
Latent Reasoning in Continuous Space for Unified Multimodal Models

Byungwoo Jeon¹, Yoonwoo Jeong², Hyunseok Lee¹, Minsu Cho^{2,3*}, Jinwoo Shin^{1,3*}
¹KAIST, ²POSTECH, ³RLWRLD
imbw2024@kaist.ac.kr

Abstract

Unified Multimodal Models (UMMs) achieve remarkable success on diverse multimodal understanding and generation tasks, yet they still struggle to effectively leverage language-based chain-of-thought reasoning for image generation. As a result, existing reasoning-augmented UMMs often fail to inherit the test-time scaling and fine-grained controllability observed in large language models. We attribute this failure to a fundamental *modality mismatch*: reasoning is performed in a discrete token space, whereas image generation operates in a continuous space. To bridge this gap, we introduce a **Latent Reasoning** framework in **Continuous** space (LARC), which alleviates the token-level bottleneck by reasoning over continuous hidden states. Specifically, the training framework consists of two stages: (i) curriculum supervised fine-tuning (SFT), and (ii) information-gain reinforcement learning (RL). First, the curriculum SFT stage gradually converts token-level reasoning into latent reasoning. Subsequently, the information-gain RL stage self-evolves latent reasoning without ground truth supervision. Experiments across diverse text-to-image generation benchmarks show that LARC consistently improves generation quality, prompt following, and compositional fidelity. Notably, LARC achieves state-of-the-art performance on GenEval with a 7%p gain over BAGEL and demonstrates clear test-time scaling behavior, outperforming token-based generation under both parallel and sequential test-time scaling strategies.

1 Introduction

Recent advances in generative models have reshaped visual content creation, achieving high quality and fine-grained controllability in text-to-image synthesis, image editing, and instruction-following generation [1, 2, 3, 4, 5]. However, generation-only models still struggle with complex semantics that require language-level planning, such as attribute binding [6, 7, 8] or world-modeling [9, 10]. Within this landscape, unified multimodal models (UMMs) [11, 12, 13, 14, 15], have emerged as a compelling paradigm which integrates an image-generation module into large language models (LLMs) that contain rich textual semantics.

Despite their success, prior UMMs fail to scale at test time: unlike in LLMs, increasing the reasoning budget does not yield the smooth performance gains observed on LLM benchmarks (Figure 1). We argue that this failure stems from a fundamental *modality mismatch*: reasoning is performed in a *discrete* token space, whereas image generation operates in a *continuous* latent space. Since discrete chains struggle to encode the dense and spatially-grounded information that image synthesis demands, reasoning collapses into surface-level prompt manipulation, which is precisely the narrow regime in which existing methods provide gains. This observation motivates our central claim: to extend the benefits of reasoning *beyond* this regime, reasoning itself must operate in a continuous space.

*Equal advising




Figure 1: **Test-time scaling.** We investigate the effect of the number of latent during test-time. Under controlled experiments, LARC performs remarkable **test-time scaling behavior**, outperforming the best-of-N sampling on image generation [22], image editing [23], and visual reasoning [24] tasks.

To this end, we introduce LARC (**Latent Reasoning in Continuous space**), a framework that mitigates the *modality gap* by aligning text and image reasoning in a continuous latent space. Building on Coconut [16], which performs latent reasoning in LLMs without explicit tokens, we extend the idea to multimodal generation. While we introduce reinforcement learning (RL) on the latent reasoning system, RL training faces two challenges: (i) Coconut-style reasoning is deterministic, whereas policy-gradient methods require stochastic rollouts; and (ii) latent reasoning traces lack an explicit evaluator, making it difficult to assess whether they are useful for generation. We address both challenges with a two-stage recipe. Stage 1 employs curriculum supervised fine-tuning that progressively replaces token-level reasoning with continuous hidden states, thereby injecting controlled stochasticity and restoring the rollout diversity required for policy gradients. Then, Stage 2 applies reinforcement learning driven by an *information-gain reward*, which quantifies how much each latent step improves the predictive distribution over the target image, together with a *collapse-prevention reward* that preserves trajectory diversity. Together, these constitute a self-evolving loop that requires neither ground-truth chains nor external evaluators.

In experiments, we instantiate LARC on BAGEL [15] with 1M interleaved samples drawn from FLUX-Reason [17], WEAVE [18], and ScaleEdit [19]. On three text-to-image benchmarks—GenEval [20], DPG-Bench [21], and OneIG-Bench [22]—LARC surpasses both the base UMMs and recent reasoning-augmented variants by substantial margins, *e.g.*, +7.6 absolute points on OneIG-Bench over Uni-CoT [13]. More importantly, to our knowledge, LARC is the first reasoning-augmented UMM that exhibits clear test-time scaling: under budget forcing, doubling the latent reasoning compute yields monotonic gains across all three benchmarks.

Our contributions are summarized as follows:

- We identify a discrete–continuous *modality mismatch* between reasoning and generation as the root cause of the limited gains and absent test-time scaling in current reasoning-augmented UMMs, and argue that reasoning must operate in a continuous space.
- We propose LARC, the first framework to perform reasoning in the same continuous latent space as image generation within a unified multimodal model, enabling test-time scaling in UMMs.
- We design a self-evolving training recipe that (i) extends Coconut-style latent reasoning with controlled stochasticity to make it amenable to RL, and (ii) introduces an *information-gain* reward coupled with a *collapse-prevention* reward, eliminating the need for ground-truth latent chains or an external critic.

2 LARC: Latent Reasoning in Continuous Space

2.1 Overview

We propose LARC, a framework that enables UMMs to perform image generation through latent reasoning on continuous space. We first describe the two-stage training pipeline of LARC. We then elaborate on the design of each component in the following sections. Specifically, in Section 2.2, we first define the interleaved trajectory format. Then, in Section 2.3, we introduce the curriculum




Figure 2: **Overview of LARC.** Our framework, LARC, consists of two stages. (1) **Stage 1** adapts supervised fine-tuning (SFT) while gradually introducing latent thought on continuous space where the last hidden state of UMM becomes the next input embedding of the model. (2) **Stage 2** adapts self-evolving reinforcement learning (RL) with information-gain reward derived from the internal knowledge. π_{ref} denotes SFT model trained in Stage 1. Dotted line is task-specific option.

supervised fine-tuning. Finally, in Section 2.4, we present reinforcement learning using information-gain reward derived from conditional mutual information.

2.2 Interleaved Latent Reasoning

To enable continuous-space reasoning compatible with policy-gradient optimization, we formulate the reasoning trajectory as an alternating sequence of text segments and continuous latent steps, both generated by the UMM backbone. Formally, given an input context $c = (c_1 \cdots c_T)$ consisting of a text prompt and optionally a source image, the model produces a reasoning trajectory followed by a target image x . Following Coconut [16] and subsequent latent reasoning work [25], the model alternates between two modes during trajectory generation: the *language mode* and the *latent mode*. In the language mode, the model emits text tokens through stochastic sampling, exposing the trajectory to the gradient signal of policy optimization. In the latent mode, the model emits a continuous reasoning step that bypasses text decoding, preserving the rich semantic content of internal representations.

