\psi(y_w | q, \mathbf{Z}^{(\leq B)}, y_{<w})$ . The training objective for those answer tokens is the cross-entropy loss:  
234

$$235 \quad \mathcal{L}_{\text{Ans}} = - \sum_{w=1}^W \log p_\psi(y_t | q, \mathbf{Z}^{(\leq B)}, y_{<w}). \quad (3)$$

236

237 **Special Token Loss.** To explicitly control the number of latent blocks, we introduce a special  
238 binary classification head on top of the same LLM transformer backbone  $\psi$ . It predicts whether  
239 the next block begins with a  $\langle \text{SOA} \rangle$  (start-of-answer) or  $\langle \text{BOT} \rangle$  (begin-of-thought) token when-  
240 ever an  $\langle \text{EOT} \rangle$  (end-of-thought) token is generated. Formally, let  $\tau$  index positions of  $\langle \text{EOT} \rangle$   
241 tokens in the output. For each  $\tau$ , the model produces a distribution  $p_\psi(s_\tau | q, \mathbf{Z}^{(\leq B)}, y_{\leq \tau})$ ,  $s_\tau \in$   
242  $\{\langle \text{SOA} \rangle, \langle \text{BOT} \rangle\}$ , and we minimize the corresponding classification loss, [which supervises the](#)  
243 [model to predict special tokens given the question  \$q\$  and latent reasoning blocks up to position  \$\tau\$](#) :  
244

$$245 \quad \mathcal{L}_{\text{Spec}} = - \sum_{\tau \in \mathcal{T}_{\text{EOT}}} \log p_\psi(s_\tau | q, \mathbf{Z}^{(\leq B)}). \quad (4)$$

246

#### 247 3.3.1 STAGE 1: TEACHER-FORCING TRAINING

248 In the first stage, the model is trained under a *teacher-forcing* regime, where it has access to oracle  
249 latent blocks produced by the VAE encoder, denoted as  $\mathbf{Z}^{(1:B)}$ . At every step, these oracle latents  
250 are concatenated between special tokens  $\langle \text{BOT} \rangle$  and  $\langle \text{EOT} \rangle$  and provided as context to the flow-  
251 matching model  $f_\psi$ . The overall training objective jointly optimizes flow matching on latent blocks  
252 and cross-entropy supervision on both final answers and special tokens:  
253

$$254 \quad \mathcal{L} = \lambda_{\text{FM}} \mathcal{L}_{\text{FM}} + \lambda_{\text{Ans}} \mathcal{L}_{\text{Ans}} + \lambda_{\text{Spec}} \mathcal{L}_{\text{Spec}}, \quad (5)$$

255 where  $\mathcal{L}_{\text{FM}}$  is defined in Eq. 2.

#### 256 3.3.2 STAGE 2: ROLLOUT TRAINING

257 After Stage 1, there is a mismatch between training and inference. During inference, the model  
258 must be conditioned on previous self-generated latents without access to oracle latents, suffering  
259 from error accumulation issue. To address this issue, Stage 2 adopts an *rollout* training. We keep  
260 the same number of blocks  $B$  as in the ground truth, but instead of conditioning on oracle latents,  
261 the model generates its own latents  $\tilde{\mathbf{Z}}^{(1:B)}$  from random noise using a fewer denoising steps (i.e.,  
262  $50 \rightarrow 10$ , following FlowGRPO (Liu et al., 2025)). We keep the gradients on  $\tilde{\mathbf{Z}}^{(1:B)}$  during denoising,  
263 allowing answer supervision to backpropagate through the trajectory and directly *shape* latent  
264 predictions. To avoid latent collapse as in Coconut w/o curriculum learning (Hao et al., 2024), we  
265 keep the flow matching loss. Therefore, the training objective is same as Eq. 5.

### 266 3.4 REASONING MODEL INFERENCE

267 At inference time, the model generates a chain of latent reasoning blocks and subsequently produces  
268 the final answer in text space. The process unfolds in two phases: (i) latent block generation via  
269 iterative denoising, and (ii) answer generation via autoregressive decoding.

270 **Iterative denoising.** Following the inference process of the standard latent diffusion model as in  
271 2.2, for each block, we initialize with a Gaussian noise, and we gradually transforms the noise into  
272 a semantically coherent latent reasoning block  $\hat{\mathbf{Z}}^{(b)}$ .  
273

274 **Stopping criterion.** Latent block generation continues until the model explicitly predicts the spe-  
275 cial token  $\langle \text{SOA} \rangle$  (*start of answer*). This token signals that sufficient reasoning has been performed  
276 and the model should transition from block diffusion to final answer generation.

277 **Answer generation.** Once reasoning terminates, conditioned on the generated latent reasoning  
278 sequence  $\hat{\mathbf{Z}}^{(1:\hat{B})}$  and the input question  $x$ , the model predicts output tokens  $y = (y_1, \dots, y_T)$  au-  
279 toregressively.

280 **Diversity improvement in parallel.** Unlike AR models that generate a single reasoning trajectory  
281 sequentially, our framework can generate multiple diverse reasoning trajectories in parallel within a  
282 batch. To encourage exploration of alternative solutions, we incorporate two complementary mech-  
283 anisms:

- 284 1. **Increased initial noise.** By sampling with an increased variance  $\tilde{\sigma}^2$  as initial noise scale, we  
285 broaden the distribution of starting points for latent trajectories. This enables the same input  
286 question to yield diverse reasoning sequences across runs, improving coverage of alternative  
287 solution strategies.
- 288 2. **Diversity gradient guidance.** At each denoising step, we enhance diversity by adding a repul-  
289 sion term to push the latent tokens in a batch apart. First, we compute a bandwidth parame-  
290 ter  $\sigma$  as the median pairwise distance between the latent tokens in a batch at the current step  
291  $\sigma = \text{median}_{i < j} \|z_i - z_j\|_2$ .

293 The repulsion force field for a latent token  $z_i$  is then defined as

$$294 \mathbf{F}(z_i) = \sum_{j \neq i} 2 \left( 1 - \frac{\|z_i - z_j\|_2^2}{\sigma^2} \right) \exp \left( - \frac{\|z_i - z_j\|_2^2}{\sigma^2} \right) (z_i - z_j), \forall j \leq B, \quad (6)$$

295

296 where  $z_j$  is any other latent token in the same batch with batch size  $B$ . We apply strong repulsion  
297 at the beginning of inference and gradually decay its effect over time. Specifically, the time-  
298 dependent scale is defined as  $\gamma_t = \gamma_{\max} \left( \frac{t}{T} \right)$ , where  $T$  is the total number of inference steps,  $t$   
299 decreases from  $T$  to 0,  $\gamma_{\max}$  is the initial repulsion strength as a hyperparameter.  
300

301 Finally, the diversity-guided prediction combines the base model output with the repulsion  
302 gradient, in a form analogous to classifier-free guidance (Ho & Salimans, 2022):  $\hat{z}_{t-1} =$   
303  $f_{\psi}(\mathbf{x}_t, t, x) + \gamma_t \mathbf{F}(z)$ , where  $f_{\psi}(\mathbf{x}_t, t, x)$  is the model’s prediction at step  $t$ .

304 Together, these mechanisms enhance the stochasticity and coverage of latent reasoning while pre-  
305 serving convergence to valid solutions.  
306

## 307 4 EXPERIMENTS

308 We evaluate LaDiR across two domains: mathematical reasoning (7 datasets) and puzzle planning  
309 (Countdown), comparing to AR, latent, and diffusion baselines. Our experiments demonstrate its  
310 effectiveness on benchmark datasets, while ablation studies in Section 4.3 and analyses 4.4 provide  
311 further insights into the contributions of individual components. See Appendix E for experimental  
312 details.  
313

### 314 4.1 MATHEMATICAL REASONING

315 We begin by assessing LaDiR on a range of mathematical reasoning benchmarks, covering both  
316 in-domain datasets, where training and test distributions are closely aligned, and out-of-domain  
317 benchmarks that require generalization to unseen problems.  
318

**Datasets** We fine-tune pretrained LLMs on the **DART-MATH** dataset (Tong et al., 2024b), a  
319 large-scale dataset synthesized to enhance mathematical reasoning. For evaluation, we adopt two in-  
320 domain benchmarks, **Math** (Hendrycks et al., 2021) and **GSM8K** (Cobbe et al., 2021), and five out-  
321 of-domain benchmarks to assess generalization: **College-Math** (Tang et al., 2024b), **DeepMind-  
322 Math** (Saxton et al., 2019), **OlympiaBench-Math** (He et al., 2024), **TheoremQA** (Chen et al.,  
323 2023), and **Fresh-Gaokao-Math-2023** (Tang et al., 2024b). Detailed dataset descriptions are pro-  
324 vided in Appendix E.  
325

Method	In-Domain			Out-of-Domain				Avg.
	MATH	GSM8K	Gaokao	DM-Math	College	Olympia	TheoremQA	
Masked Diffusion Methods - LLaDA 8B								
Base Model	36.2	77.3	10.2	29.8	30.2	3.0	16.6	29.0
CoT SFT	39.0	82.3	20.1	43.7	38.9	5.9	20.9	35.8
Autoregressive Methods - LLaMA 3.1 8B								
Sol-Only SFT	13.3	16.4	0.0	18.2	15.9	4.7	16.9	12.2
CoT SFT	43.1	<b>84.5</b>	30.7	<b>47.8</b>	45.7	10.1	21.2	40.4
iCoT	35.2	61.8	30.0	30.6	37.6	8.3	19.5	31.8
Pause Token	42.1	83.9	28.3	42.4	31.5	3.5	8.3	34.2
Coconut	37.3	68.3	26.8	33.5	40.2	5.8	11.4	31.9
Discrete Latent	43.2	83.9	33.3	44.7	47.1	<b>13.3</b>	20.3	40.8
Latent Diffusion Methods - LLaMA 3.1 8B								
LD4LG	<b>32.9</b>	<b>9.1</b>	-	-	-	-	-	-
PLANNER	18.7	<b>5.7</b>	-	-	-	-	-	-
<b>LaDiR</b> (ours)	<b>45.2</b>	84.2	<b>33.4</b>	46.3	<b>48.6</b>	11.9	<b>22.9</b>	<b>41.8</b>
<b>–w/o Stage 2</b>	30.7	57.8	24.7	32.0	32.8	5.9	11.9	32.6

Table 1: Results for pass@1 accuracy across in-domain and out-of-domain math benchmarks.

