REPROMPT: REASONING-AUGMENTED REPROMPTING FOR TEXT-TO-IMAGE GENERATION VIA REINFORCEMENT LEARNING

Anonymous authorsPaper under double-blind review

ABSTRACT

Despite recent progress in text-to-image (T2I) generation, existing models often struggle to faithfully capture user intentions from short and under-specified prompts. While prior work has attempted to enhance prompts using large language models (LLMs), these methods frequently generate stylistic or unrealistic content due to insufficient grounding in visual semantics and real-world composition. Inspired by recent advances in reasoning for language model, we propose RePrompt, a novel reprompting framework that introduces explicit reasoning into the prompt enhancement process via reinforcement learning. Instead of relying on handcrafted rules or stylistic rewrites, our method trains a language model to generate structured, self-reflective prompts by optimizing for image-level outcomes. The tailored reward models assesse the generated images in terms of human preference, semantic alignment, and visual composition, providing indirect supervision to refine prompt generation. Our approach enables end-to-end training without human-annotated data. Experiments on GenEval and T2I-Compbench show that RePrompt significantly boosts spatial layout fidelity and compositional generalization across diverse T2I backbones, establishing new state-of-the-art results. Code is available at: https://anonymous.4open.science/r/ RePrompt-CD21.

1 Introduction

Text-to-image (T2I) generation has made rapid progress with the rise of large-scale generative models Labs (2024); Esser et al. (2024); Podell et al. (2023); Chen et al. (2024b), yet a persistent challenge remains: users typically provide concise and under-specified prompts, which often result in images that fail to reflect the intended semantics or visually coherent compositions. Generated outputs may misrepresent object counts, overlook spatial relations, or violate real-world plausibility. This misalignment arises from the gap between natural language descriptions and the structured visual logic required for faithful image generation Yang et al. (2024b).

Previous work on prompt enhancement in T2I can be divided into two main approaches. The first approach focuses on iterative refinement: an image is generated from an initial prompt, and subsequent feedback, derived either from human preference models or from automated scoring systems, is used to improve the prompt or intermediate representations over multiple rounds Yang et al. (2024d); Wu et al. (2024); Wang et al. (2024); Guo et al. (2025b). Although this approach can progressively improve image quality, it suffers from high latency and computational overhead due to repeated image generation, and it rarely incorporates explicit scene semantics or compositional reasoning. The second approach enriches prompts in a single pass by leveraging large language models (LLMs) to inject additional detail and context Betker et al. (2023); Hao et al. (2023). While these methods produce linguistically fluent and expressive descriptions, they frequently generate prompts that produce images with semantically inconsistent or visually implausible content, such as conflicting object placements or unrealistic interactions, because the underlying LLMs lack grounding in physical reality and do not incorporate feedback from downstream visual tasks. As a result, they frequently hallucinate content or miss critical spatial and attribute-level relationships (see Figure 1).

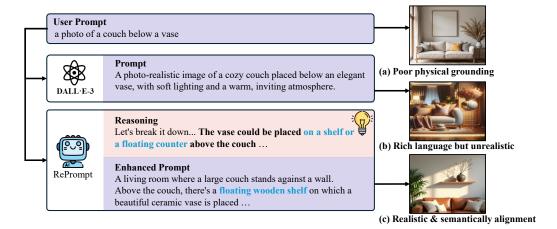


Figure 1: Given the user prompt "a photo of a couch below a vase", existing models like DELL-E3 generate rich language descriptions but often produce unrealistic or physically implausible compositions. In contrast, our RePrompt performs explicit chain-of-thought reasoning to resolve spatial relations, resulting in enhanced prompts that guide text-to-image models towards realistic and semantically aligned generations.

In contrast, we propose RePrompt, a reasoning-augmented prompt refinement framework trained via reinforcement learning. Rather than relying on stylistic rewriting or black-box completions, RePrompt trains a language model to generate structured and semantically grounded prompts through self-reflection and step-by-step decomposition. Motivated by recent advances in reasoning-augmented language models Guo et al. (2025a); Trung et al. (2024); Team et al. (2025); Jaech et al. (2024), RePrompt enables the model to internally simulate the visual implications of a prompt—much like how humans mentally visualize a scene before drawing. This structured, logic-driven process anticipates potential errors (e.g., conflicting object positions, missing entities, or spatial incoherence) during prompt construction, thereby reducing the need for multiple rounds of image generation.