Specifically, the UMM backbone first processes the input embedding $E_T = [e(c_1), \cdots, e(c_T)]$ to produce the last hidden state $H_T = \text{Transformer}(E_T)$, where e is the token embedding function. Once the model produces the special token `<latent>` at position t , the mode turns latent and the embedding of the next position is taken to be h_t the hidden state at position t in the last transformer layer. The mode returns to language after the fixed steps, at which point the trajectory ends and the image module is invoked. Concretely, with E_t denoting the input embedding sequence up to position t and $h_t = H_t[t, :]$ the corresponding hidden state, the next-position embeddings is

$$E_{t+1} = \begin{cases} [E_t; h_t] & \text{if latent mode,} \\ [E_t; e(x_{t+1})] & \text{if language mode,} \end{cases}$$

$$H_{t+1} = \text{Transformer}(E_{t+1}).$$

At inference, the mode alternation proceeds until the model emits the `<eos>` token, with each `<latent>` token marking the transition from a language-mode span to the next latent-mode step. After termination, the full trajectory of hidden states $h(\tau)$ is read by the image module through cross-attention to produce the target image x . The resulting reasoning trajectory naturally takes the form

$$\tau = (s_1, l_1, s_2, l_2, \cdots, s_N, l_N),$$

where s_n denotes the n -th language-mode span and l_n denotes the n -th latent-mode hidden state, each consisting of multiple tokens. Here, the trajectory length N and the length of each language-mode span $|s_n|$ are determined by when the model issues `<eos>` and `<latent>`, respectively, rather than fixed by an external schedule.

Table 1: **LARC achieves strong text-to-image generation performance on GenEval.** We report accuracy on GenEval subcategories and the overall score. Bold and underline indicate the best and second-best results within each group, respectively.

Type	Method	Single Obj. $\uparrow$	Two Obj. $\uparrow$	Counting $\uparrow$	Colors $\uparrow$	Position $\uparrow$	Attr. Binding $\uparrow$	Overall $\uparrow$
Reference	GPT-4o [26]	0.99	0.92	<u>0.85</u>	<u>0.92</u>	0.75	0.61	<u>0.84</u>
	SD3-medium [27]	0.99	0.94	0.72	0.89	0.33	0.46	<u>0.74</u>
	FLUX.1-dev [28]	0.99	0.81	0.79	0.74	0.20	0.47	0.67
Discrete UMM								
UMM	Janus-Pro [29]	0.99	0.89	0.59	0.90	0.79	0.66	0.80
	Emu3 [14]	0.98	0.71	0.34	0.81	0.17	0.21	0.54
	Continuous UMM							
UMM	Show-o2 [12]	1.00	0.87	0.58	<u>0.92</u>	0.52	0.62	0.76
	BAGEL [15]	0.99	0.94	0.80	<u>0.87</u>	0.64	0.63	0.81
	Uni-CoT [13]	0.99	<u>0.96</u>	0.84	<u>0.92</u>	0.57	<u>0.71</u>	0.83
	LARC (ours)	1.00	0.98	0.86	0.93	<u>0.77</u>	0.78	0.88

2.3 Stage 1: Supervised Fine-Tuning with Curriculum Learning

As observed in prior work [16], introducing all N latent steps from the beginning of training is unstable, since the latent stream lacks any direct supervision while the model has yet to learn the basic reasoning format. To progressively guide the model toward reasoning in latent space, we perform supervised fine-tuning (SFT) under a curriculum that gradually converts token-level reasoning into latent reasoning. We therefore decompose this stage into a curriculum that activates latent steps progressively, allowing each newly introduced latent step to build upon the trained prefix.

Data construction. For SFT with a curriculum, we collect 1M interleaved chain-of-thought (CoT) samples, providing the input context $\mathbf{c}$, a step-by-step reasoning chain, and the target image $\mathbf{x}$. We process each reasoning chain by inserting the special token $\langle \text{latent} \rangle$ at its step boundaries, marking the positions where latent steps are introduced.

Curriculum learning. To stabilize training, we adopt a curriculum learning that exposes the model to deeper interleaving over phase [16]. We first suppress all latent steps and train on the text-only trajectory $\tau^{(0)} = (s_1, \dots, s_N)$, establishing the reasoning format under standard next-token prediction. Subsequently, in phase k for $k = 1, \dots, N - 1$, we activate only the first k latent steps, ensuring that each newly introduced latent is built on top of the stable prefix. The final phase activates all latent steps, with the learning rate reduced at each phase transition.

Training loss. We optimize the model π_θ with a objective that combines a standard next-token prediction loss $\mathcal{L}_{\text{SFT}}(\cdot)$ over text segments and a flow-matching loss $\mathcal{L}_{\text{flow}}(\cdot)$ on the target image,

$$\mathcal{L}_{\text{SFT}}^{(k)} = - \sum_{n=1}^N \log \pi_\theta(s_n | \mathbf{c}, s_{<n}, l_{<\min(n,k+1)}) + \mathcal{L}_{\text{flow}}(\mathbf{x} | \mathbf{c}, s_{1:N}, l_{1:k}),$$

where $s_{<n}$ and $l_{<\min(n,k+1)}$ denote the text segments and the active latent steps preceding s_n at phase k , respectively. Note that we propagate the gradient through the latent steps without explicit supervision target. This implicit supervision allows the latent steps to encode useful representations automatically, imposing no external bias.

2.4 Stage 2: Self-Evolving Reinforcement Learning

Building on the latent reasoning capability established in Stage 1, we further refine the latent of reasoning trace through reinforcement learning (RL), guiding the model toward latent thoughts that are demonstrably useful for image generation. While supervised fine-tuning (SFT) effectively guides the model to produce interleaved trajectories, it remains fundamentally limited at refining the latent as it depends highly on the data quality, even the reasoning chain does not have ground-truth latent annotations. In this stage, we therefore adapt self-rewarding RL with an intrinsic reward derived from the UMM itself, yielding a fully self-evolving framework without any external critic.

Information-Gain reward. We define the reward as the *Information Gain* that a reasoning trace provides to a frozen reference model, measuring how much the latent reasoning improves the prediction of the target image. Concretely, letting π_{ref} denote the model after Stage 1 and marginal

Table 2: **LARC improves prompt-following fidelity on DPG-Bench.** We report scores across global, entity-level, attribute, relation, and other categories. Higher is better.

Type	Method	Global $\uparrow$	Entity $\uparrow$	Attribute $\uparrow$	Relation $\uparrow$	Other $\uparrow$	Overall $\uparrow$
Reference	SD3-medium [27]	87.90	91.01	88.83	80.70	88.68	84.08
	FLUX.1-dev [28]	82.10	89.50	88.70	91.10	89.40	84.00
UMM	Discrete UMM						
	Janus-Pro [29]	86.90	88.90	89.40	89.32	89.48	84.19
	Emu3 [14]	85.21	86.68	86.84	90.22	83.15	80.60
	Continuous UMM						
	Show-o2 [12]	89.00	91.78	89.96	91.81	91.64	86.14
	BAGEL [15]	88.94	90.37	91.29	90.82	88.67	85.07
	Uni-CoT [13]	87.58	89.14	91.62	91.09	86.97	85.56
	LARC (ours)	92.25	91.89	93.00	94.70	92.35	90.04

Table 3: **LARC substantially improves instruction-guided image generation on OneIG-Bench.** We report performance across prompt types and prompt lengths. NP refers natural language prompt. T&P denotes tag-based and phrase-based prompt. Higher is better.

Method	NP $\uparrow$	T&P $\uparrow$	Short $\uparrow$	Medium $\uparrow$	Long $\uparrow$	Overall $\uparrow$
Discrete UMM						
Janus-Pro [29]	0.557	0.533	0.609	0.548	0.515	0.552
BLIP-o3 [30]	0.719	0.671	0.754	0.712	0.674	0.706
Continuous UMM						
Show-o2 [12]	0.825	0.778	0.825	0.819	0.807	0.811
BAGEL [15]	0.776	0.734	0.782	0.769	0.759	0.764
Uni-CoT [13]	0.803	0.748	0.817	0.805	0.776	0.790
LARC (ours)	0.892	0.840	0.887	0.869	0.840	0.866

density of its image module $p_{\pi_{\text{ref}}}(\mathbf{x}|\cdot)$, the reward for a sampled trajectory τ is

$$r_{\text{IG}}(\tau; \mathbf{x}, \mathbf{c}) \triangleq \log p_{\pi_{\text{ref}}}(\mathbf{x}|\mathbf{c}, h(\tau)) - \log p_{\pi_{\text{ref}}}(\mathbf{x}|\mathbf{c}), \quad (1)$$

where $p_{\pi_{\text{ref}}}(\mathbf{x}|\mathbf{c}, h(\tau))$ is the predictive density when the trace is included and $p_{\pi_{\text{ref}}}(\mathbf{x}|\mathbf{c})$ is the reasoning-free baseline of the same model. Equation (1) assigns reward exactly when the trace raises likelihood of the target image in reference model, and zero reward when the trace carries no information beyond the context c .