**Baselines** We compare our approach against a diverse set of reasoning methods spanning both autoregressive and diffusion-based paradigms. For autoregressive models, we include **Sol-Only**, trained only on question–solution pairs without intermediate reasoning steps, and **CoT**, trained with full chain-of-thought supervision. We further evaluate several latent reasoning methods: **Implicit CoT (iCoT)** (Deng et al., 2023), which gradually removes explicit CoT tokens through curriculum learning; **Pause Token** (Goyal et al., 2023), which introduces a learnable pause token to provide additional computation before answering; **Coconut** (Hao et al., 2024), which leverages hidden states as latent reasoning tokens through curriculum learning; and **Discrete Latent** (Su et al., 2025), which compresses a span of text (e.g., 16 tokens) into discrete latent codes via VQ-VAE for reasoning. For diffusion-based language models, we compare with the open-sourced **LLaDA 8B** (Nie et al., 2025), evaluated both with and without SFT. To directly compare with prior latent diffusion methods for language generation, we also evaluate **LD4LG** (Lovelace et al., 2024) and **PLANNER** (Zhang et al., 2023). For fair comparison, we use FLAN-T5 as the encoder and LLaMA-3.1-8B as the decoder, and train them on the same reasoning datasets. We utilize LLaMA-3.1 8B (Dubey et al., 2024) as the backbone model for our framework as well as for the AR baselines for fair comparison.

**Results** As shown in Table 1, our method achieves the strongest overall performance on average, improving the average pass@1 accuracy by 2% over the best prior latent approach. Also, LaDiR achieves higher Pass@100 across all benchmarks, with a 6.1% absolute gain over AR CoT SFT on average (see Table 8 in Appendix). Compared to text-based CoT baselines, our latent reasoning consistently yields more robust solutions, particularly on harder benchmarks such as DM-Math and College-level datasets, where direct text reasoning often struggles with long-horizon consistency. This suggests that reasoning in a latent space at semantic level learns more abstract reasoning patterns. Compared to prior latent approaches (e.g., Coconut), the latent diffusion objective provides a more principled objective for modeling continuous trajectories, leading to stronger generalization to out-of-domain settings such as TheoremQA. Also, incorporating the stage 2 rollout training notably improves performance across all benchmarks, showing its effectiveness in mitigating error accumulation. Moreover, prior latent diffusion methods of LD4LG and PLANNER perform poorly on reasoning tasks, indicating that effective latent reasoning requires more than architectural changes, such as blockwise variable-length diffusion and rollout training for answer alignment. Taken together, these results indicate that LaDiR combines the interpretability benefits of CoT-style reasoning (see Appendix D) and expressiveness of continuous latent space, producing generalizable reasoning traces.

## 4.2 PUZZLE PLANNING – COUNTDOWN

We evaluate the planning ability of our method using *Countdown*, a combinatorial arithmetic game. Given a set of input numbers, the goal is to reach a target in [10, 100] by applying basic operations  $\{+, -, \times, \div\}$ . Solving a problem thus demands decomposing the target into intermediate subgoals and chaining them correctly. For example, given input numbers  $\{97, 38, 3, 17\}$  and target 14, one valid solution is:  $97 - 38 = 59$ ,  $59 - 17 = 42$ ,  $42 \div 3 = 14$ . Following Gandhi et al. (2024), we construct a dataset of 500k examples, holding out 10% of target numbers for *out-of-distribution*

Model	CD-4 P@1	CD-4 P@100	CD-4 Div.	CD-5 P@1	CD-5 P@100	CD-5 Div.
Dream 7B Base*	16.0	24.7	4.1	4.2	10.3	5.6
MGDM <sup>†</sup>	<b>91.5</b>	<u>95.2</u>	3.2	<b>46.6</b>	70.4	4.9
LLaDA 8B SFT	<u>51.2</u>	<u>75.2</u>	<u>5.4</u>	<u>34.4</u>	<u>45.2</u>	<u>6.2</u>
LLaMA 8B SFT	46.7	65.3	3.0	8.9	15.4	3.5
<b>LaDiR</b>	<b>76.6</b>	<b>96.4</b>	<b>7.3</b>	<b>38.5</b>	<b>75.2</b>	<b>8.9</b>

Table 2: Results on Countdown tasks. We report Pass@1, Pass@100, and Diversity (Div.). Best results are in **bold**, and second-best are underlined. \*Dream 7B Base refers to the open-sourced base model without finetuning on this task. <sup>†</sup>MGDM is a task-specific small discrete diffusion model rather than a general-purpose language model.

evaluation as test set. We study two settings of growing complexity: CD-4 and CD-5, which use four and five input numbers, respectively.

**Baselines** We compare our method against both autoregressive and diffusion-based approaches. For the autoregressive setting, we include (1) LLaMA 8B SFT, which shares the same base model as ours and is finetuned on the same dataset. For diffusion-based baselines, we consider (2) LLaDA 8B SFT (Nie et al., 2025) and (3) Dream 7B Base (Ye et al., 2025b), two diffusion-based general-purpose language models (the latter evaluated without finetuning), as well as (4) MGDM (Ye et al., 2024a), a small task-specific multinomial-guided diffusion model trained for Countdown.

**Metrics.** We report Pass@1 and Pass@100 accuracy, using an exact string match between the generated arithmetic equations and the ground-truth solution. Pass@ $k$  reflects the accuracy that at least one valid solution is found among  $k$  samples. In addition, we report Diversity, measured as the number of unique valid solutions discovered among 100 samples. All models are evaluated with a decoding temperature of 1.0.

**Implementation Details** In this setting, we deliberately *disable* the answer generation and restrict the reasoning process to a single latent block, which is compressed into a fixed-size representation (4 tokens). The model is trained under a teacher-forcing regime and evaluated on decoded text tokens from our VAE, thereby isolating the latent diffusion model’s capacity to capture planning dynamics without autoregressive supervision. During inference, we set the initial noise scale to 2 and the maximum diversity guidance scale to 0.8.

**Results** On the Countdown tasks, as shown in Table 2, our method outperforms autoregressive baselines and remains competitive with specialized diffusion models. In CD-4, it improves Pass@1 by more than 25 points over LLaMA 8B SFT and over 20 points over LLaDA SFT, demonstrating stronger *planning* ability beyond token-by-token generation, while also delivering the best Pass@100 and over two points higher *diversity* than any baseline. On the more challenging CD-5 task, our model surpasses AR baselines by nearly 30 points in Pass@1 and over 30 points in Pass@100, again with the highest accuracy. In addition, as shown in Figure 3, our pass@ $k$  curve rises steeply with  $k$ , surpassing MGDM at larger  $k$ . This high pass@ $k$  reflects both diverse trajectory exploration and strong potential for reinforcement learning for post-training (Yue et al., 2025).

#### 4.3 ABLATION STUDY

**Diversity Scale and Initial Noise** We study how inference-time stochasticity and diversity guidance affect solution diversity and accuracy by varying (i) the *initial noise scale*, which controls the variance of Gaussian initialization, and (ii) the *maximum diversity scale*  $\gamma_{\max}$ , which regulates the repulsion strength among latent tokens (see Sec. 3.4). We evaluate both the average number of unique solutions and best-of-100 accuracy. Table 4 shows that increasing noise from 1 to 2 improves both diversity and accuracy, but excessive noise (scale 3) harms convergence despite higher diversity. For diversity guidance, removing repulsion ( $\gamma_{\max} = 0$ ) yields the lowest diversity, while moderate values (0.3–0.5) strike the best trade-off. Stronger repulsion ( $\gamma_{\max} \geq 1.0$ ) further boosts

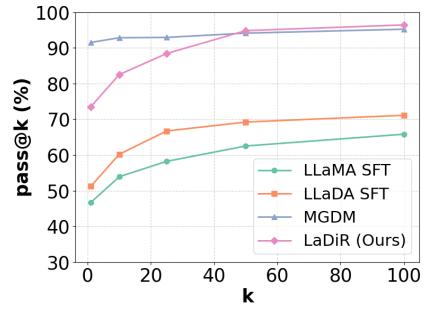
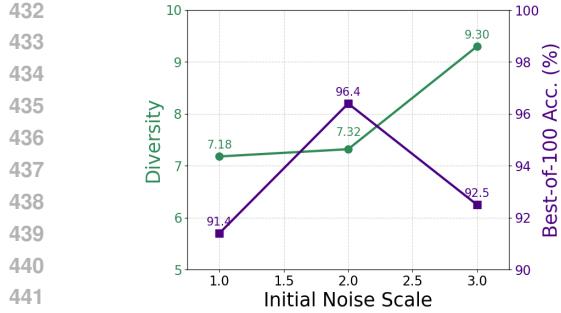
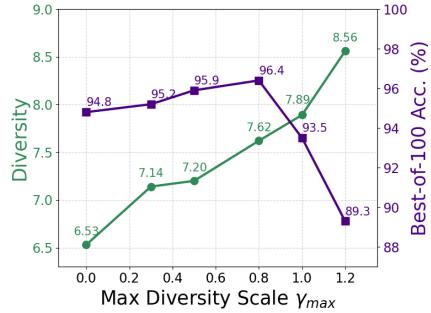


Figure 3: Results for  $\text{pass}@k$  performance on Countdown-4 with  $k \in \{1, 10, 25, 50, 100\}$ .



(a) Effect of Initial Noise Scale  $\sigma^2$  ( $\gamma_{max} = 0.8$ ).



(b) Effect of Max Diversity Scale  $\gamma_{max}$  ( $\sigma^2 = 2$ ).

Figure 4: Ablation study on the hyperparameters for diversity during inference on Countdown-4.

diversity but causes accuracy to drop, suggesting that over-dispersing latents destabilizes reasoning. See Appendix C for additional ablation studies.

#### 4.4 ANALYSIS

**Iterative Refinement** Table 3 shows how the our flow-matching model refines its reasoning across denoising steps. From pure noise at  $T = 1$ , the model quickly produces structured equations, though early steps contain arithmetic errors (e.g., off-by-one mistakes). As denoising progresses, partial results stabilize—such as  $43 + 9 = 52$  appearing consistently from  $T = 0$  onward—and later steps are gradually corrected. By  $T = 0.25$ , the full reasoning matches the ground truth and remains stable through  $T = 0$ . This demonstrates that our method exhibits the same iterative refinement ability as reasoning models (Shao et al., 2024), progressively correcting previously generated steps. See Table 10 for an example on GSM8K.

**Adaptive Test-Time Compute.** As shown in Figure 5, using more denoising steps consistently improves accuracy across different math benchmarks. For example, increasing from 5 to 10 steps (a  $2\times$  compute increase) yields a large jump of +11.7 points in accuracy on average of 7 benchmarks. Starting from 10 steps, tripling the compute to 30 steps provides an additional +4.8 points on average, while a  $5\times$  compute increase to 50 steps brings a total gain of +9.8 points on average. These results demonstrate that our method can flexibly trade test-time compute for higher performance as an alternative paradigm in reasoning for long CoT of existing reasoning models (Jaech et al., 2024; Muennighoff et al., 2025; Liu et al., 2024a; Li et al., 2025). This may motivate adaptive policies that dynamically assign more denoising steps to harder queries, maximizing the overall accuracy–compute trade-off.