A core component of RePrompt is a **T2I RePrompt Reward Model** tailored for text-to-image generation. Instead of relying on pre-labeled reasoning traces or handcrafted prompt templates, RePrompt learns from downstream visual feedback by optimizing prompt generation through reinforcement learning. To capture the multifaceted nature of image quality, we design a ensemble reward that evaluates generated images along three dimentions: human preference, visual realism, and semantic alignment with the input. By learning from diverse and grounded feedback signals, the model develops a more robust reasoning strategy that transfers across prompt types, scene structures, and T2I backbones, enabling stronger performance on unseen inputs without overfitting to specific linguistic patterns.

Experiments on GenEval Ghosh et al. (2023) and T2I-Compbench Huang et al. (2023) demonstrate that RePrompt significantly improves compositional accuracy, semantic fidelity, and spatial coherence. Notably, on the GenEval benchmark, RePrompt surpasses Qwen2.5 3B-enhanced baselines by +77.1% (FLUX Labs (2024)), +78.8% (SD3 Esser et al. (2024)) and +122.2% (Pixart- Σ Chen et al. (2024b)) in the position category, highlighting its superior capability in grounding spatial relations. Furthermore, RePrompt achieves the best overall accuracy (0.76) among all evaluated methods while maintaining an order-of-magnitude lower latency than optimization-heavy baselines like Idea2Img (30s vs. 140s per image), offering a scalable and inference-efficient solution. These findings validate the effectiveness of explicit reasoning in prompt construction for closing the semantic-visual alignment gap in text-to-image generation, without relying on larger language models or expensive optimization at inference time.

2 Related Work

Text to Image Generation. Recently, large-scale diffusion models Labs (2024); Esser et al. (2024); Podell et al. (2023); Chen et al. (2024b); Xie et al. (2025); Liu et al. (2024); Betker et al. (2023);

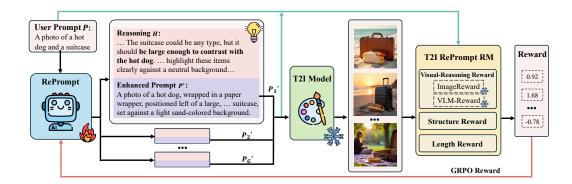


Figure 2: Overview of the proposed RePrompt. For each input prompt, RePrompt generates multiple reasoning trace and enhanced prompt pairs. The reasoning trace guides the model to produce more detailed, image-grounded prompts. These are used to synthesize candidate images via a T2I model, which are then scored by a reward model. Feedback is used to update RePrompt via GRPO.

Ma et al. (2025a); Rombach et al. (2022); Saharia et al. (2022) have achieved impressive progress in generating high-resolution, photorealistic images from complex textual prompts. To enhance alignment between text and visuals, prior work has explored prompt engineering Hao et al. (2023); Mo et al. (2024); Yeh et al. (2024); Mañas et al. (2024); Yun et al. (2025); Cao et al. (2023); Qin et al. (2024); Yang et al. (2024d); Wu et al. (2024); Wang et al. (2024), often relying on manual or heuristic strategies with limited generalization. We propose a reinforcement learning-based framework that automatically refines prompts through iterative reasoning, achieving better semantic alignment than static or rule-based methods.

Reinforcement Learning. Reinforcement learning (RL) has proven effective in scenarios where iterative feedback is essential for task optimization. In the realm of generative models, RL-based approaches Guo et al. (2025a); Wallace et al. (2024); Yang et al. (2024c); Guo et al. (2024); Gupta et al. (2025); Lee et al. (2024); Zhao et al. (2024); Nabati et al. (2024); Black et al. (2023); Liang et al. (2024); Kirstain et al. (2023); Lee et al. (2023); Shen et al. (2025); Zhang et al. (2024); Wang et al. (2025a); Jiang et al. (2025); Xue et al. (2025) have been employed to fine-tune outputs by maximizing a reward function that encapsulates desired attributes such as image realism, diversity, and semantic fidelity. In our framework, we adopt RL techniques to drive the automatic optimization of text prompts. By defining a multi-faceted reward that not only evaluates the visual quality of the generated image but also the interpretability and relevance of the prompt, our approach enables the model to learn an optimal prompt refinement strategy over successive iterations. Notably, T2I-R1 Jiang et al. (2025) is closely related to our work, but it targets Janus-Pro Chen et al. (2025), a unified vision-language model. In contrast, our method trains an auxiliary LLM that generalizes across diverse text-to-image models, offering a more flexible and model-agnostic solution.