Specifically, the Information-Gain reward r_{IG} is the per-sample form of conditional mutual information, with $\mathbb{E}[r_{\text{IG}}] = I_{\pi_{\text{ref}}}(\mathbf{x}; h(\tau)|\mathbf{c}) \geq 0$ under the joint distribution. Maximizing r_{IG} is therefore equivalent to maximizing the mutual information between the latent reasoning trace and the target image, with the bare context c playing the role of the marginal baseline. However, since the image module of UMM does not provide direct access to log-likelihoods, we thereby estimate r_{IG} as a paired difference of conditional flow-matching losses under shared noise,

$$r_{\text{IG}}(\tau; \mathbf{x}, \mathbf{c}) \approx \mathcal{L}_{\text{flow}}^{\pi_{\text{ref}}}(\mathbf{x}, \mathbf{c}) - \mathcal{L}_{\text{flow}}^{\pi_{\text{ref}}}(\mathbf{x}, \mathbf{c}, h(\tau)), \quad (2)$$

where $\mathcal{L}_{\text{flow}}^{\pi_{\text{ref}}}(\mathbf{x}, \xi)$ denotes the conditional flow-matching loss of the reference image module under conditioning ξ , with the two evaluations sharing the same noise sample for sample-wise constant cancellation. We provide the full derivation in Section C.2.

Total reward. To prevent the latent collapse where all N latent steps within a trajectory converge into identical vectors, we add an inter-step regularization for latent diversity,

$$r_{\text{div}}(\tau) = \frac{1}{N(N-1)} \sum_{n \neq n'} (1 - \cos(l_n, l_{n'})),$$

where $\cos(\cdot, \cdot)$ denotes the cosine similarity between two latent steps, l_n is the n -th latent-mode hidden state in trajectory τ , and N is the number of latent steps. The total reward used for policy optimization is then

$$R(\tau; \mathbf{x}, \mathbf{c}) = r_{\text{IG}}(\tau; \mathbf{x}, \mathbf{c}) + \lambda_{\text{div}} \cdot r_{\text{div}}(\tau). \quad (3)$$

Policy optimization. We optimize the policy π_θ through reinforcement learning with the following objective:

$$\max_{\pi_\theta} \mathbb{E}_{(\mathbf{c}, \mathbf{x}) \sim \mathcal{D}} \left[\mathbb{E}_{\tau \sim \pi_\theta(\cdot | \mathbf{c})} [R(\tau; \mathbf{x}, \mathbf{c})] \right]. \quad (4)$$

We adopt GRPO [31] as the optimization algorithm, which estimates advantages from group rollouts and removes an explicit value model. Its simplicity and stability allow a single set of hyperparameters to transfer across all benchmarks. For more optimization details, please see Appendix A.

3 Experiments

We provide an empirical evaluation of LARC by investigating the following questions:

- **RQ1:** Does LARC improve the performance on various tasks, *i.e.*, image generation, image editing, visual reasoning? (Table 1, 2, 3, 4, 5)
- **RQ2:** Does LARC enable test-time scaling of UMMs? (Figure 1)
- **RQ3:** Which component mostly contributes to LARC? (Table 6, 7, 8)

Evaluation Setup. We evaluate Unified Multimodal Models (UMMs) on various tasks including Text-to-Image (T2I) generation, image editing, and multimodal understanding. Specifically, we choose GenEval [20], DPG-Bench [21], and OneIG-Bench [22] for T2I benchmarks and CompBench [23] for image editing benchmark. For multimodal understanding, we choose MME [32], MMB [33], MMMU [34], and MIRA [24]. For T2I generation and image editing tasks, we follow the evaluation setup of target benchmarks, which mainly reports model-based metrics, including MLLM-as-a-Judge. We provide detailed evaluation setups in Section A.2.

Training Setup. In the main experiments, we trained LARC on BAGEL [15]. We perform supervised fine-tuning and reinforcement learning on a fixed 1M scale of Chain-of-Thought (CoT) data. We construct the dataset sampled from open-source datasets, *e.g.*, FLUX-Reason [17], WEAVE [18], ScaleEdit [19]. We provide more details in Section A.1.

Baselines. We compare LARC with image generation models, *e.g.*, Stable Diffusion3 [27] and FLUX.1-dev [17]) and foundation models in unified multimodal modeling, *e.g.*, BAGEL [15], Show-o2 [12], Emu3 [14], and Janus-Pro [29], and reasoning-enhanced UMM, Uni-CoT [13].

3.1 Image Generation Experiments

We evaluate the performance on various Text-to-Image (T2I) and image-editing benchmarks. Here, we mainly compare LARC with other UMMs and generation models, reporting score for each subcategory and overall averages. Specifically, we evaluate our LARC in (i) GenEval [20], (ii) DPG-Bench [21], (iii) OneIG-Bench [22], and (iv) CompBench [23].

LARC improves text-to-image generation performance. As shown in Table 1, LARC outperforms or consistently matches competing baselines across the evaluated objectives, demonstrating broad gains in T2I performance. The improvement is most pronounced on Attribute Binding, with a +15%p increase over the baseline. This subtask requires the joint handling of visual realization and compositional planning, suggesting that multimodal generation-aware latent reasoning is particularly effective for such compositionally demanding scenarios.

LARC improves prompt-following fidelity. We validate the prompt-following fidelity of LARC against competing baselines on DPG-Bench [21] and OneIG-Bench [22]. As shown in Table 2 and 3, LARC achieves consistent improvements over competing baselines. On DPG-Bench, the overall score of LARC improves by +5.8% (85.07 $\rightarrow$ 90.04) over the base model BAGEL. On OneIG-Bench, LARC yields a +10.2%p gain over the same base model. These improvements indicate that our latent reasoning is particularly effective in prompt-following generation. The consistent gains across every subcategory further suggest that LARC generalizes broadly across diverse prompt types.

LARC improves image editing. As shown in Table 4, LARC consistently outperforms all baselines across every metric. LARC achieves a +10.2% gain on LC-T over the base model BAGEL (20.434 $\rightarrow$ 22.509), indicating substantially stronger adherence to editing instructions. Pixel-level quality also improves markedly, with PSNR increasing from 24.3 to 25.7 and LPIPS decreasing from 0.069 to 0.049. These results suggest that continuous latent reasoning enables the model to plan spatially grounded edits, particularly for localized attribute modification.

Table 4: **LARC improves image editing on CompBench.** We report performance across score of CLIP for text (LC-T) and image (LC-I), PSNR, SSIM and LPIPS. LC-T refers local CLIP score between predicted foreground and descriptions while LC-I refers CLIP similarity between predicted foreground and ground-truth image. Higher is better except for LPIPS.

Method	LC-T $\uparrow$	LC-I $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Janus-Pro [29]	19.769	0.801	21.192	0.856	0.097
Emu3 [14]	19.536	0.803	22.036	0.868	0.092
Show-o2 [12]	20.350	0.831	24.661	0.919	0.065
BAGEL [15]	<u>20.434</u>	<u>0.842</u>	<u>24.370</u>	<u>0.917</u>	<u>0.069</u>
Uni-CoT [13]	20.215	0.837	23.925	0.904	0.088
LARC (ours)	22.509	0.880	25.710	0.953	0.049

Table 5: **LARC consistently enhances multimodal understanding of UMM.** We report matching score for MME, MMB, MMMU, and MIRA.