**Inference Efficiency** Table 4 compares inference latency on the MATH dataset using identical hardware and batch size (8). LaDiR with 10 diffusion steps matches the latency of the AR baseline while achieving comparable Pass@1 and higher Pass@100. With 30 steps, LaDiR offers a flexible accuracy–compute trade-off. This efficiency stems from our compact latent representation: each latent block contains only 4 latent tokens, representing on average 22 text tokens, reducing per-step

<b>Input</b>	43, 9, 54, 25, 81
<b>GT Answer</b>	43+9=52, 54-25=29, 52+29=81
Decode( $\hat{Z}_{t=1}$ )	“I .. ex1 ...” (random noise)
Decode( $\hat{Z}_{t=0.8}$ )	43+8=51, 54-24=30, 51+30=81
Decode( $\hat{Z}_{t=0.7}$ )	43+10=53, 54-25=28, 53+28=81
Decode( $\hat{Z}_{t=0.6}$ )	43+9=52, 54-27=27, 52+27=79
Decode( $\hat{Z}_{t=0.5}$ )	43+9=52, 54-25=29, 52+28=80
Decode( $\hat{Z}_{t=0.4}$ )	43+9=52, 54-25=29, 52+30=82
Decode( $\hat{Z}_{t=0.25}$ )	43+9=52, 54-25=29, 52+29=81
Decode( $\hat{Z}_{t=0}$ )	43+9=52, 54-25=29, 52+29=81

Table 3: Examples of iterative self-refinement of decoded text from the VAE decoder on the Countdown-4 dataset across different denoising timesteps ( $t$ ).

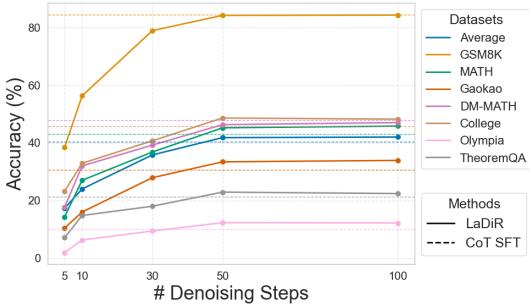


Figure 5: Effect of number of denoising steps on downstream reasoning performance on the math reasoning tasks.

486 computation and context length compared to autoregressive decoding over long text sequences. We  
487 provide additional analyses on interpretability and semantic-level reasoning in Appendix D.  
488

## 489 5 RELATED WORKS

490 **Latent Reasoning** Latent reasoning  
491 methods address token-level limits of  
492 chain-of-thought by enabling reasoning  
493 in a latent space, yielding more abstract  
494 representations through discrete special  
495 tokens that expand internal computation or  
496 capture unstated steps (Herel & Mikolov,  
497 2024; Pfau et al., 2024; Wang et al., 2024;  
498 Zelikman et al., 2024; Jin et al., 2025).

499 Prior works show that reasoning in latent  
500 space, rather than discrete tokens, improves performance by allowing LLMs to generate continuous  
501 tokens, either self-generated or provided by an auxiliary model (Cheng & Durme, 2024; Hao et al.,  
502 2024; Liu et al., 2024b; Shen et al., 2025; Tack et al., 2025; Zhu et al., 2025; Butt et al., 2025; Zhang  
503 et al., 2025; Wu et al., 2025). This has been further extended to recurrent or looped architectures  
504 that induce latent reasoning internally, removing the need to represent reasoning steps explicitly as  
505 tokens (Chen et al., 2025c; Geiping et al., 2025; Mohtashami et al., 2025; Saunshi et al., 2025; Yu  
506 et al., 2025b). However, prior latent reasoning approaches lack interpretability, as their continuous  
507 states are opaque and difficult to understand or control. whereas our method structures the latent  
508 space with a VAE, making each step explicit and thus more transparent.

509 **Latent Diffusion for Language Generation** Generative text modeling has recently expanded  
510 from autoregressive paradigms to diffusion-based approaches that allow for global iterative refine-  
511 ment. One of the first, Diffusion-LM (Li et al., 2022), frames generation as denoising continuous  
512 word embeddings to enable fine-grained control, a concept Lovelace et al. (2023); Lovelace et al.  
513 extended by performing diffusion in a compressed latent space for improved quality and diverse  
514 generation modes. For sequence-to-sequence tasks, DiffuSeq (Gong et al., 2022) enables parallel  
515 generation with high diversity, while PLANNER (Zhang et al., 2023) addresses long-form text by  
516 combining a latent semantic diffusion planner with an autoregressive decoder to reduce repetition.  
517 Similarly, Cosmos (Meshchaninov et al., 2025) learns a compressed latent space for diffusion, en-  
518 abling parallel text generation with robust semantic grounding. In specialized domains, Diffusion-  
519 Dialog (Xiang et al., 2024) utilizes latent variables to handle open-ended conversations, whereas  
520 CodeFusion (Singh et al., 2023) and TreeDiff (Zeng et al., 2025) apply diffusion to code synthesis.  
521 While prior latent diffusion models focus on text generation, they lack the granularity to model the  
522 multi-step causal dependencies required for reasoning tasks. We address this by introducing block-  
523 wise variable-length diffusion and rollout training, explicitly shifting the objective from generating  
524 fluent text to optimizing reasoning trajectories that lead to correct answers. Due to the page limit,  
525 we discuss further related works in Appendix A.

## 526 6 CONCLUSION

527 We introduced LaDiR, a latent diffusion reasoner that utilizes the iterative refinement capability  
528 of latent diffusion models to perform reasoning at the semantic level, our framework offers three  
529 key benefits: (1) better tradeoff between accuracy and test-time compute through iterative denois-  
530 ing steps with self-refinement, (2) parallel and diverse exploration of reasoning trajectories beyond  
531 the limitations of sequential autoregression, and (3) enhanced interpretability through semantically  
532 meaningful latent representations. Our experiments on mathematical reasoning and planning bench-  
533 marks show that LaDiR consistently outperforms AR and diffusion baselines, achieving both higher  
534 accuracy and greater diversity in reasoning.

## 535 REFERENCES

536 Marianne Arriola, Aaron Gokaslan, Justin T Chiu, Zhihan Yang, Zhixuan Qi, Jiaqi Han, Sub-  
537 ham Sekhar Sahoo, and Volodymyr Kuleshov. Block diffusion: Interpolating between autore-  
538 gressive and diffusion language models. *arXiv preprint arXiv:2503.09573*, 2025.  
539 Gregor Bachmann and Vaishnav Nagarajan. The pitfalls of next-token prediction, 2025. URL  
540 <https://arxiv.org/abs/2403.06963>.

---

540 Zhenni Bi, Kai Han, Chuanjian Liu, Yehui Tang, and Yunhe Wang. Forest-of-thought: Scaling test-  
541 time compute for enhancing llm reasoning, 2025. URL <https://arxiv.org/abs/2412.09078>.

542

543 Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo  
544 Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke,  
545 Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi,  
546 James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky.  $\pi_0$ : A vision-  
547 language-action flow model for general robot control, 2024. URL <https://arxiv.org/abs/2410.24164>.

548

549

550 Natasha Butt, Ariel Kwiatkowski, Ismail Labiad, Julia Kempe, and Yann Ollivier. Soft tokens, hard  
551 truths. *arXiv preprint arXiv:2509.19170*, 2025.

552

553 Edoardo Cetin, Tianyu Zhao, and Yujin Tang. Large language models to diffusion finetuning. *arXiv  
554 preprint arXiv:2501.15781*, 2025.

555

556 Shoufa Chen, Chongjian Ge, Yuqi Zhang, Yida Zhang, Fengda Zhu, Hao Yang, Hongxiang Hao,  
557 Hui Wu, Zhichao Lai, Yifei Hu, Ting-Che Lin, Shilong Zhang, Fu Li, Chuan Li, Xing Wang,  
558 Yanghua Peng, Peize Sun, Ping Luo, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Xiaobing Liu.  
559 Goku: Flow based video generative foundation models, 2025a. URL <https://arxiv.org/abs/2502.04896>.

560

561 Wenhui Chen, Ming Yin, Max Ku, Pan Lu, Yixin Wan, Xueguang Ma, Jianyu Xu, Xinyi Wang,  
562 and Tony Xia. Theoremqa: A theorem-driven question answering dataset. *arXiv preprint  
563 arXiv:2305.12524*, 2023.

564

565 Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and  
566 Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model  
567 scaling, 2025b. URL <https://arxiv.org/abs/2501.17811>.

568

569 Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuwei Liu,  
570 Mengfei Zhou, Zhuosheng Zhang, et al. Do not think that much for 2+ 3=? on the overthinking  
571 of o1-like llms. *arXiv preprint arXiv:2412.21187*, 2024.

572

573 Yilong Chen, Junyuan Shang, Zhenyu Zhang, Yanxi Xie, Jiawei Sheng, Tingwen Liu, Shuohuan  
574 Wang, Yu Sun, Hua Wu, and Haifeng Wang. Inner thinking transformer: Leveraging dynamic  
575 depth scaling to foster adaptive internal thinking, 2025c. URL <https://arxiv.org/abs/2502.13842>.

576

577 Jeffrey Cheng and Benjamin Van Durme. Compressed chain of thought: Efficient reasoning through  
578 dense representations, 2024. URL <https://arxiv.org/abs/2412.13171>.

579

580 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
581 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to  
582 solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

583

584 Yuntian Deng, Kiran Prasad, Roland Fernandez, Paul Smolensky, Vishrav Chaudhary, and Stu-  
585 art Shieber. Implicit chain of thought reasoning via knowledge distillation. *arXiv preprint  
586 arXiv:2311.01460*, 2023.

587

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

591

592 Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jiang, Bill Yuchen Lin, Peter  
593 West, Chandra Bhagavatula, Ronan Le Bras, Jena D. Hwang, Soumya Sanyal, Sean Welleck, Xi-  
594 ang Ren, Allyson Ettinger, Zaid Harchaoui, and Yejin Choi. Faith and fate: Limits of transformers  
595 on compositionality, 2023. URL <https://arxiv.org/abs/2305.18654>.

596

597 Lijie Fan, Tianhong Li, Siyang Qin, Yuanzhen Li, Chen Sun, Michael Rubinstein, Deqing Sun,  
598 Kaiming He, and Yonglong Tian. Fluid: Scaling autoregressive text-to-image generative models  
599 with continuous tokens, 2024. URL <https://arxiv.org/abs/2410.13863>.