Reasoning in LLM. Reasoning in LLMs Wei et al. (2022) improves complex task solving by decomposing problems into intermediate steps Feng et al. (2025); Huang et al. (2025b); Yang et al. (2025); Zhang et al. (2025a); Huang et al. (2025a); Yu et al. (2025); Ma et al. (2025b); Li et al. (2025); Lu et al. (2025). In multimodal generation, reasoning mechanisms Yang et al. (2024b); Guo et al. (2025b); Wang et al. (2025b); Zhang et al. (2025b); Sahili et al. (2024); Chen et al. (2024a) enhance prompt understanding and semantic alignment. Our method integrates reasoning with reinforcement learning, enabling step-wise prompt refinement that improves text-image alignment, and image quality—offering a new perspective for prompt optimization in T2I generation.

3 Method

We present **RePrompt**, a reasoning-augmented reprompting framework for text-to-image (T2I) generation. RePrompt decouples prompt generation from image generation, i.e., training a language model to produce structured, semantically rich prompts, while keeping the T2I backbone fixed. We optimize RePrompt via reinforcement learning (RL) to directly improve downstream image quality, compositional correctness, and usability.

3.1 Framework Overview

RePrompt comprises three main modules (Figure 2): 1) a **Prompting Policy** π_{θ} , which produces a reasoning trace H and an enhanced prompt P'; 2) a fixed **T2I Synthesizer** f_{ϕ} , which renders an image I from P'; 3) a **T2I RePrompt Reward Model** $R_{\text{total}}(I, P, P')$, which scores the image on realism, semantic alignment, and prompt structure.

Given an input prompt P, the policy samples: $y=(H,P')\sim\pi_{\theta}(P)$, the synthesizer then generates: $I=f_{\phi}(P')$. Since f_{ϕ} is non-differentiable, we formulate prompt generation as a single-step Markov Decision Process (MDP): **State**: the original prompt P. **Action**: sampling $y=(H,P')\sim\pi_{\theta}(y\mid P)$. **Transition**: deterministic mapping $P'\mapsto I=f_{\phi}(P')$. **Reward** r: the reasoning reward $r=R_{\text{total}}(I,P,P')$. **Objective**: maximize $E_{P\sim\mathcal{D}}[E_{y\sim\pi_{\theta}(y\mid P)}[r]]$.

We train π_{θ} via reinforcement learning (Group Relative Policy Optimization, GRPO Guo et al. (2025a)) to improve downstream image quality. By keeping the T2I model f_{ϕ} fixed, RePrompt learns backbone-specific reasoning strategies that enhance semantic fidelity and visual realism without requiring any manually annotated reasoning traces.

3.2 T2I REPROMPT REWARD MODEL

A central component of our framework is the **T2I RePrompt Reward Model**—an ensemble, imagegrounded reward function specifically designed for the prompt refinement task in T2I generation. In contrast to generic reward functions used in natural language or vision tasks, our reward model is co-developed with the objective of enhancing reasoning-driven prompt construction. It explicitly evaluates whether a generated prompt yields an image that is *realistic*, *semantically faithful to the user intent*, and *compositionally coherent*.



Figure 3: The Visual-Reasoning Reward.

This reward framework is critical to training Re-Prompt effectively and serves three key goals: (1)

Stable optimization: Each component provides dense and structured feedback, mitigating the challenges of sparse or noisy reward signals during early-stage learning. (2) **Multi-faceted supervision:** The reward captures complementary aspects of T2I quality, including human preference, visual quality, and semantic alignment, ensuring holistic prompt improvement. (3) **Cross-model generalization:** Because the reward depends only on the prompt-image output pair and not on any specific T2I architecture, it generalizes naturally across different generation backbones and unseen prompt distributions. Together, these properties enable our reward model to not only guide the training of the reprompting policy but also ensure broad applicability and stable learning across varying settings in text-to-image generation.

Visual-Reasoning Reward (R_{vis}). This component (as shown in Figure 3) captures both useraligned preferences and semantic correctness at the image level:

$$R_{\rm vis} = \alpha R_{\rm pref}^{\rm IMG} + \gamma R_{\rm sem}^{\rm VLM}. \tag{1}$$

Here, $R_{\rm pref}^{\rm IMG}$ is derived from ImageReward Xu et al. (2023), which used to evaluate whether generated images align with human preferences. $R_{\rm sem}^{\rm VLM}$ is obtained from VLM-Reward Achiam et al. (2023), which evaluates semantic consistency and visual quality using a vision-language model. The weights α and γ allow us to control the trade-off between perceived image quality and factual alignment.