Method	MME [32] $\uparrow$	MMB [33] $\uparrow$	MMMU [34] $\uparrow$	MIRA [24] $\uparrow$
Janus-Pro [29]	1567.1	79.2	41.0	4.9
Emu3 [14]	1487.3	67.4	31.6	4.1
Show-o2 [12]	1620.5	79.3	48.9	6.9
BAGEL [15]	1687.0	85.0	55.3	7.5
Uni-CoT [13]	<u>1691.7</u>	<u>85.4</u>	<u>55.5</u>	<u>10.6</u>
LARC (ours)	1725.3	88.1	56.2	11.7

3.2 Text Generation Experiments

While the previous section established the benefits of latent reasoning for image generation and editing, an open question is whether such continuous reasoning also benefits text generation tasks that share the same UMM backbone. We address this question by evaluating LARC on perception-heavy benchmarks (MME, MMB) and a visual reasoning benchmark (MIRA), comparing against the base model BAGEL and Uni-CoT, a recent UMM with discrete token-level reasoning.

As shown in Table 5, LARC improves over both BAGEL and Uni-CoT across all benchmarks. The most pronounced gain appears on MIRA (7.5 $\rightarrow$ 11.7), a benchmark requiring visual chain-of-thought reasoning that explicitly demands multi-step intermediate inference. The substantial margin over Uni-CoT, which performs reasoning through discrete tokens, supports our hypothesis that latent reasoning provides richer intermediate representations than its discrete counterpart. We attribute this to the fact that continuous latent steps preserve information that discretization through a fixed vocabulary inevitably collapses, an effect amplified in tasks where reasoning must integrate fine-grained visual cues. Gains on MME and MMB further confirm that the continuous reasoning capability transfers positively to perception-heavy tasks without sacrificing general visual understanding. This indicates that the additional latent reasoning supervision introduced during training does not degrade the original visual perception ability of the UMM backbone, an important property for unified models that serve both generation and understanding.

3.3 Test-time Scaling

As shown in Figure 1, we investigate the effect of latent reasoning at test-time. LARC exhibits monotonically increasing performance on image generation, image editing, and visual reasoning as the latent budget grows, a behavior that is absent in prior reasoning-augmented UMMs. Furthermore, LARC consistently outperforms best-of-N sampling under matched compute budgets across all three tasks. We attribute this scaling behavior to continuous latent reasoning operating in the same space as image generation, allowing additional compute to refine useful representations rather than producing redundant discrete tokens.

3.4 Component Ablations

We conduct ablation studies to assess the contribution of each design choice in LARC, using GenEval, OneIG-Bench, CompBench, and MIRA as representative benchmarks.

Table 6: **Effect of latent reasoning compared to token-level reasoning.** We report scores on GenEval for T2I generation, OneIG-Bench for compositional generation, CompBench for image editing, and MIRA for visual reasoning. Token-level Reasoning denotes standard Chain-of-Thought on discrete reasoning space. Latent Reasoning refers LARC in Stage 1 trained by SFT.

Method	Reasoning space	T2I gen. $\uparrow$	Comp. gen. $\uparrow$	Image editing $\uparrow$	Visual reasoning $\uparrow$
Baseline [15]	Discrete	0.81	0.76	20.43	7.5
Token-level Reasoning	Discrete	0.82	0.79	20.81	7.1
Latent Reasoning	Continuous	0.83	0.81	21.46	9.7

Table 7: **Ablation on latent reasoning strategy.** We report scores on GenEval for T2I generation, OneIG-Bench for compositional generation, CompBench for image editing, and MIRA for visual reasoning. Naive refers latent reasoning without text segments proposed by the prior work [16].

Method	T2I gen. $\uparrow$	Comp. gen. $\uparrow$	Image editing $\uparrow$	Visual reasoning $\uparrow$
Baseline [15]	0.81	0.76	20.43	7.5
Naive	0.83	0.79	21.02	9.9
Interleaved (ours)	0.83	0.81	21.46	9.7

Effect of reasoning space. In Table 6, We compare token-level reasoning and latent reasoning initialized from the same base model. Switching from token-level to latent reasoning yields consistent gains across all four benchmarks. The improvement is particularly pronounced on visual reasoning (+2.6 on MIRA) and image editing (+0.65 on LC-T), confirming that continuous representations better capture the dense, spatially-grounded information required for generation tasks. Notably, token-level reasoning slightly degrades visual reasoning performance compared to the baseline, suggesting that discrete reasoning suffers the modality gap.

Effect of latent reasoning strategy In Table 7, we compare two latent reasoning strategies: Naive latent reasoning without interleaved text segments and our proposed interleaved design. While the Naive approach provides reasonable gains, interleaved latent reasoning achieves better compositional generation (+0.02 on OneIG-Bench) and image editing (+0.44 on LC-T) with comparable performance elsewhere. We attribute this to the text segments maintaining stochastic rollout diversity required for policy gradient optimization, while the latent steps encode the continuous representations that image generation demands.

Ablation study for training stages. In Table 8, we conduct ablation study for each training stage. While Stage 1 (SFT with curriculum) yields consistent improvements over the baseline across all tasks, Stage 2 (self-evolving RL) further boosts performance substantially, with T2I generation improving from 0.83 to 0.87 and visual reasoning from 9.7 to 11.7. This confirms that Stage 2 successfully refines latent traces beyond what SFT can achieve, and validates the self-evolving design that requires neither ground-truth reasoning chains nor an external critic.

4 Related Work

4.1 Test-time Scaling in LLMs

Test-time scaling is a model behavior where additional inference compute improves task performance in large language models (LLMs). Chain-of-Thought (CoT) prompting [35] has emerged as a primary driver of test-time scaling in LLMs by decomposing problems into intermediate textual reasoning steps. Subsequent work has internalized this capability through supervised fine-tuning [36] and reinforcement learning [31]. Inference-time search methods further expand the reasoning space through tree-based [37] and trajectory-based [38] exploration. Recognizing that natural language fails to encapsulate all forms of reasoning, recent paradigms shift reasoning directly into latent space [16, 25, 39]. Despite these advances, prior efforts to transfer test-time scaling to UMM image generation [12, 13, 14, 15, 29] have relied on discrete reasoning tokens and failed to exhibit the predictable scaling observed in language reasoning. We attribute this failure to the modality mismatch between discrete reasoning tokens and continuous image generation. Motivated by this observation, we propose LARC, which performs reasoning entirely in continuous space and becomes the first UMM to achieve test-time scaling for image generation.

Table 8: **Ablation on training stages.** We report scores on GenEval for T2I generation, OneIG-Bench for compositional generation, CompBench for image editing, and MIRA for visual reasoning.

Method	T2I gen. $\uparrow$	Comp. gen. $\uparrow$	Image editing $\uparrow$	Visual reasoning $\uparrow$
Baseline [15]	0.81	0.76	20.43	7.5
Stage 1	0.83	0.81	21.46	9.7
Stage 1 & 2	0.87	0.86	22.51	11.7

4.2 Image Generation Models

Image generation has emerged as one of the most influential research directions in modern deep learning, advancing along two complementary paradigms. Diffusion models [40, 41] have become the dominant paradigm for text-to-image generation, synthesizing images through iterative denoising in systems such as Stable Diffusion [1], DALL-E 2 [2], and Imagen [3]. In parallel, autoregressive approaches predict visual tokens from discrete vocabularies established by vector-quantized tokenizers [42, 43], scaling from DALL-E [4] to recent transformer-based models including LlamaGen [5] and VAR [44]. Building on these advances, unified multimodal models (UMMs) integrate visual understanding and generation within a single architecture, typically extending a pretrained LLM with image tokenization to inherit strong language priors and enable cross-task knowledge transfer. Early work pursues fully discrete tokenization with next-token prediction: Chameleon [45] pioneers an early-fusion design trained from scratch, and Emu3 [14] scales this paradigm across understanding and generation. Hybrid formulations combine autoregressive text decoding with diffusion-based image generation: Show-o [12, 46] uses discrete diffusion for images, and Transfusion [47] applies continuous diffusion over image patches. A separate line of work decouples visual representations for the two tasks: Janus [11] and Janus-Pro [29] introduce separate visual encoders under a shared transformer backbone, and BAGEL [15] extends decoupling to the transformer itself through a Mixture-of-Transformer-Experts architecture combined with rectified flow for image generation. Our work is orthogonal to these architectural advances, introducing a reasoning capability that operates over an existing UMM backbone, with BAGEL being adopted as our base models.