---

594 Kanishk Gandhi, Denise Lee, Gabriel Grand, Muxin Liu, Winson Cheng, Archit Sharma, and  
595 Noah D Goodman. Stream of search (sos): Learning to search in language. *arXiv preprint*  
596 *arXiv:2404.03683*, 2024.

597 Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, Nathan Lile, and Noah D. Goodman. Cognitive  
598 behaviors that enable self-improving reasoners, or, four habits of highly effective stars, 2025.  
599 URL <https://arxiv.org/abs/2503.01307>.

600 Silin Gao, Mete Ismayilzada, Mengjie Zhao, Hiromi Wakaki, Yuki Mitsufuji, and Antoine Bosselut.  
601 Diffucomet: Contextual commonsense knowledge diffusion, 2024. URL <https://arxiv.org/abs/2402.17011>.

602 Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R. Bartoldson,  
603 Bhavya Kailkhura, Abhinav Bhatele, and Tom Goldstein. Scaling up test-time compute with latent  
604 reasoning: A recurrent depth approach, 2025. URL <https://arxiv.org/abs/2502.05171>.

605 Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and LingPeng Kong. Diffuseq: Sequence to  
606 sequence text generation with diffusion models. *arXiv preprint arXiv:2210.08933*, 2022.

607 Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An,  
608 Peilin Zhao, Wei Bi, Jiawei Han, Hao Peng, and Lingpeng Kong. Scaling diffusion language  
609 models via adaptation from autoregressive models, 2025. URL <https://arxiv.org/abs/2410.17891>.

610 Sachin Goyal, Ziwei Ji, Ankit Singh Rawat, Aditya Krishna Menon, Sanjiv Kumar, and Vaishnavh  
611 Nagarajan. Think before you speak: Training language models with pause tokens. *arXiv preprint*  
612 *arXiv:2310.02226*, 2023.

613 Ishaan Gulrajani and Tatsunori B Hashimoto. Likelihood-based diffusion language models. *Advances in Neural Information Processing Systems*, 36:16693–16715, 2023.

614 Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong  
615 Tian. Training large language models to reason in a continuous latent space, 2024. URL <https://arxiv.org/abs/2412.06769>.

616 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu,  
617 Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for  
618 promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint*  
619 *arXiv:2402.14008*, 2024.

620 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,  
621 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv*  
622 *preprint arXiv:2103.03874*, 2021.

623 David Herel and Tomas Mikolov. Thinking tokens for language modeling, 2024. URL <https://arxiv.org/abs/2405.08644>.

624 Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick,  
625 Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a  
626 constrained variational framework. In *International conference on learning representations*, 2017.

627 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint*  
628 *arXiv:2207.12598*, 2022.

629 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*  
630 *neural information processing systems*, 33:6840–6851, 2020.

631 Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song,  
632 and Denny Zhou. Large language models cannot self-correct reasoning yet, 2024. URL <https://arxiv.org/abs/2310.01798>.

633 Daniel Israel, Guy Van den Broeck, and Aditya Grover. Accelerating diffusion llms via adaptive  
634 parallel decoding, 2025. URL <https://arxiv.org/abs/2506.00413>.

---

648 Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec  
649 Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv*  
650 *preprint arXiv:2412.16720*, 2024.

651 Mingyu Jin, Weidi Luo, Sitalo Cheng, Xinyi Wang, Wenyue Hua, Ruixiang Tang, William Yang  
652 Wang, and Yongfeng Zhang. Disentangling memory and reasoning ability in large language  
653 models, 2025. URL <https://arxiv.org/abs/2411.13504>.

654 Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish  
655 Sabharwal. Decomposed prompting: A modular approach for solving complex tasks, 2023. URL  
656 <https://arxiv.org/abs/2210.02406>.

657 Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint*  
658 *arXiv:1312.6114*, 2013.

659 Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-  
660 lm improves controllable text generation. *Advances in neural information processing systems*, 35:  
661 4328–4343, 2022.

662 Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian  
663 Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, et al. From system 1 to system 2: A survey of  
664 reasoning large language models. *arXiv preprint arXiv:2502.17419*, 2025.

665 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching  
666 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

667 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,  
668 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint*  
669 *arXiv:2412.19437*, 2024a.

670 Jie Liu, Gongye Liu, Jiajun Liang, Yangguang Li, Jiaheng Liu, Xintao Wang, Pengfei Wan,  
671 Di Zhang, and Wanli Ouyang. Flow-grpo: Training flow matching models via online rl. *arXiv*  
672 *preprint arXiv:2505.05470*, 2025.

673 Tianqiao Liu, Zui Chen, Zitao Liu, Mi Tian, and Weiqi Luo. Expediting and elevating large language  
674 model reasoning via hidden chain-of-thought decoding, 2024b. URL <https://arxiv.org/abs/2409.08561>.

675 Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion language modeling by estimating  
676 the ratios of the data distribution. *arXiv preprint arXiv:2310.16834*, 2023.

677 Justin Lovelace, Christian K Belardi, Sofian Zalouk, Adhitya Polavarapu, Srivatsa R Kundurthy, and  
678 Kilian Q Weinberger. Stop-think-autoregress: Language modeling with latent diffusion planning.  
679 In *Second Conference on Language Modeling*.

680 Justin Lovelace, Varsha Kishore, Chao Wan, Eliot Shekhtman, and Kilian Q Weinberger. Latent dif-  
681 fusion for language generation. *Advances in Neural Information Processing Systems*, 36:56998–  
682 57025, 2023.

683 Justin Lovelace, Varsha Kishore, Yiwei Chen, and Kilian Q Weinberger. Diffusion guided language  
684 modeling. *arXiv preprint arXiv:2408.04220*, 2024.

685 Viacheslav Meshchaninov, Egor Chimbulatov, Alexander Shabalin, Aleksandr Abramov, and  
686 Dmitry Vetrov. Compressed and smooth latent space for text diffusion modeling. *arXiv preprint*  
687 *arXiv:2506.21170*, 2025.

688 Amirkeivan Mohtashami, Matteo Pagliardini, and Martin Jaggi. CoTFormer: A chain of thought  
689 driven architecture with budget-adaptive computation cost at inference. In *The Thirteenth In-  
690 ternational Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=7igPXQFupX>.

691 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke  
692 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time  
693 scaling. *arXiv preprint arXiv:2501.19393*, 2025.

---

702 Ranjita Naik, Varun Chandrasekaran, Mert Yuksekgonul, Hamid Palangi, and Besmira Nushi. Di-  
703 versity of thought improves reasoning abilities of llms, 2024. URL <https://arxiv.org/abs/2310.07088>.  
704

705 Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin,  
706 Ji-Rong Wen, and Chongxuan Li. Large language diffusion models, 2025. URL <https://arxiv.org/abs/2502.09992>.  
707

708 Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David  
709 Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Au-  
710 gustus Odena. Show your work: Scratchpads for intermediate computation with language models,  
711 2021. URL <https://arxiv.org/abs/2112.00114>.  
712

713 Xichen Pan, Satya Narayan Shukla, Aashu Singh, Zhuokai Zhao, Shlok Kumar Mishra, Jialiang  
714 Wang, Zhiyang Xu, Juhai Chen, Kunpeng Li, Felix Juefei-Xu, Ji Hou, and Saining Xie. Trans-  
715 fer between modalities with metaqueries, 2025. URL <https://arxiv.org/abs/2504.06256>.  
716

717 Jacob Pfau, William Merrill, and Samuel R. Bowman. Let’s think dot by dot: Hidden computation  
718 in transformer language models, 2024. URL <https://arxiv.org/abs/2404.15758>.  
719

720 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
721 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-  
722 ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.  
723

724 Swarnadeep Saha, Xian Li, Marjan Ghazvininejad, Jason Weston, and Tianlu Wang. Learning to  
725 plan & reason for evaluation with thinking-llm-as-a-judge, 2025. URL <https://arxiv.org/abs/2501.18099>.  
726

727 Subham Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin Chiu,  
728 Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language  
729 models. *Advances in Neural Information Processing Systems*, 37:130136–130184, 2024.  
730

731 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv  
732 preprint arXiv:2202.00512*, 2022.  
733

734 Nikunj Saunshi, Nishanth Dikkala, Zhiyuan Li, Sanjiv Kumar, and Sashank J. Reddi. Reasoning  
735 with latent thoughts: On the power of looped transformers, 2025. URL <https://arxiv.org/abs/2502.17416>.  
736

737 David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. Analysing mathematical rea-  
738 soning abilities of neural models. *arXiv preprint arXiv:1904.01557*, 2019.  
739

740 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,  
741 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathe-  
742 matical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300>.  
743

744 Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. Codi: Com-  
745 pressing chain-of-thought into continuous space via self-distillation, 2025. URL <https://arxiv.org/abs/2502.21074>.  
746

747 Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis K. Titsias. Simplified and gener-  
748 alized masked diffusion for discrete data, 2025a. URL <https://arxiv.org/abs/2406.04329>.  
749

750 Weijia Shi, Xiaochuang Han, Chunting Zhou, Weixin Liang, Xi Victoria Lin, Luke Zettlemoyer,  
751 and Lili Yu. Lmfusion: Adapting pretrained language models for multimodal generation, 2025b.  
752 URL <https://arxiv.org/abs/2412.15188>.  
753

754 Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and  
755 Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning, 2023. URL  
<https://arxiv.org/abs/2303.11366>.

756 Mukul Singh, José Cambronero, Sumit Gulwani, Vu Le, Carina Negreanu, and Gust Verbruggen.  
757 Codefusion: A pre-trained diffusion model for code generation. In *Proceedings of the 2023*  
758 *Conference on Empirical Methods in Natural Language Processing*, pp. 11697–11708, 2023.

759

760 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv*  
761 *preprint arXiv:2010.02502*, 2020.

762

763 Yuxuan Song, Zheng Zhang, Cheng Luo, Pengyang Gao, Fan Xia, Hao Luo, Zheng Li, Yuehang  
764 Yang, Hongli Yu, Xingwei Qu, Yuwei Fu, Jing Su, Ge Zhang, Wenhao Huang, Mingxuan Wang,  
765 Lin Yan, Xiaoying Jia, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Yonghui Wu, and Hao Zhou.  
766 Seed diffusion: A large-scale diffusion language model with high-speed inference, 2025. URL  
767 <https://arxiv.org/abs/2508.02193>.

768

769 DiJia Su, Hanlin Zhu, Yingchen Xu, Jiantao Jiao, Yuandong Tian, and Qinqing Zheng. Token  
770 assorted: Mixing latent and text tokens for improved language model reasoning. *arXiv preprint*  
771 *arXiv:2502.03275*, 2025.