Structure Reward ($R_{\rm struc}$). To ensure that the generated output maintains a clear reasoning-to-prompt format, we enforce a structured syntax:

We apply a binary reward:

$$R_{\text{struc}} = \begin{cases} +1, & \text{if format is correct} \\ -1, & \text{otherwise} \end{cases}$$
 (2)

This reward encourages models to adhere to a standardized output layout, simplifying downstream parsing and ensuring the interpretability of reasoning traces.

Length Reward (R_{len}). To ensure compatibility with real-world T2I models such as SDXL and FLUX.1, which impose token-length limits, we apply a constraint on prompt length:

$$R_{\rm len} = \begin{cases} +1, & L_{\rm min} \le L \le L_{\rm max}, \\ -1, & \text{otherwise} \end{cases}$$
 (3)

where L is the token length of P', and $[L_{\min}, L_{\max}]$ is empirically set to 15 and 77 tokens. This ensures the prompt is both concise and informative.

Ensemble Reward and Optimization. All reward components are normalized to unit variance and summed:

$$R_{\text{total}} = R_{\text{vis}} + R_{\text{struc}} + R_{\text{len}}.$$
 (4)

The RePrompt policy π_{θ} is trained to maximize expected downstream reward:

$$\theta^* = \arg\max_{\theta} \mathbb{E}_{P \sim \mathcal{D}, y \sim \pi_{\theta}(y|P)} \left[R_{\text{total}}(I, P, P') \right], \tag{5}$$

where $I = f_{\phi}(P')$ and f_{ϕ} is a fixed text-to-image generator.

Generalization and Flexibility. Since the reward depends only on the input-output behavior of the system, not the internals of the image model, our framework generalizes well across unseen prompts and new T2I backbones. This modularity also enables RePrompt to adapt to model-specific strengths and failure patterns, while maintaining stable and interpretable training signals throughout.

3.3 REPROMPT OPTIMIZATION

Prompt generation for T2I involves a non-differentiable, black-box image renderer f_{ϕ} . RL allows us to optimize π_{θ} directly with respect to downstream image outcomes rather than proxy losses on text. Moreover, because RePrompt is individualized to each T2I backbone, it can adapt its reasoning and prompt style to the specific strengths and limitations of a given model—improving generalization and image fidelity without retraining f_{ϕ} .

At each update step, for a given user prompt P, we sample a set of G candidate outputs $\{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(y \mid P)$. Each candidate $y_i = (H_i, P_i')$ yields an image $I_i = f_{\phi}(P_i')$ and receives a scalar reward $r_i = R_{\text{total}}(I_i, P, P')$. We then compute normalized advantages:

$$A_i = \frac{r_i - \mu_r}{\sigma_r}, \quad \mu_r = \frac{1}{G} \sum_{j=1}^G r_j, \quad \sigma_r = \sqrt{\frac{1}{G} \sum_{j=1}^G (r_j - \mu_r)^2}.$$
 (6)

We use Group Relative Policy Optimization (GRPO) Guo et al. (2025a) as a practical and stable update method for training the reprompting policy based on group-wise reward comparisons. The objective is defined as:

$$\mathcal{J}_{GRPO}(\theta) = E_{P,\{y_i\} \sim \pi_{\theta_{old}}} \left[\frac{1}{G} \sum_{i=1}^{G} \min \left(\rho_i A_i, \operatorname{clip}(\rho_i, 1 - \varepsilon, 1 + \varepsilon) A_i \right) - \beta_{KL} KL \left(\pi_{\theta}(y \mid P) \parallel \pi_{ref}(y \mid P) \right) \right], \tag{7}$$

where $\rho_i = \frac{\pi_{\theta}(y_i|P)}{\pi_{\theta_{\text{old}}}(y_i|P)}$, ε and β_{KL} are clipping and penalty coefficients, π_{ref} is a reference policy (e.g., the initial or distillation policy). By fixing f_{ϕ} and optimizing only π_{θ} , RePrompt can be

applied *universally* to any off-the-shelf T2I model, learning backbone, specific reasoning and prompt strategies without retraining the image generator. We further validate RePrompt with a variance-reduction analysis, showing that structured reasoning reduces reward uncertainty and lowers the sample complexity for accurate estimation. This leads to faster and more stable GRPO training. Full analysis and proof are provided in Appendix B.