4.3 Reinforcement Learning in LLMs

Reinforcement learning has long served as the standard machinery for policy optimization. Trust-region methods such as TRPO [48] and its first-order successor PPO [49] stabilize policy updates through constrained or clipped objectives. More recent work adapts policy optimization to language modeling. GRPO [31] eliminates the value network by estimating advantages from group rollouts, and follow-up methods [50, 51, 52] further improve this group-relative formulation. These advances have powered substantial capability gains in LLMs, ranging from instruction-following in InstructGPT [53] to answer-verifiable reasoning in DeepSeek-R1 [54], where GRPO drives chain-of-thought reasoning at scale. However, for reasoning in image generation, no well-defined ground-truth reasoning trajectory exists. This precludes direct supervision over the reasoning chain and motivates the RL-based formulation adopted in LARC. We instantiate this formulation with GRPO, which is widely used in reinforcement learning tasks, to train reasoning chains using outcome-level rewards. This bypasses the need for chain-level annotations standard in text-based CoT, yet absent in our setting.

5 Conclusion

We introduced LARC, a latent reasoning framework that performs reasoning in continuous space for unified multimodal models (UMMs). By representing reasoning steps as continuous hidden states rather than discrete tokens, LARC alleviates the modality gap that has prevented prior reasoning-augmented UMMs from inheriting the test-time scaling of large language models. Our two-stage pipeline pairs curriculum supervised fine-tuning with a self-evolving reinforcement learning stage driven by an information-gain reward, requiring no external critic. Instantiated on BAGEL with 1M interleaved samples, LARC achieves state-of-the-art text-to-image performance with a 7%p gain over BAGEL and is the first reasoning-augmented UMM to exhibit clear test-time scaling for image generation. This finding bridges a long-standing modality gap in UMMs and opens a path toward UMMs that match the inference-time scalability of modern reasoning LLMs.

Broader Impact Statement

LARC advances the reasoning capability of unified multimodal models, with the goal of producing more controllable and inference-time-scalable image generation systems. We discuss the broader impacts below.

Positive impacts. By enabling test-time scaling for image generation, LARC allows users to allocate additional inference compute to obtain higher-fidelity outputs, particularly for prompts that require fine-grained compositional reasoning such as attribute binding and spatial layout. This controllability supports applications in design assistance, accessibility tools for visual communication, education, and scientific visualization, where prompt fidelity directly impacts utility. The latent reasoning paradigm also offers a unified mechanism for transferring reasoning capability across modalities, potentially reducing the need for modality-specific reasoning pipelines and lowering the engineering cost of building multimodal systems.

Potential risks. As with other text-to-image generation systems, LARC inherits risks common to large generative models, including the potential for generating misleading or harmful imagery, amplifying biases present in pretraining data, and being misused for the production of deceptive content such as misinformation or non-consensual depictions. The improved prompt-following fidelity demonstrated in our experiments increases generation quality across legitimate uses, and also amplifies these dual-use concerns, since outputs that more faithfully reflect adversarial prompts are correspondingly more capable of producing targeted harmful content. We do not introduce new training data or new modalities beyond those already established in BAGEL and the public datasets we use, so risks specific to dataset provenance are inherited rather than newly created.

Mitigation and responsible deployment. We recommend that downstream applications of LARC combine the model with established safeguards, including content moderation classifiers, watermarking of generated images, prompt-level safety filters, and auditing of training data for licensing and bias. For research deployments, releasing only model weights without bundled deployment infrastructure helps preserve the ability of downstream practitioners to apply context-appropriate safety measures. We will release LARC under a license that prohibits use for generating disinformation, non-consensual intimate imagery, or content targeting protected groups, and we encourage the community to develop evaluation protocols specifically for reasoning-augmented generative models, where the relationship between reasoning fidelity and harmful output remains understudied.

Environmental considerations. Training LARC requires reinforcement learning over a multi-billion-parameter UMM backbone, which incurs non-trivial compute and associated carbon footprint. We mitigate this by initializing from a pretrained BAGEL checkpoint rather than training from scratch and by using a curriculum that converges in fewer steps than naive joint training of all latent reasoning steps. Future work on parameter-efficient adaptation of latent reasoning, such as adapter-based or low-rank variants, would further reduce the resource cost of deploying this paradigm at scale.

References

- [1] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. [1](#), [9](#)
- [2] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022. [1](#), [9](#)
- [3] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022. [1](#), [9](#)
- [4] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International conference on machine learning*, pages 8821–8831. Pmlr, 2021. [1](#), [9](#)
- [5] Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*, 2024. [1](#), [9](#)
- [6] Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM transactions on Graphics (TOG)*, 42(4): 1–10, 2023. [1](#)
- [7] Weixi Feng, Xuehai He, Tsu-Jui Fu, Varun Jampani, Arjun Akula, Pradyumna Narayana, Sugato Basu, Xin Eric Wang, and William Yang Wang. Training-free structured diffusion guidance for compositional text-to-image synthesis. *arXiv preprint arXiv:2212.05032*, 2022. [1](#)
- [8] Royi Rassin, Eran Hirsch, Daniel Glickman, Shauli Ravfogel, Yoav Goldberg, and Gal Chechik. Linguistic binding in diffusion models: Enhancing attribute correspondence through attention map alignment. *Advances in Neural Information Processing Systems*, 36:3536–3559, 2023. [1](#)
- [9] David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2(3):440, 2018. [1](#)
- [10] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019. [1](#)
- [11] Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified multimodal understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 12966–12977, 2025. [1](#), [9](#)
- [12] Jinheng Xie, Zhenheng Yang, and Mike Zheng Shou. Show-o2: Improved native unified multimodal models. *arXiv preprint arXiv:2506.15564*, 2025. [1](#), [4](#), [5](#), [6](#), [7](#), [8](#), [9](#)
- [13] Luozheng Qin, Jia Gong, Yuqing Sun, Tianjiao Li, Mengping Yang, Xiaomeng Yang, Chao Qu, Zhiyu Tan, and Hao Li. Uni-cot: Towards unified chain-of-thought reasoning across text and vision. *arXiv preprint arXiv:2508.05606*, 2025. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [14] Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiyang Yu, et al. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*, 2024. [1](#), [4](#), [5](#), [6](#), [7](#), [8](#), [9](#)
- [15] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, et al. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [8](#), [9](#), [15](#), [17](#), [18](#), [19](#), [20](#)
- [16] Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. Training large language models to reason in a continuous latent space. *arXiv preprint arXiv:2412.06769*, 2024. [2](#), [3](#), [4](#), [8](#)
- [17] Rongyao Fang, Aldrich Yu, Chengqi Duan, Linjiang Huang, Shuai Bai, Yuxuan Cai, Kun Wang, Si Liu, Xihui Liu, and Hongsheng Li. Flux-reason-6m and prism-bench: A million-scale text-to-image reasoning dataset and comprehensive benchmark. *arXiv preprint arXiv:2509.09680*, 2025. [2](#), [6](#), [16](#)
- [18] Wei Chow et al. Weave: A benchmark for evaluating multimodal editing models. *arXiv preprint arXiv:2511.15738*, 2024. [2](#), [6](#), [16](#)