772

773 Jihoon Tack, Jack Lanchantin, Jane Yu, Andrew Cohen, Ilia Kulikov, Janice Lan, Shibo Hao, Yuan-  
774 dong Tian, Jason Weston, and Xian Li. Llm pretraining with continuous concepts, 2025. URL  
775 <https://arxiv.org/abs/2502.08524>.

776

777 Haotian Tang, Yecheng Wu, Shang Yang, Enze Xie, Junsong Chen, Junyu Chen, Zhuoyang Zhang,  
778 Han Cai, Yao Lu, and Song Han. Hart: Efficient visual generation with hybrid autoregressive  
779 transformer, 2024a. URL <https://arxiv.org/abs/2410.10812>.

780

781 Zhengyang Tang, Xingxing Zhang, Benyou Wang, and Furu Wei. Mathscale: Scaling instruction  
782 tuning for mathematical reasoning. *arXiv preprint arXiv:2403.02884*, 2024b.

783

784 Shengbang Tong, David Fan, Jiachen Zhu, Yunyang Xiong, Xinlei Chen, Koustuv Sinha, Michael  
785 Rabbat, Yann LeCun, Saining Xie, and Zhuang Liu. Metamorph: Multimodal understanding and  
786 generation via instruction tuning, 2024a. URL <https://arxiv.org/abs/2412.14164>.

787

788 Yuxuan Tong, Xiwen Zhang, Rui Wang, Ruidong Wu, and Junxian He. Dart-math: Difficulty-aware  
789 rejection tuning for mathematical problem-solving. *Advances in Neural Information Processing  
790 Systems*, 37:7821–7846, 2024b.

791

792 Emile van Krieken, Pasquale Minervini, Edoardo Ponti, and Antonio Vergari. Neurosymbolic dif-  
793 fusion models, 2025. URL <https://arxiv.org/abs/2505.13138>.

794

795 Xinyi Wang, Lucas Caccia, Oleksiy Ostapenko, Xingdi Yuan, William Yang Wang, and Alessan-  
796 dro Sordoni. Guiding language model reasoning with planning tokens. In *First Conference on  
797 Language Modeling*, 2024. URL <https://openreview.net/forum?id=wi9IffRhVM>.

798

799 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc  
800 Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models,  
801 2023. URL <https://arxiv.org/abs/2201.11903>.

802

803 Ashen Weligalle. Discrete diffusion models for language generation. *arXiv preprint*  
804 *arXiv:2507.07050*, 2025.

805

806 Chünhung Wu, Jinliang Lu, Zixuan Ren, Gangqiang Hu, Zhi Wu, Dai Dai, and Hua Wu. Llms are  
807 single-threaded reasoners: Demystifying the working mechanism of soft thinking. *arXiv preprint*  
808 *arXiv:2508.03440*, 2025.

809

810 Jianxiang Xiang, Zhenhua Liu, Haodong Liu, Yin Bai, Jia Cheng, and Wenliang Chen. Diffu-  
811 siondialog: A diffusion model for diverse dialog generation with latent space. *arXiv preprint*  
812 *arXiv:2404.06760*, 2024.

813

814 Shitao Xiao, Yueze Wang, Junjie Zhou, Huaying Yuan, Xingrun Xing, Ruiran Yan, Chaofan Li,  
815 Shuting Wang, Tiejun Huang, and Zheng Liu. Omnigen: Unified image generation, 2024. URL  
816 <https://arxiv.org/abs/2409.11340>.

---

810 Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu,  
811 Zhirong Wu, and Chong Luo. Logic-rl: Unleashing llm reasoning with rule-based reinforcement  
812 learning, 2025. URL <https://arxiv.org/abs/2502.14768>.

813

814 Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, Xu Zhao, Min-Yen Kan, Junxian He, and Qizhe Xie.  
815 Self-evaluation guided beam search for reasoning, 2023. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2305.00633)  
816 [2305.00633](https://arxiv.org/abs/2305.00633).

817

818 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik  
819 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models, 2023.  
820 URL <https://arxiv.org/abs/2305.10601>.

821

822 Jiacheng Ye, Jiahui Gao, Shansan Gong, Lin Zheng, Xin Jiang, Zhenguo Li, and Lingpeng Kong.  
823 Beyond autoregression: Discrete diffusion for complex reasoning and planning. *arXiv preprint*  
*arXiv:2410.14157*, 2024a.

824

825 Jiacheng Ye, Shansan Gong, Liheng Chen, Lin Zheng, Jiahui Gao, Han Shi, Chuan Wu, Xin Jiang,  
826 Zhenguo Li, Wei Bi, and Lingpeng Kong. Diffusion of thoughts: Chain-of-thought reasoning in  
827 diffusion language models, 2024b. URL <https://arxiv.org/abs/2402.07754>.

828

829 Jiacheng Ye, Jiahui Gao, Shansan Gong, Lin Zheng, Xin Jiang, Zhenguo Li, and Lingpeng Kong.  
830 Beyond autoregression: Discrete diffusion for complex reasoning and planning, 2025a. URL  
<https://arxiv.org/abs/2410.14157>.

831

832 Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng  
833 Kong. Dream 7b: Diffusion large language models. *arXiv preprint arXiv:2508.15487*, 2025b.

834

835 Jiasheng Ye, Zaixiang Zheng, Yu Bao, Lihua Qian, and Quanquan Gu. Diffusion language mod-  
836 els can perform many tasks with scaling and instruction-finetuning, 2025c. URL <https://arxiv.org/abs/2308.12219>.

837

838 Fangxu Yu, Lai Jiang, Haoqiang Kang, Shibo Hao, and Lianhui Qin. Flow of reasoning: Efficient  
839 training of llm policy with divergent thinking. *arXiv preprint arXiv:2406.05673*, 1(2):6, 2024.

840

841 Fangxu Yu, Lai Jiang, Haoqiang Kang, Shibo Hao, and Lianhui Qin. Flow of reasoning: Training  
842 llms for divergent reasoning with minimal examples, 2025a. URL [https://arxiv.org/](https://arxiv.org/abs/2406.05673)  
[abs/2406.05673](https://arxiv.org/abs/2406.05673).

843

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

847

848 Qifan Yu, Zhenyu He, Sijie Li, Xun Zhou, Jun Zhang, Jingjing Xu, and Di He. Enhancing auto-  
849 regressive chain-of-thought through loop-aligned reasoning, 2025b. URL [https://arxiv.org/](https://arxiv.org/abs/2502.08482)  
[abs/2502.08482](https://arxiv.org/abs/2502.08482).

850

851 Runpeng Yu, Xinyin Ma, and Xinchao Wang. Dimple: Discrete diffusion multimodal large language  
852 model with parallel decoding. *arXiv preprint arXiv:2505.16990*, 2025c.

853

854 Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Shiji Song, and Gao Huang. Does re-  
855 enforcement learning really incentivize reasoning capacity in llms beyond the base model? *arXiv*  
856 *preprint arXiv:2504.13837*, 2025.

857

858 Eric Zelikman, Georges Raif Harik, Yijia Shao, Varuna Jayasiri, Nick Haber, and Noah Good-  
859 man. Quiet-STar: Language models can teach themselves to think before speaking. In *First*  
860 *Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=oRXPiSOGH9>.

861

862 Yiming Zeng, Jinghan Cao, Zexin Li, Yiming Chen, Tao Ren, Dawei Xiang, Xidong Wu, Shangqian  
863 Gao, and Tingting Yu. Treediff: Ast-guided code generation with diffusion llms. *arXiv preprint*  
*arXiv:2508.01473*, 2025.

---

864 Dan Zhang, Sining Zhoubian, Ziniu Hu, Yisong Yue, Yuxiao Dong, and Jie Tang. ReST-MCTS\*:  
865 LLM self-training via process reward guided tree search. In *The Thirty-eighth Annual Confer-*  
866 *ence on Neural Information Processing Systems*, 2024. URL [https://openreview.net/](https://openreview.net/forum?id=8rcFOqEud5)  
867 [forum?id=8rcFOqEud5](https://openreview.net/forum?id=8rcFOqEud5).

868 Yizhe Zhang, Jiatao Gu, Zhuofeng Wu, Shuangfei Zhai, Joshua Susskind, and Navdeep Jaitly. Plan-  
869 ner: Generating diversified paragraph via latent language diffusion model. *Advances in Neural*  
870 *Information Processing Systems*, 36:80178–80190, 2023.

871 Zhen Zhang, Xuehai He, Weixiang Yan, Ao Shen, Chenyang Zhao, Shuohang Wang, Yelong Shen,  
872 and Xin Eric Wang. Soft thinking: Unlocking the reasoning potential of llms in continuous  
873 concept space. *arXiv preprint arXiv:2505.15778*, 2025.

874 Lin Zheng, Jianbo Yuan, Lei Yu, and Lingpeng Kong. A reparameterized discrete diffusion model  
875 for text generation, 2024. URL <https://arxiv.org/abs/2302.05737>.

876 Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Ja-  
877 cob Kahn, Xuezhe Ma, Luke Zettlemoyer, and Omer Levy. Transfusion: Predict the next token  
878 and diffuse images with one multi-modal model, 2024a. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2408.11039)  
879 [2408.11039](https://arxiv.org/abs/2408.11039).

880 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuur-  
881 mans, Claire Cui, Olivier Bousquet, Quoc Le, and Ed Chi. Least-to-most prompting enables com-  
882 plex reasoning in large language models, 2023. URL [https://arxiv.org/abs/2205.](https://arxiv.org/abs/2205.10625)  
883 [10625](https://arxiv.org/abs/2205.10625).

884 Zixuan Zhou, Xuefei Ning, Ke Hong, Tianyu Fu, Jiaming Xu, Shiyao Li, Yuming Lou, Luning  
885 Wang, Zhihang Yuan, Xiuhong Li, Shengen Yan, Guohao Dai, Xiao-Ping Zhang, Yuhan Dong,  
886 and Yu Wang. A survey on efficient inference for large language models, 2024b. URL <https://arxiv.org/abs/2404.14294>.

887 Hanlin Zhu, Shibo Hao, Zhiting Hu, Jiantao Jiao, Stuart Russell, and Yuandong Tian. Reason-  
888 ing by superposition: A theoretical perspective on chain of continuous thought. *arXiv preprint*  
889 *arXiv:2505.12514*, 2025.