4 EXPERIMENTS

4.1 SETTINGS

Implementation Details. We use Qwen2.5-3B Yang et al. (2024a) as our base language model. For the text-to-image model used for training, we use the FLUX.1-dev Labs (2024) model, which generates images at a resolution of 512×512 pixels. Our model is trained using the TRL ¹ reinforcement learning framework for 3 epochs, with 4 outputs generated per instance, the weight of ImageReward Xu et al. (2023) and VLM-Reward are both 0.5. The VLM used for computing VLM-Reward is GPT-4V. All experiments were conducted on 8 NVIDIA A100 (80GB) GPUs, and the entire training process required about 6 hours. More details are in the Appendix C.

Training Data. Inspired by the prompt construction strategy in GenEval Ghosh et al. (2023), we adapt six object-centric templates to a newly curated list of 288 common daily objects generated via GPT-4 Achiam et al. (2023). This results in a training corpus of 9,000 prompts, carefully filtered to avoid overlap with the GenEval. We use 8,000 prompts to fine-tune our RePrompt via supervised learning, and 1,000 prompts for reinforcement learning.

Evaluation Setup. To assess RePrompt, we evaluate on two benchmarks: GenEval Ghosh et al. (2023) and T2I-Compbench Huang et al. (2023). GenEval focuses on instance-level alignment with user intent using concise prompts, while T2I-Compbench measures compositional generation under complex scenarios involving multiple objects, attributes, and spatial relations.

4.2 COMPARISION

Table 1 presents the performance comparison across various text-to-image generation models evaluated on the GenEval benchmark, which assesses six fine-grained composition capabilities: single-object, two-object, counting, color, spatial position, and attribute binding.

Notably, our method demonstrates exceptional gains in spatial layout understanding (Position). For instance, when built upon FLUX Labs (2024), our approach achieves a 0.62 score on Position, representing a +77.1% relative improvement over the Qwen2.5 3B baseline. Similarly, for SD3 Esser et al. (2024), we observe a +78.8% gain, and for Pixart- Σ Chen et al. (2024b), the relative improvement reaches an impressive +122.2%. This strong enhancement highlights the strength of our compositional training strategy in explicitly grounding spatial relations between objects. Beyond the Position metric, our method also achieves substantial improvements in Counting (+22.2% for FLUX, +13.2% for SD3, +16.7% for Pixart- Σ) and moderate gains in attribute binding (+2.1% to +9.4%). As a result, we observe consistent boosts in overall GenEval scores: +11.8% (FLUX), +10.3% (SD3), and +6.9% (Pixart- Σ). These results verify the effectiveness of our reinforcement learning–based reprompting strategy, which enables the language model to iteratively reason about visual composition and generate more precise, image-aligned prompts, and brings state-of-the-art advances in spatial understanding.

4.3 GENERALIZATION PERFORMANCE ON T2I-COMPBENCH

We further assess the robustness of our method on the T2I-Compbench benchmark, which tests compositional generalization across six dimensions: color, shape, texture, spatial reasoning, numeracy, and complex attribute combinations. As shown in Table 2, RePrompt consistently improves performance across all evaluated backbones. In particular, it significantly boosts spatial compositional scores (e.g., from 0.2494 to 0.3301 on FLUX and from 0.2815 to 0.3315 on SD3) and enhances numeracy understanding, two long-standing challenges in text-to-image generation. Moreover, our method outperforms stronger LLM-enhanced baselines, for instance, on Pixart- Σ ,

¹https://github.com/huggingface/trl

Table 1: Evaluation of text-to-image generation on the GenEval benchmark. Our method consistently outperforms strong baselines, achieving the best overall scores. Notably, our approach shows substantial gains in spatial position understanding over Qwen2.5 3B-enhanced baselines, demonstrating its superior capability in grounding spatial relations.