- [19] Guanzhou Chen, Erfei Cui, Changyao Tian, Danni Yang, Ganlin Yang, Yu Qiao, Hongsheng Li, Gen Luo, and Hongjie Zhang. Scaleedit-12m: Scaling open-source image editing data generation via multi-agent framework. *arXiv preprint arXiv:2603.20644*, 2026. 2, 6, 16
- [20] Dhruva Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems*, 36:52132–52152, 2023. 2, 6, 15
- [21] Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024. 2, 6, 15
- [22] Jingjing Chang, Yixiao Fang, Peng Xing, Shuhan Wu, Wei Cheng, Rui Wang, Xianfang Zeng, Gang Yu, and Hai-Bao Chen. Oneig-bench: Omni-dimensional nuanced evaluation for image generation. *arXiv preprint arxiv:2506.07977*, 2025. 2, 6, 15
- [23] Bohan Jia, Wenxuan Huang, Yuntian Tang, Junbo Qiao, Jincheng Liao, Shaosheng Cao, Fei Zhao, Zhaopeng Feng, Zhouhong Gu, Zhenfei Yin, et al. Compbench: Benchmarking complex instruction-guided image editing. *arXiv preprint arXiv:2505.12200*, 2025. 2, 6, 15
- [24] Yiyang Zhou, Haoqin Tu, Zijun Wang, Zeyu Wang, Niklas Muennighoff, Fan Nie, Yejin Choi, James Zou, Chaorui Deng, Shen Yan, et al. When visualizing is the first step to reasoning: Mira, a benchmark for visual chain-of-thought. *arXiv preprint arXiv:2511.02779*, 2025. 2, 6, 7, 15
- [25] Byungwoo Jeon, Yoonwoo Jeong, Hyunseok Lee, Minsu Cho, and Jinwoo Shin. Vision-aligned latent reasoning for multi-modal large language model. In *International Conference on Machine Learning (ICML)*, 2026. 3, 8
- [26] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 4, 15
- [27] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024. 4, 5, 6
- [28] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, Sumith Kulal, Kyle Lacey, Yam Levi, Cheng Li, Dominik Lorenz, Jonas Müller, Dustin Podell, Robin Rombach, Harry Saini, Axel Sauer, and Luke Smith. Flux.1 kontext: Flow matching for in-context image generation and editing in latent space, 2025. URL <https://arxiv.org/abs/2506.15742>. 4, 5
- [29] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025. 4, 5, 6, 7, 8, 9
- [30] Jiu hai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi Zhou, Saining Xie, Silvio Savarese, et al. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*, 2025. 5
- [31] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024. 6, 8, 9, 16, 17
- [32] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023. 6, 7, 15
- [33] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? In *European conference on computer vision*, pages 216–233. Springer, 2024. 6, 7, 15
- [34] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoyi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9556–9567, 2024. 6, 7, 15
- [35] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, pages 24824–24837, 2022. 8

- [36] Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen. Mammoth: Building math generalist models through hybrid instruction tuning. *arXiv preprint arXiv:2309.05653*, 2023. 8
- [37] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pages 11809–11822, 2023. 8
- [38] Kanishk Gandhi, Denise Lee, Gabriel Grand, Muxin Liu, Winson Cheng, Archit Sharma, and Noah D Goodman. Stream of search (sos): Learning to search in language. *arXiv preprint arXiv:2404.03683*, 2024. 8
- [39] Bangzheng Li, Ximeng Sun, Jiang Liu, Ze Wang, Jialian Wu, Xiaodong Yu, Hao Chen, Emad Barsoum, Muhao Chen, and Zicheng Liu. Latent visual reasoning. *arXiv preprint arXiv:2509.24251*, 2025. 8
- [40] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 9
- [41] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 9
- [42] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 9
- [43] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. 9
- [44] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *Advances in neural information processing systems*, 37:84839–84865, 2024. 9
- [45] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024. 9
- [46] Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin, Yuchao Gu, Zhijie Chen, Zhenheng Yang, and Mike Zheng Shou. Show-o: One single transformer to unify multimodal understanding and generation. *arXiv preprint arXiv:2408.12528*, 2024. 9
- [47] Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Jacob Kahn, Xuezhe Ma, Luke Zettlemoyer, and Omer Levy. Transfusion: Predict the next token and diffuse images with one multi-modal model. *arXiv preprint arXiv:2408.11039*, 2024. 9
- [48] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015. 9
- [49] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 9
- [50] Shih-Yang Liu, Xin Dong, Ximing Lu, Shizhe Diao, Peter Belcak, Mingjie Liu, Min-Hung Chen, Hongxu Yin, Yu-Chiang Frank Wang, Kwang-Ting Cheng, et al. Gdpo: Group reward-decoupled normalization policy optimization for multi-reward rl optimization. *arXiv preprint arXiv:2601.05242*, 2026. 9
- [51] Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025. 9
- [52] Chiyu Ma, Shuo Yang, Kexin Huang, Jinda Lu, Haoming Meng, Shangshang Wang, Bolin Ding, Soroush Vosoughi, Guoyin Wang, and Jingren Zhou. Fipo: Eliciting deep reasoning with future-kl influenced policy optimization. *arXiv preprint arXiv:2603.19835*, 2026. 9
- [53] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022. 9
- [54] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-rl: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 9

- [55] Haodong Duan, Junming Yang, Yuxuan Qiao, Xinyu Fang, Lin Chen, Yuan Liu, Xiaoyi Dong, Yuhang Zang, Pan Zhang, Jiaqi Wang, et al. Vlmevalkit: An open-source toolkit for evaluating large multi-modality models. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 11198–11201, 2024. [15](#)
- [56] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36:53728–53741, 2023. [16](#)
- [57] Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam, Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. Diffusion model alignment using direct preference optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8228–8238, 2024. [16](#)
- [58] Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36:15903–15935, 2023. [16](#)
- [59] Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *Advances in neural information processing systems*, 36:36652–36663, 2023. [16](#)
- [60] Chengzhuo Tong, Ziyu Guo, Renrui Zhang, Wenyu Shan, Xinyu Wei, Zhenghao Xing, Hongsheng Li, and Pheng-Ann Heng. Delving into rl for image generation with cot: A study on dpo vs. grpo. *arXiv preprint arXiv:2505.17017*, 2025. [16](#)
- [61] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. [16](#)

A Implementation Details

A.1 Training Details

We adopt BAGEL [15] as the base model and perform two-stage learning, namely (1) Supervised Fine-tuning (SFT) with curriculum and (2) Self-evolving Reinforcement Learning (RL). In Stage 1, we freeze the vision tokenizer and jointly train the understanding expert and the generation expert of the UMM backbone. We follow the curriculum schedule described in Section 2.3, where the number of activated latent steps is increased from 0 to N over $N+1$ phases with the learning rate decayed at every phase transition. In Stage 2, we initialize the policy π_θ from $\mathcal{M}_{\text{SFT}}$ and freeze a copy of the same model as the reference π_{ref} . Only the understanding expert is updated during RL, since the generation expert is reused for reward estimation through the shared-noise flow-matching loss. Detailed hyperparameters are summarized in Table 9. All training experiments are conducted with $8 \times$ NVIDIA A100s, taking approximately 300 GPU hours in Stage 1 and 96 GPU hours in Stage 2.

Table 9: **Hyperparameters for Stage 1 and Stage 2.**

Hyperparameter	Stage 1 (SFT)	Stage 2 (RL)
Optimizer	Adam	Adam
Algorithm	-	GRPO
Learning rate	$1\text{e-}5$	$1\text{e-}6$
Global batch size	-	64
Max number of tokens	6×10^5	-
Weight decay	10^{-2}	0.0
Training steps	1.2×10^6	10^6
Group size G	-	8
Clipping ϵ_{clip}	-	2×10^{-1}
KL coefficient β	-	10^{-3}
Diversity weight λ_{div}	-	10^{-1}

A.2 Evaluation Details

We evaluate LARC across three categories of benchmarks following the official protocol of each benchmark. For inference, we apply greedy sampling and generate one image per prompt unless specified otherwise.