890

891

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

---

918    **A ADDITIONAL RELATED WORKS**  
919

920    **Diffusion Language Models for Text Reasoning** Masked diffusion language models attempt to  
921    address some common limitations of autoregressive LLMs—such as rigid left-to-right decoding and  
922    inefficiency—by iteratively denoising masked tokens, enabling parallel and order-agnostic text genera-  
923    tion. Prior studies show that these models achieve better inference efficiency compared to AR  
924    models while maintaining comparable performance on both general tasks (Zheng et al., 2024; Gong  
925    et al., 2025; Nie et al., 2025; Shi et al., 2025a; Song et al., 2025; Ye et al., 2025c;b) and reasoning  
926    benchmarks with chain-of-thought (Gao et al., 2024; Ye et al., 2024b; van Krieken et al., 2025; Ye  
927    et al., 2025a). However, these approaches remain constrained to language space, unable to capture  
928    reasoning at an abstract semantic level or revise previously generated tokens as continuous diffu-  
929    sion models (Ho et al., 2020; Song et al., 2020) do, and they require training on massive datasets  
930    rather than leveraging a well-trained LLM. In contrast, our method overcomes these limitations by  
931    structuring reasoning in an interpretable continuous latent space, producing abstract CoT representa-  
932    tions with self-correction ability for an existing LLM, while keeping diffusion’s strengths in parallel  
933    generation to enhance exploration and diversity.

933    **CoT Reasoning** Chain-of-thought reasoning refers to methods which elicit LLMs to generate in-  
934    termediate reasoning steps in language prior to outputting a final answer in order to improve perfor-  
935    mance on reasoning tasks. This can be accomplished via prompting methods (Nye et al., 2021; Wei  
936    et al., 2023; Khot et al., 2023; Zhou et al., 2023) or through training LLMs (by SFT, RL, or a combi-  
937    nation of the two) to output the intermediate reasoning steps (Yu et al., 2023; Shao et al., 2024; Yu  
938    et al., 2025a). Works have also extended CoT to allow LLMs to mimic various tree search algorithms  
939    such as BFS or MCTS, which especially improves performance on more complex tasks (Xie et al.,  
940    2023; Yao et al., 2023; Zhang et al., 2024; Bi et al., 2025). Beyond following specific algorithms,  
941    works that implement long chain-of-thought (combining extensive reasoning, exploration, and re-  
942    flection) have also demonstrated improved reasoning performance (Shinn et al., 2023; Gandhi et al.,  
943    2025; Saha et al., 2025; Xie et al., 2025). One overarching limiting factor with these CoT methods is  
944    that they fundamentally work at a next-token-prediction level, constraining the outputs to the token  
945    space and limiting the model’s horizon.

946    **Hybrid AR+Diffusion Model Architecture** Other AR-Diffusion hybrid models have shown suc-  
947    cessful results in rivaling their AR and diffusion counterparts, particularly in multimodal generation  
948    and image understanding. The Transfusion (Zhou et al., 2024a) architecture demonstrated that  
949    hybrid models could outperform standard AR models and compete with state-of-the-art diffusion  
950    models in image-generation benchmarks, a phenomenon further reinforced by other studies of hy-  
951    brid models (Fan et al., 2024; Tang et al., 2024a; Xiao et al., 2024). This extends beyond image  
952    generation, with several works demonstrating the effectiveness of hybrid AR-diffusion models in  
953    other domains such as image understanding, video generation, and robot control (Black et al., 2024;  
954    Tong et al., 2024a; Chen et al., 2025a;b). Furthermore—similar to our model architecture—works have  
955    demonstrated successful adaptations of frozen models for these hybrid AR-diffusion architectures in  
956    multimodal domains (Pan et al., 2025; Shi et al., 2025b). Aside from the difference in domain from  
957    these works, many do not use block diffusion for variable-length generations as in LaDiR and we  
958    critically introduce CE loss to guide better latent predictions.

959    **B ADDITIONAL PRELIMINARIES AND BACKGROUND**

960    We provide more details about the background information of VAE and Diffusion models in this  
961    section.

962    **B.1 VARIATIONAL AUTOENCODER AND  $\beta$ -VAE**

963    The Variational Autoencoder (VAE) (Kingma & Welling, 2013) is a latent-variable model that learns  
964    a compressed representation of data  $x$  through an encoder–decoder pair. The encoder  $q_\phi(z|x)$  maps  
965    inputs into a distribution over latent variables  $z$ , typically parameterized as a diagonal Gaussian. The  
966    decoder  $p_\theta(x|z)$  reconstructs the input from  $z$ , enabling generative sampling. Training maximizes  
967    the evidence lower bound (ELBO):

$$\mathcal{L}_{\text{VAE}} = \mathbb{E}_{q_\phi(z|x)}[-\log p_\theta(x|z)] + \text{KL}(q_\phi(z|x) \parallel p(z)), \quad (7)$$

968    where the first term ensures faithful reconstruction and the second term regularizes the posterior  
969    toward a simple prior  $p(z)$ , usually  $\mathcal{N}(0, I)$ . The reparameterization trick,  
970

$$z = \mu_\phi(x) + \sigma_\phi(x) \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I),$$

enables low-variance gradient estimates for stochastic optimization.

**$\beta$ -VAE.** The  $\beta$ -VAE (Higgins et al., 2017) introduces a hyperparameter  $\beta$  to control the KL weight:

$$\mathcal{L}_{\beta\text{-VAE}} = \mathbb{E}_{q_\phi(z|x)}[-\log p_\theta(x|z)] + \beta \text{KL}(q_\phi(z|x) \parallel p(z)). \quad (8)$$

When  $\beta > 1$ , the model enforces stronger alignment to the prior, which encourages disentangled and interpretable latent variables at the expense of reconstruction fidelity. This property is desirable when latent codes are later used as the substrate for generative modeling.

**Why VAE for Latent Diffusion.** Latent diffusion models (LDMs) (Rombach et al., 2022) operate not on raw high-dimensional inputs (e.g., images or sequences), but in a compressed latent space learned by a VAE. This design provides three key advantages:

1. **Efficiency.** Operating in latent space reduces dimensionality, leading to faster training and inference while maintaining semantic richness.
2. **Semantic abstraction.** The VAE learns to discard imperceptible details and retain high-level structure, making diffusion steps focus on meaningful features rather than pixel-level noise.
3. **Flexibility.** The decoder  $p_\theta(x|z)$  ensures that even when denoising occurs in latent space, the final output remains in the original input domain. This separation enables diffusion to generalize across modalities with a shared latent backbone.

## B.2 LATENT DIFFUSION: TRAINING AND INFERENCE

Latent diffusion operates in the compressed latent space  $z_0$  of a pretrained VAE.

**Forward process.** Noise is added gradually:

$$q(z_t|z_0) = \mathcal{N}(z_t; \sqrt{\bar{\alpha}_t} z_0, (1 - \bar{\alpha}_t)I),$$

with  $\bar{\alpha}_t = \prod_{s=1}^t (1 - \beta_s)$ .

**Training objective.** The denoiser  $\epsilon_\theta(z_t, t)$  predicts the injected noise:

$$\mathcal{L}_{\text{LDM}} = \mathbb{E}_{z_0, \epsilon, t} [\|\epsilon - \epsilon_\theta(z_t, t)\|^2].$$

**Inference.** Generation starts from  $z_T \sim \mathcal{N}(0, I)$  and denoises iteratively:

$$z_{t-1} = \frac{1}{\sqrt{\bar{\alpha}_t}} \left( z_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(z_t, t) \right) + \sigma_t \epsilon.$$

## B.3 COMPARISON OF PARAMETERIZATIONS

Diffusion training can be expressed through different target parameterizations, all of which can be interpreted as variants of the same continuous-time flow. Below we summarize the most common forms:

### B.3.1 $\epsilon$ -PREDICTION

The denoiser directly predicts the added Gaussian noise  $\epsilon$ :

$$\mathcal{L}_\epsilon = \mathbb{E}_{z_0, \epsilon, t} [\|\epsilon - \epsilon_\theta(z_t, t)\|^2], \quad (9)$$

where  $z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ . This is the standard DDPM formulation (Ho et al., 2020). It is stable but sometimes less efficient for long horizons.

### B.3.2 $x_0$ -PREDICTION (DDIM- $x_0$ )

Instead of noise, the model predicts the clean latent  $z_0$ :

$$\mathcal{L}_{x_0} = \mathbb{E}_{z_0, t} [\|z_0 - x_{0,\theta}(z_t, t)\|^2]. \quad (10)$$

This corresponds to the DDIM formulation (Song et al., 2020), enabling deterministic sampling and fewer inference steps, but can overfit to data scale.

1026 B.3.3  $v$ -PREDICTION  
 1027 Proposed by Salimans & Ho (2022),  $v$  is defined as a linear combination of noise and clean latent:  
 1028  
 1029  $v = \sqrt{\bar{\alpha}_t} \epsilon - \sqrt{1 - \bar{\alpha}_t} z_0,$  (11)

1030 with objective

$$1031 \mathcal{L}_v = \mathbb{E}_{z_0, \epsilon, t} [\|v - v_\theta(z_t, t)\|^2]. \quad (12)$$

1032  $v$ -prediction is numerically better conditioned, often improving stability across timesteps.  
 1033

1034 All four parameterizations can be viewed as different instantiations of the same underlying generative  
 1035 flow.  $\epsilon$ -prediction,  $x_0$ -prediction, and  $v$ -prediction specify *which quantity* the denoiser regresses  
 1036 on. Flow matching directly learns the continuous velocity field, avoiding discretization artifacts.