Text-to-Image generation. We evaluate the models on GenEval [20], DPG-Bench [21], and OneIG-Bench [22]. GenEval automatically verifies compositional properties of generated images via an object-detection pipeline, reporting accuracy across six sub-tasks, *e.g.*, single object, two objects, counting, colors, position, and attribute binding. DPG-Bench evaluates 1,065 dense prompts through an MLLM-as-a-Judge protocol, providing scores across global, entity, attribute, relation, and other categories. OneIG-Bench spans multiple dimensions of text-to-image (T2I) evaluation, including prompt-image alignment, text rendering, and reasoning, with scores partitioned by prompt type (NP, T&P) and prompt length (short, medium, long). Specifically, we report alignment score for OneIG-Bench.

Image editing. We evaluate the models on CompBench [23], which features 3,200 complex instruction-guided editing tasks built through an MLLM-human collaborative framework. Specifically, we report multi-object editing tasks in Table 4. Following the official protocol, we report local CLIP-Text (LC-T) score between the predicted foreground and the description, local CLIP-Image (LC-I) similarity between the predicted foreground and the ground-truth image, along with pixel-level metrics PSNR, SSIM, and LPIPS.

Multimodal understanding. We evaluate on MME [32], MMBench [33], MMMU [34], and MIRA [24]. MME measures perception and cognition across 14 sub-tasks, while MMBench and MMMU benchmark broad multimodal reasoning across diverse disciplines. MIRA targets visual chain-of-thought reasoning with 546 carefully curated problems spanning 20 task types, where successful reasoning requires generating intermediate visual representations.

For API-based baselines, we report results using the model version gpt-4o-2024-08-06 for GPT-4o [26]. We adopt VLMEvalKit [55] as the unified evaluation toolkit for multimodal understanding tasks, and use the official codebase released by each benchmark for T2I generation and image editing evaluation.

B Data Statistics

We construct a 1M scale reasoning dataset by sampling from three open-source datasets, namely FLUX-Reason [17], WEAVE [18], and ScaleEdit [19]. The detailed composition is summarized in Table 10. We balance the mixture so that text-to-image, image editing, and interleaved reasoning are jointly represented, ensuring that the latent reasoning trajectory generalizes across diverse generation tasks. For each sample, we insert the special token `<latent>` at every reasoning step boundary as described in Section 2.3, marking the positions where latent steps are progressively activated by the curriculum.

Table 10: Composition of the 1M reasoning dataset.

Source dataset	Task	# samples
FLUX-Reason [17]	Text-to-Image	500K
WEAVE [18]	Text-to-Image	100K
ScaleEdit [19]	Image Editing	400K
Total	–	1M

C Self-Evolving Reinforcement Learning

C.1 Why Self-Rewarding RL?

In this section, we discuss why standard reward-based reinforcement learning (RL) paradigms are not directly applicable to latent reasoning in UMMs, motivating the self-rewarding formulation adopted in LARC.

Inapplicability of preference-based methods. Direct Preference Optimization (DPO) [56] and its diffusion variants such as Diffusion-DPO [57] optimize the policy from paired preference data (τ^w, τ^l) , where τ^w is preferred over τ^l . This paradigm presumes the existence of a meaningful preference order between trajectories. In our setting, however, each trajectory τ contains continuous latent steps $h(\tau)$ that lie outside any human-interpretable space, ruling out direct human annotation over the reasoning trace. External reward models for image generation [58, 59] also operate on the rendered image x rather than the latent chain, providing no signal that distinguishes between two latent traces that yield comparable images. A recent study [60] further reports that DPO-based optimization in autoregressive image generation tends to overfit to in-domain preference distributions, limiting generalization beyond the curated preference set. These observations indicate that preference-based methods are structurally misaligned with the latent reasoning regime, which lacks both pairwise preference labels and a natural notion of pairwise dominance over hidden states.

Limitations of critic-based methods. Proximal Policy Optimization (PPO) [61] sidesteps the preference assumption but introduces a separate value network that estimates the expected return at each position, which roughly doubles the trainable parameters and adds non-trivial training overhead at the scale of UMMs. More fundamentally, fitting an accurate critic over the latent reasoning trajectory requires a dense external supervision signal, which is precisely what is absent in our setting. Without ground-truth latent annotations, the critic must rely on an external image-quality estimator or a learned reward model, both of which inherit the same evaluation gap described above.

Self-rewarding via information gain. We point out two requirements that the optimization framework must satisfy. (1) The reward should be derivable from the model itself, since no external evaluator can faithfully assess latent reasoning traces. (2) The policy gradient estimator should avoid an additional value network to keep training tractable on top of the UMM backbone. In LARC, the information-gain reward satisfies the first requirement by reusing the frozen reference π_{ref} obtained from Stage 1, scoring each trajectory through the change in target-image likelihood. Additionally, GRPO [31] satisfies the second by replacing the value network with group-relative advantages estimated from multiple rollouts.

C.2 Derivation of the Information-Gain Reward

We provide the full derivation of the Information-Gain reward r_{IG} introduced in Section 2.4. Recall that the reward is defined as

$$r_{\text{IG}}(\tau; \mathbf{x}, \mathbf{c}) \triangleq \log p_{\pi_{\text{ref}}}(\mathbf{x} \mid \mathbf{c}, h(\tau)) - \log p_{\pi_{\text{ref}}}(\mathbf{x} \mid \mathbf{c}), \quad (5)$$

where $h(\tau)$ denotes the sequence of latent hidden states produced along trajectory τ , and $\mathbf{x}$ is the target image. Taking the expectation of Equation 5 under the joint distribution $(\mathbf{x}, \mathbf{c}, \tau)$ recovers the conditional mutual information,

$$\mathbb{E}[r_{\text{IG}}] = \mathbb{E}\left[\log \frac{p_{\pi_{\text{ref}}}(\mathbf{x} | \mathbf{c}, h(\tau))}{p_{\pi_{\text{ref}}}(\mathbf{x} | \mathbf{c})}\right] = I_{\pi_{\text{ref}}}(\mathbf{x}; h(\tau) | \mathbf{c}) \geq 0. \quad (6)$$

Maximizing r_{IG} is therefore equivalent to maximizing the mutual information between the latent reasoning trace and the target image, with the bare context $\mathbf{c}$ playing the role of the marginal baseline. The non-negativity in Equation 6 follows from Jensen’s inequality applied to the log-ratio, which guarantees that the expected reward is bounded below by zero regardless of the policy.

Since the generation expert of UMM does not provide closed-form access to the marginal density $p_{\pi_{\text{ref}}}(\mathbf{x} | \cdot)$, we approximate the log-likelihood through the conditional flow-matching objective. Following the rectified flow formulation of BAGEL [15], the negative log-likelihood admits the bound

$$-\log p_{\pi_{\text{ref}}}(\mathbf{x} | \xi) = \mathcal{L}_{\text{flow}}^{\pi_{\text{ref}}}(\mathbf{x}, \xi) + C(\mathbf{x}), \quad (7)$$

where ξ denotes the conditioning variable, $\mathcal{L}_{\text{flow}}^{\pi_{\text{ref}}}$ is the conditional flow-matching loss evaluated under π_{ref} , and $C(\mathbf{x})$ is a sample-wise constant that depends only on the target image and the noise schedule. Substituting Equation 7 into Equation 5 for both terms yields

$$\begin{aligned} r_{\text{IG}}(\tau; \mathbf{x}, \mathbf{c}) &= [-\mathcal{L}_{\text{flow}}^{\pi_{\text{ref}}}(\mathbf{x}, \mathbf{c}, h(\tau)) - C(\mathbf{x})] - [-\mathcal{L}_{\text{flow}}^{\pi_{\text{ref}}}(\mathbf{x}, \mathbf{c}) - C(\mathbf{x})] \\ &= \mathcal{L}_{\text{flow}}^{\pi_{\text{ref}}}(\mathbf{x}, \mathbf{c}) - \mathcal{L}_{\text{flow}}^{\pi_{\text{ref}}}(\mathbf{x}, \mathbf{c}, h(\tau)). \end{aligned} \quad (8)$$

Crucially, the constant $C(\mathbf{x})$ cancels exactly when the two flow-matching losses are evaluated under the same noise sample $\eta \sim \mathcal{N}(0, I)$ and the same time step $t \sim \mathcal{U}[0, 1]$, since both terms inherit the identical perturbation. We therefore implement the reward as a paired evaluation under shared randomness, as detailed in Algorithm 1. This shared-noise estimator yields an unbiased and low-variance approximation of Equation 5, requiring no additional reward model and reusing the frozen π_{ref} that is already available from Stage 1.

Algorithm 1 Self-Evolving Reinforcement Learning of LARC

Require: Frozen reference $\pi_{\text{ref}} = \mathcal{M}_{\text{SFT}}$, dataset $\mathcal{D}$, hyperparameters $G, \epsilon_{\text{clip}}, \beta, \lambda_{\text{div}}$

- 1: Initialize $\pi_{\theta} \leftarrow \pi_{\text{ref}}$
 - 2: **for** each RL step **do**
 - 3: $\pi_{\theta_{\text{old}}} \leftarrow \pi_{\theta}$
 - 4: Sample $(\mathbf{c}, \mathbf{x}) \sim \mathcal{D}$ and rollout $\{\tau_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot | \mathbf{c})$
 - 5: **for** $i = 1, \dots, G$ **do**
 - 6: $r_{\text{IG}}(\tau_i) \leftarrow \mathcal{L}_{\text{flow}}^{\pi_{\text{ref}}}(\mathbf{x}, \mathbf{c}) - \mathcal{L}_{\text{flow}}^{\pi_{\text{ref}}}(\mathbf{x}, \mathbf{c}, h(\tau_i))$ $\triangleright$ Shared noise
 - 7: $R_i \leftarrow r_{\text{IG}}(\tau_i) + \lambda_{\text{div}} \cdot r_{\text{div}}(\tau_i)$
 - 8: **end for**
 - 9: $A_i \leftarrow (R_i - \mu_R) / (\sigma_R + \epsilon)$ for $i = 1, \dots, G$
 - 10: Update θ by maximizing $\mathcal{J}_{\text{GRPO}}(\theta)$ $\triangleright$ See Equation (9)
 - 11: **end for**
 - 12: **return** $\mathcal{M}_{\text{LARC}} \leftarrow \pi_{\theta}$
-

C.3 Policy Optimization

We instantiate the policy gradient optimization in Equation 4 using GRPO [31], which estimates advantages from group rollouts and removes the need for a separate value network. For each context $\mathbf{c}$ sampled from $\mathcal{D}$, we draw G trajectories $\{\tau_i\}_{i=1}^G$ from the old policy $\pi_{\theta_{\text{old}}}$, compute the total reward R_i for each rollout, and normalize within the group to obtain the advantage A_i . The policy is then updated by maximizing the clipped surrogate objective

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}\left[\frac{1}{G} \sum_{i=1}^G \min\left(\rho_i(\theta) A_i, \text{clip}(\rho_i(\theta), 1 - \epsilon_{\text{clip}}, 1 + \epsilon_{\text{clip}}) A_i\right) - \beta \text{KL}[\pi_{\theta} \| \pi_{\text{ref}}]\right], \quad (9)$$

where $\rho_i(\theta) = \frac{\pi_{\theta}(\tau_i | \mathbf{c})}{\pi_{\theta_{\text{old}}}(\tau_i | \mathbf{c})}$ is the importance ratio. The full training procedure is summarized in Algorithm 1.

D Additional Experimental Results

D.1 Detailed Analysis on Test-time Scaling

We provide an additional analysis on performance and computational cost incurred by latent reasoning at test-time. We measure the per-sample clock-time on a single NVIDIA A100 with batch size 1, varying the number of images N and latent steps L from 1 to 10. We compare LARC against best-of- N sampling on the base model BAGEL [15], which represents the standard test-time scaling baseline that draws N independent samples and selects the best one. We analyze clock time (seconds) of test-time scaling in Table 11.

Table 11: Clock-time analysis under test-time scaling.

Method	$N=1$	$N=2$	$N=4$	$N=6$	$N=8$	$N=10$
BAGEL [15] + Best-of- N	46.5	92.8	187.1	275.5	367.7	469.0
	$L=1$	$L=2$	$L=4$	$L=6$	$L=8$	$L=10$
LARC (Ours)	46.8	46.9	47.0	47.3	47.6	48.0

Note that LARC only requires single-pass while best-of- N demands iterative sampling. As shown in Table 11, LARC incurs only a marginal overhead as the number of latent steps increases, requiring additional 1.2 seconds at $L = 10$ compared to $L = 1$. Notably, combined with the results shown in Figure 1, LARC outperforms best-of- N on every scale with highly efficient computational cost, performing consistent test-time scaling behavior. We attribute this efficiency to the fact that latent reasoning extends only a small portion of the forward pass through the understanding expert, whereas best-of- N requires repeated invocation of the full generation expert. These results confirm that continuous latent reasoning offers a practically efficient route to test-time scaling for UMMs.

D.2 Additional Ablation Study

Ablation study for latent diversity reward. We investigate the role of the collapse-prevention reward r_{div} introduced in Equation 3. Without r_{div} , we observe that the policy quickly collapses to trivial trajectories where all N latent steps converge to nearly identical hidden states, eliminating the benefit of multi-step reasoning. To quantify this phenomenon, we report the average pairwise cosine similarity among latent steps along with task performance under varying λ_{div} in Table 12.

Table 12: Effect of the latent diversity reward. We report the average pairwise cosine similarity (Cos. Sim.) among latent steps within a trajectory, and accuracy on GenEval, OneIG-Bench, and MIRA.

λ_{div}	Cos. Sim.	GenEval $\uparrow$	OneIG-Bench $\uparrow$	MIRA $\uparrow$
0.0	0.97	0.83	0.81	9.8
0.05	0.87	0.85	0.84	10.9
0.1 (Ours)	0.79	0.87	0.86	11.7
0.2	0.71	0.85	0.85	11.4
0.5	0.65	0.82	0.82	10.6

As shown in Table 12, removing the diversity reward ($\lambda_{\text{div}}=0$) leads to a pairwise cosine similarity of 0.97 across latent steps, indicating severe collapse where every step encodes nearly redundant information. This collapse degrades all three benchmarks, with MIRA dropping from 11.7 to 9.8. Setting $\lambda_{\text{div}}=0.1$ yields the best trade-off. These results validate that the diversity regularizer is essential for preventing latent collapse and confirm that the chosen weight balances exploration of distinct latent directions with exploitation of informative ones.

D.3 Qualitative Results

In this section, we provide qualitative results of LARC.




Figure 3: **Qualitative comparison during test-time.** We compare BAGEL [15] with best-of- N sampling to LARC on test-time scaling with L latent steps on text-to-image generation task. First row presents $N = L = 1$, second row presents $N = L = 4$, and third row presents $N = L = 10$.




Figure 4: **Qualitative comparison during test-time.** We compare BAGEL [15] with best-of- N sampling to LARC on test-time scaling with L latent steps on image editing task. First row presents $N = L = 1$, second row presents $N = L = 4$, and third row presents $N = L = 10$.

E Limitations and Future Work

While LARC demonstrates strong empirical gains and clear test-time scaling on various generation tasks, several directions remain open for future exploration. First, our framework is currently instantiated on UMMs whose image module is a flow-matching head, which generates the image in continuous space. Extending the same principle to UMMs with discrete space generation or other generative heads is a natural next step. Second, although LARC enables effective test-time scaling, the interleaved latent reasoning trajectory introduces additional computation at inference, leading to latency overhead compared to the base UMM. Reducing this overhead through efficient decoding strategy, distillation, or adaptive latent allocation is a promising research direction. We leave these extensions to future work.