#### 1037 B.4 BLOCK DIFFUSION

1038 Suppose we are using  $\epsilon$ -prediction, and let a sequence be segmented into  $M$  blocks  $\{B_1, \dots, B_M\}$ ,  
 1039 where  $B_m \in \mathbb{R}^{k \times d}$  contains  $k$  latent tokens of dimension  $d$ . The forward noising process for block  
 1040  $B_m$  is

$$1041 q(B_{m,t} | B_{m,0}) = \mathcal{N}(\sqrt{\bar{\alpha}_t} B_{m,0}, (1 - \bar{\alpha}_t) I), \quad (13)$$

1042 and the denoiser  $f_\theta$  is trained to predict the noise at the block level:

$$1043 \mathcal{L}_{\text{block}} = \mathbb{E}_{m,t,\epsilon} [\|\epsilon - f_\theta(B_{m,t}, t)\|^2]. \quad (14)$$

1044 Blocks are generated autoregressively, i.e.,

$$1045 p(B_m | B_{<m}) = \int q(B_{m,0} | x) \prod_t p_\theta(B_{m,t-1} | B_{m,t}, B_{<m}) dB_{m,0}, \quad (15)$$

1046 so that each block is denoised iteratively while conditioning on all previously generated blocks.  
 1047

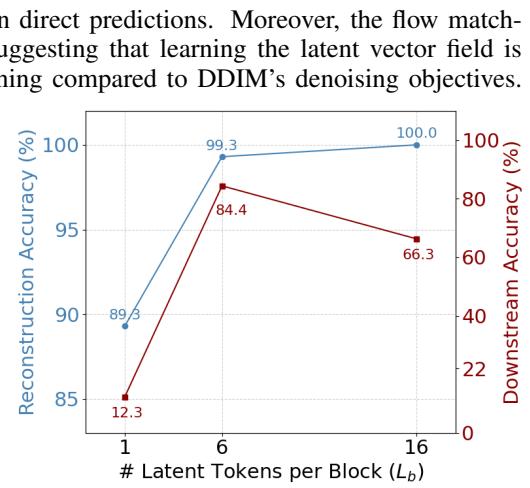
## 1048 C ADDITIONAL ABLATION STUDIES

1049 **Latent Prediction Objective** To assess the impact of different training objectives for latent prediction,  
 1050 we compare several widely used formulations.  
 1051 The first baseline is *MSE loss*, which directly minimizes the mean-squared error between predicted  
 1052 and ground-truth latents but yields the weakest results.  
 1053 We then adopt three DDIM-based (Song et al.,  
 1054 2020) objectives: predicting the clean latent state  
 1055 ( $x_0$ ), the added noise vector ( $\epsilon$ ), and the velocity ( $v$ ).  
 1056 These diffusion objectives consistently improve accuracy,  
 1057 highlighting the advantage of explicitly modeling the denoising process rather than relying on direct predictions.  
 1058 Moreover, the flow matching ( $u$ ) objective achieves the strongest gains, suggesting that learning the latent vector field is  
 1059 more effective for capturing the pattern of reasoning compared to DDIM’s denoising objectives.  
 1060

1061 **Effect of the Block Size.** We investigate how  
 1062 the number of latent tokens per block ( $L_b$ ) influences  
 1063 reconstruction quality and downstream reasoning  
 1064 accuracy on GSM8K. As shown in Figure 6, too few tokens (i.e., 1 token) limit  
 1065 the model’s ability to capture necessary information,  
 1066 harming reconstruction. Performance improves as the number of tokens increases, reaching  
 1067 near-perfect reconstruction at  $n = 6$ . Beyond  
 1068 this point, however, adding more tokens introduces  
 1069 redundancy, which makes the latent diffusion  
 1070 model harder to predict accurately and leads  
 1071 to diminished reasoning accuracy. This reveals a  
 1072 trade-off between compact latent representations  
 1073 and effective downstream reasoning.  
 1074

Objective	CD-4 Pass@1 (%)
MSE Loss	46.0
$x_0$	53.0
$\epsilon$	58.0
$v$	62.0
<b><math>u</math> (ours)</b>	<b>73.5</b>

1075 Table 5: Ablation study on latent prediction  
 1076 objectives on the Countdown-4 dataset.



1077 Figure 6: Ablation analysis of block size on the  
 1078 GSM8K benchmark.

---

1080  
1081 **VAE Robustness Augmentations.** When training our VAE, to improve the robustness of the VAE  
1082 latent space, we apply two augmentation strategies during VAE training: (1) adding Gaussian noise  
1083 to latent representations with standard deviation  $k$ , and (2) randomly substituting input tokens with  
1084 probability  $p$ . Table 6 shows the impact of these augmentations on GSM8K accuracy. Best per-  
1085 formance is achieved at  $k = 3$  and  $p = 0.3$ . Too little augmentation ( $k = 0$  or  $p = 0$ ) results in  
1086 overfitting to clean inputs, while excessive augmentation ( $k = 5$  or  $p > 0.5$ ) degrades the latent  
1087 space quality.

Latent Gaussian Noise (p=0.3)		Token Substitution (k=3)	
$k$ (std)	GSM8K Acc (%)	$p$ (prob.)	GSM8K Acc (%)
0	68.3	0.0	70.2
1	73.4	0.1	78.3
<b>3</b>	<b>84.2</b>	<b>0.3</b>	<b>84.2</b>
5	79.4	0.5	64.0
–	–	0.7	32.4

1095  
1096 Table 6: Ablation study on VAE robustness augmentations on GSM8K.  
1097  
1098

1099 **Blockization Strategy** We investigate the sensitivity of the model to different blockization strate-  
1100 gies by varying the number of sentences per block. Table 7 shows results for 1, 2, and 3 sentences per  
1101 block. Using more sentences per block requires more latent tokens to maintain reconstruction qual-  
1102 ity and significantly increases difficulty for the diffusion model. We find that **1 sentence per block**  
1103 with 4 latent tokens offers the best balance between latent compactness and reasoning accuracy.

# Sentences	# Latent Tokens	GSM8K	MATH
<b>1</b>	4	<b>84.2</b>	<b>45.2</b>
2	8	78.4	39.6
3	12	72.0	36.1

1109  
1110 Table 7: Ablation study on blockization strategy (sentences per block).  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119  
1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133

Method	MATH	GSM8K	Gaokao	DM	College	Olympia	TheoremQA	Avg.
DIFFUSION LANGUAGE MODELS — LLADA 8B								
Base Model	—	—	—	—	—	—	—	—
CoT SFT	59.9	92.4	43.0	50.7	58.3	10.2	26.3	48.7
AUTOREGRESSIVE MODELS — LLAMA 3.1 8B								
Sol-Only SFT	—	—	—	—	—	—	—	—
CoT SFT	49.8	89.0	37.9	52.0	54.9	12.9	25.0	45.9
iCoT	—	—	—	—	—	—	—	—
Pause Token	—	—	—	—	—	—	—	—
Coconut	39.3	74.3	29.3	36.9	42.9	6.3	14.9	34.8
Discrete Latent	47.3	88.6	39.7	49.5	53.7	17.8	28.5	46.4
<b>LaDiR</b>	<b>63.7</b>	<b>93.7</b>	<b>45.8</b>	<b>54.2</b>	<b>60.3</b>	<b>15.3</b>	<b>30.7</b>	<b>52.0</b>

Table 8: Pass@100 accuracy across in-domain and out-of-domain math benchmarks.

## D ADDITIONAL RESULTS AND ANALYSIS

### Pass@k Results on Math Reasoning

**Interpretability** In addition to achieving superior or competitive accuracy across each benchmark, as shown in Table 9, LaDiR also benefits from being more interpretable by nature compared to standard diffusion-based methods.

Block	Text
Question	<i>Billy sells DVDs. He has 8 customers on Tuesday. The first 3 buy one DVD each, the next 2 buy two DVDs each, the last 3 buy none. How many DVDs did Billy sell?</i>
Decode( $Z^{(1)}$ ):	Billy’s first 3 customers buy one DVD each, so that’s $3 * 1 = \langle\langle 3 * 1 = 3 \rangle\rangle$ 3 DVDs.
Decode( $Z^{(2)}$ ):	His next 2 customers buy 2 DVDs each, so that’s $2 * 2 = 4$ DVDs.
Decode( $Z^{(3)}$ ):	His last 3 customers don’t buy any DVDs, so that’s 0 DVDs sold.
Decode( $Z^{(4)}$ ):	Therefore, Billy sold a total of $3 + 4 + 0 = 7$ DVDs on Tuesday.
Answer	The answer is: 7.

Table 9: Example of interpretable continuous thought tokens: each latent block  $Z^i$  is able to be decoded to human-readable text through the VAE decoder. Each latent block is decoded individually, so the entire latent thought is represented by the block in isolation. This allows for clear interpretability of each latent thought, while still allowing for a model to reason in a latent space.

**Reasoning at Semantic Level** Table 10 demonstrates that LaDiR refines its reasoning through semantic information rather than lexical connections. The **pink segments** trace how the model incrementally refines its intermediate reasoning steps—transforming rough or incomplete arithmetic expressions into logically consistent equations. For example, the model first produces an approximate relation such as  $2 = 1 + 1$  and later reconstructs the correct multiplicative form  $2 * 1 = 2$ , eventually combining these into a complete causal chain that leads to  $2 * 4 = 8$ . Meanwhile, the **green segments** serve as stable anchors capturing the essential reasoning cues—phrases like “two times older than” and “four times as old”—which define the causal dependencies between Suzy, Ben, and Brandon. In contrast, the **yellow segments** correspond to non-reasoning or noisy parts (e.g., “Alice,” “Instagram”) that the model learns to “ignore” during early denoising, prioritizing semantically coherent and logical grounded reasoning over irrelevant lexical content. In contrast, the SFT baseline remains syntactically fluent but semantically brittle—producing text that appears correct on the surface yet fails to capture the true causal logic. These observations indicate that LaDiR’s denoising process operates on the level of semantics and relational consistency, enabling reasoning that aligns with human-like conceptual understanding.

**Consistency Across Long Reasoning Blocks** As shown in Table 11, LaDiR maintains coherent and semantically consistent reasoning across a long sequence of blocks. Each block builds logically on the previous one—starting from variable definition ( $x$  as *silver coins*), to constraint formulation ( $x + (x + 30) = 110$ ), simplification, and finally the correct numerical conclusion ( $x = 40$ , *gold*

---

1188	<b>Question</b>	Brandon's iPhone is four times as old as Ben's iPhone. Ben's iPhone is two times older than Suzy's iPhone. If Suzy's iPhone is 1 year old, how old is Brandon's iPhone?
1189	<b>SFT (Baseline)</b>	Suzy's iPhone is 1 year old. Ben's iPhone would be $1 + 1 = 2$ years old. Since Brandon's iPhone is four times as old as Ben's, we compute $4 \times 1 = 4$ . Therefore, Brandon's iPhone is 4 years old. The answer is: 4.
1190	<b>LaDiR (Ours)</b>	<b>Block 1</b>
1191	Decode( $\hat{Z}_{t=1.0}$ )	10 and 10 10....
1192	Decode( $\hat{Z}_{t=0.9}$ )	If Alice is 1 year old, then Bob's age is 2 years older than Alice, which means Bob is $2 = 1 + 1$ years old.
1193	Decode( $\hat{Z}_{t=0.8}$ )	If Suzy's age is 1 year old, then Ben's age is two times older than Suzy's age, which is $2 * 1 = 2$ years old.
1194	Decode( $\hat{Z}_{t \leq 0.7}$ )	If Suzy's iPhone is 1 year old, then Ben's iPhone is two times older than Suzy's iPhone, so Ben's iPhone is $2 * 1 = 2$ years old.
1195	<b>Block 2</b>	[garbled tokens] & ^ * - // @natechandra's Instagram is four times as old as Ben's Instagram, so Ben's Instagram is $1 * 4 = 4$ years old.
1196		If Brandon's phone is four times as old as Ben's phone, then Brandon's phone is $2 * 4 = 8$ years old. Innbsp;nnbsp;nnbsp...
1197		If Brandon's iPhone is four times as old as Ben's iPhone, then Brandon's iPhone is $2 * 4 = 8$ years old.
1198	<b>Answer</b>	The answer is: 8.

---

1210 Table 10: An example of self-refinement during inference on the GSM8k dataset, showing how  
1211 reasoning becomes progressively clearer as  $t$  decreases. Later denoising steps correct arithmetic  
1212 errors while maintaining earlier structure, demonstrating semantic self-refinement. Pink segments  
1213 highlight refined reasoning portions, yellow segments indicate non-reasoning or noisy parts that the  
1214 model gradually corrects, and green segments denote key reasoning cues essential for correct logic.

1215  
1216 = 70). Unlike the SFT baseline, which produces a single-step approximation that conflates intermediate  
1217 relations, LaDiR preserves arithmetic and causal consistency throughout the reasoning chain,  
1218 demonstrating stable multi-step inference even with a large number of reasoning blocks.

---

1221	<b>Question</b>	Gretchen has 110 coins. There are 30 more gold coins than silver coins. How many gold coins does Gretchen have?
1222	<b>SFT (Baseline)</b>	Let's assume Gretchen has 110 coins in total and 30 more gold than silver. Half of the coins plus 30 should be gold, so $110/2 + 30 = 85$ gold coins.
1223	<b>LaDiR (Ours)</b>	
1224	Decode( $Z^{(1)}$ )	Let's assume the number of silver coins Gretchen has is $x$ silver coins.
1225	Decode( $Z^{(2)}$ )	We also know that there are 30 more gold coins than silver coins, so the number of silver coins is $x + 30$ gold coins.
1226	Decode( $Z^{(3)}$ )	The total number of coins Gretchen has is $x + (x + 30) = 110$ .
1227	Decode( $Z^{(4)}$ )	Combining like terms, we get $2x + 30 = 110$ .
1228	Decode( $Z^{(5)}$ )	Subtracting 30 from both sides, we get $2x = 80$ .
1229	Decode( $Z^{(6)}$ )	Dividing both sides by 2, we get $x = 40$ .
1230	Decode( $Z^{(7)}$ )	Therefore, Gretchen has $30 + 40 = 70$ gold coins.
1231	Answer	The answer is: 70.

---

1237 Table 11: An qualitative example in GSM8K illustrating the long reasoning blocks generated by our  
1238 method compared to the baseline SFT.

1239  
1240  
1241

---

## 1242 E EXPERIMENTAL DETAILS

### 1244 E.1 MATH REASONING

1245 **Implementation Details** The CoT data is segmented into thought-level blocks, where each block  
 1246 corresponds to a single sentence and is represented by 6 latent thought tokens. For VAE training,  
 1247 we set the  $\beta$  to  $10^{-5}$ , in which the encoder is finetuned from the backbone model while the  
 1248 decoder remains frozen. The flow-matching model is trained with the objective in Eq. 5, using  
 1249  $\lambda_{FM} = 5$ ,  $\lambda_{Ans} = 1$ ,  $\lambda_{Spec} = 2$ . During inference, we initialize the Gaussian noise of scale 2, and  
 1250 apply diversity guidance with a maximum scale of 0.8.

1251 **Datasets** We train on only the DART-MATH dataset, holding out all other benchmarks for evalua-  
 1252 tion only. Table 12 summarizes the datasets. While training is limited to mathematical reasoning  
 1253 problems, our evaluations also include out-of-domain tasks such as engineering and physics, pro-  
 1254 viding both in-domain and out-of-domain benchmarks to assess the reasoning and generalization  
 1255 capabilities of LaDiR and the baselines.

1258 Dataset	1259 # Samples	1259 Domain / Level	1259 Notes
<b>1260 DART-MATH</b>	585k	Mixed math (train)	Synthesized for reasoning, based on GSM8K/MATH
<b>1261 MATH</b>	500	High school / competition	In-domain benchmark
<b>1262 GSM8K</b>	1.3k	Grade school arithmetic	In-domain benchmark
<b>1263 College-Math</b>	2.8k	College-level	Linear algebra, differential equations, etc.
<b>1264 DM-Math</b>	1k	K-12 curriculum	Out-of-domain generalization
<b>1265 OlympiaBench-Math</b>	675	Olympiad-level	Advanced competition problems
<b>TheoremQA</b>	800	STEM / theorem-driven	Math, physics, engineering
<b>Fresh-Gaokao-Math-2023</b>	30	Gaokao exam	Real-world test distribution

1266 Table 12: Summary of datasets used in our experiments. We use DART-MATH (Tong et al., 2024b),  
 1267 MATH (Hendrycks et al., 2021), GSM8K (Cobbe et al., 2021), College-Math (Tang et al., 2024b),  
 1268 DeepMind-Math (Saxton et al., 2019), OlympiaBench-Math (He et al., 2024), TheoremQA (Chen  
 1269 et al., 2023), and Fresh-Gaokao-Math-2023 (Tang et al., 2024b).

### 1272 E.2 HYPERPARAMETERS

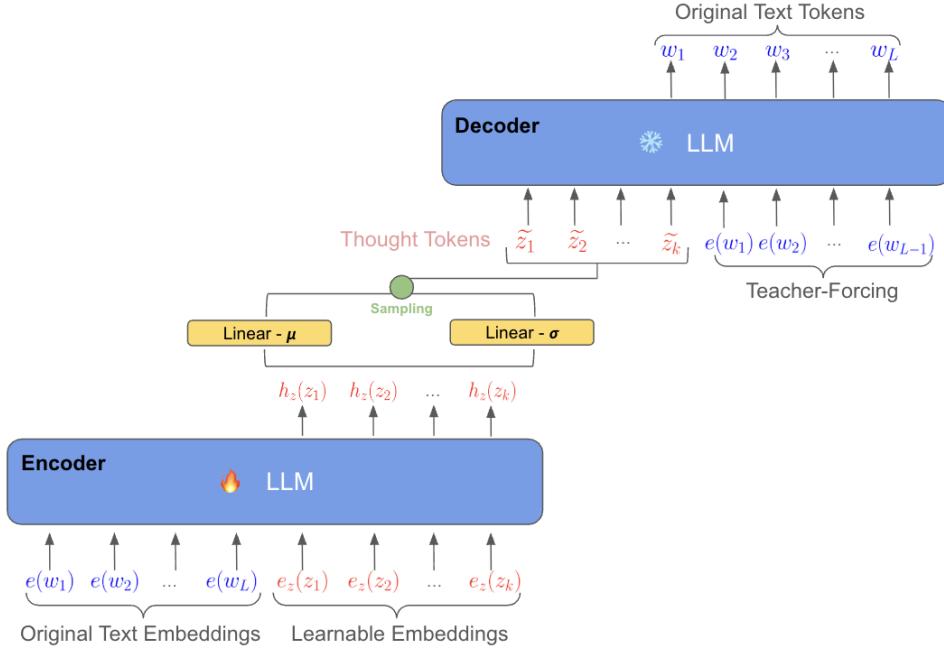
1273 [Table 13 provides a complete summary of all hyperparameters used in our experiments, covering](#)  
 1274 [VAE pretraining, Stage-1 teacher-forcing training, Stage-2 rollout training, and inference settings.](#)

1277 Component	1277 Hyperparameter	1277 Value
<b>1278 VAE Pretraining</b>	Latent dimension ( $d_z$ )	512
	# latent tokens per block	4
	KL weight $\beta$	$1 \times 10^{-5}$
	Learning rate	$2 \times 10^{-5}$
	Batch size	128
	# of Epochs	2
<b>1284 Stage-1 Teacher-Forcing</b>	Flow-matching loss weight ( $\lambda_{FM}$ )	5
	CE loss weight ( $\lambda_{Ans}$ )	1
	Special-token loss weight ( $\lambda_{Spec}$ )	1
	Learning rate	$1 \times 10^{-5}$
	Batch size	64
	# of Epochs	20
<b>1289 Stage-2 Rollout Training</b>	Learning rate	$1 \times 10^{-5}$
	Batch size	12
	# of Epochs	20
<b>1293 Inference</b>	Classifier-free guidance scale	4
	Answer token decoding temperature	0.7

1294 Table 13: Complete hyperparameter settings for all training and inference stages.

---

1296 **F ADDITIONAL MODEL DETAILS**
1297

1298 **VAE Training Architecture** Figure 7 is a more in-depth diagram of the training of the VAE. The  
1299 encoder LLM first maps the input sequence of original text embeddings and learnable embeddings  
1300 into hidden states. These hidden states are then projected through two linear layers to produce  
1301 the mean and variance of the latent distribution, from which we sample the thought tokens via  
1302 the reparameterization trick. The sampled latent tokens  $\tilde{z}_1, \tilde{z}_2, \dots, \tilde{z}_k$  are passed to the decoder  
1303 LLM, which reconstructs the original text tokens under a teacher-forcing setup. This design enables  
1304 the model to compress high-dimensional text into a smaller set of semantically meaningful latent  
1305 variables, while still maintaining faithful reconstruction of the original reasoning process.

1327 Figure 7: Detailed architecture of the variational autoencoder for latent reasoning. The encoder is  
1328 a finetuned LLM that takes both original text embeddings  $e(w_i)$  and learnable embeddings  $e_z(z_i)$ ,  
1329 producing mean and variance vectors through linear projections of the last hidden state  $h$ . Latent  
1330 thought tokens  $\tilde{z}_i$  are then sampled from  $\mathcal{N}(\mu, \sigma^2)$ . The decoder is a frozen LLM that reconstruc-  
1331 the original CoT text under teacher forcing, conditioned on both the sampled thought tokens and the  
1332 original text embeddings.

1333  
1334  
1335  
1336  
1337  
1338  
1339  
1340  
1341  
1342  
1343  
1344  
1345  
1346  
1347  
1348  
1349

---

1350  
1351

## G USE OF LARGE LANGUAGE MODELS (LLMs)

1352  
1353  
1354  
1355  
1356  
1357  
1358  
1359  
1360  
1361  
1362  
1363  
1364  
1365  
1366  
1367  
1368  
1369  
1370  
1371  
1372  
1373  
1374  
1375  
1376  
1377  
1378  
1379  
1380  
1381  
1382  
1383  
1384  
1385  
1386  
1387  
1388  
1389  
1390  
1391  
1392  
1393  
1394  
1395  
1396  
1397  
1398  
1399  
1400  
1401  
1402  
1403

Throughout this project, Large Language Models were utilized as coding tools and as grammar checkers to support the writing of the paper. They did **not** play a significant role in research ideation or writing to the extent of being listed as a contributor. LLMs were used strictly as a general purpose tool.