Method	Single	Two	Counting	Colors	Position	Attribute	Overall ↑
	object	object				binding	
FLUX Labs (2024)	0.99	0.79	0.75	0.78	0.18	0.45	0.66
+ Promptist Hao et al. (2023)	0.98	0.72	0.70	0.78	0.21	0.44	0.66
+ PAG Yun et al. (2025)	0.97	0.74	0.73	0.80	0.36	0.46	0.69
+ GPT4	0.99	0.79	0.68	0.84	0.51	0.52	0.72
+ Deepseek-r1	1.00	0.81	0.56	0.78	0.47	0.43	0.67
+ Qwen2.5 7B	0.99	0.83	0.62	0.84	0.36	0.51	0.69
+ Qwen2.5 3B	0.99	0.84	0.63	0.81	0.35	0.48	0.68
+ Ours (train w/ FLUX)	0.98	0.87	0.77	0.85	0.62	0.49	0.76
Improvement over Qwen2.5 3B	-1.0%	+3.6%	+22.2%	_+4.9%	_+77.1%	+2.1%	+17.8%
SD3 Esser et al. (2024)	1.00	0.85	0.62	0.88	0.22	0.58	0.69
+ Promptist	0.99	0.84	0.66	0.84	0.45	0.52	0.69
+ PAG	0.99	0.85	0.68	0.85	0.49	0.53	0.71
+ GPT4	1.00	0.84	0.51	0.85	0.48	0.54	0.70
+ Deepseek-r1	0.99	0.82	0.53	0.80	0.44	0.46	0.67
+ Qwen2.5 7B	1.00	0.82	0.49	0.85	0.34	0.58	0.68
+ Qwen2.5 3B	1.00	0.86	0.53	0.84	0.33	0.55	0.68
+ Ours (train w/ FLUX)	0.99	0.86	0.60	0.86	0.59	0.60	0.75
Improvement over Qwen2.5 3B	-1.0%	0.0%	+13.2%	+2.4%	+78.8%	+9.1%	+10.3%
Pixart- Σ Chen et al. (2024b)	0.99	0.60	0.47	0.81	0.10	0.26	0.54
+ Promptist	0.98	0.60	0.49	0.80	0.20	0.27	0.55
+ PAG	0.98	0.63	0.52	0.80	0.28	0.29	0.56
+ GPT4	0.96	0.67	0.48	0.84	0.36	0.31	0.60
+ Deepseek-r1	0.99	0.63	0.43	0.78	0.24	0.27	0.56
+ Qwen2.5 7B	0.96	0.67	0.43	0.83	0.20	0.32	0.57
+ Qwen2.5 3B	0.99	0.68	0.48	0.82	0.18	0.32	0.58
+ Ours (train w/ FLUX)	0.98	0.64	0.56	0.81	0.40	0.35	0.62
Improvement over Qwen2.5 3B	-1.0%	-5.9%	+16.7%	-1.2%	_+ <u>1</u> 2 <u>2</u> . <u>2</u> <u>%</u>	+9.4%	+6.9%

Table 2: Evaluation of text-to-image generation on the T2I-Compbench. We report the baseline results, their variants enhanced with Qwen2.5 3B, and our method trained with FLUX. Our approach consistently improves performance across most aspects, particularly in Spatial compositions.

Method	Color	Shape	Texture	Spatial	Numeracy	Complex
FLUX	0.7223	0.4796	0.6522	0.2494	0.6101	0.3616
+ Qwen2.5 3B	0.7149	0.5103	0.6012	0.2579	0.5982	0.3325
+ Ours (train w/ FLUX)	0.7501	0.5276	0.6515	0.3301	0.6499	0.3721
SD3	0.7941	0.5812	0.7224	0.2815	0.5871	0.3714
+ Qwen2.5 3B	0.7227	0.5478	0.6581	0.2549	0.5934	0.3307
+ Ours (train w/ FLUX)	0.7866	0.5891	0.7184	0.3315	0.6289	0.3744
Pixart-Σ	0.5682	0.4717	0.5622	0.2497	0.5366	0.3655
+ Qwen2.5 3B	0.6618	0.4814	0.5662	0.2481	0.5443	0.3335
+ Ours (train w/ FLUX)	0.6665	0.5011	0.6190	0.2913	0.5716	0.3680

RePrompt achieves notable gains in both texture and spatial dimensions. These results demonstrate that our approach generalizes well across diverse models and compositional skills, validating its effectiveness as a versatile and plug-and-play enhancement for real-world T2I systems.

4.4 ABLATION STUDY

Ablation on SFT and RL. Table 3 shows the effect of supervised fine-tuning (SFT) and reinforcement learning (RL) on GenEval. SFT brings modest gains (+0.01 overall), mainly improving positional understanding ($0.35 \rightarrow 0.43$), suggesting it helps inject object-attribute knowledge but strug-

Table 3: Ablation study of SFT and RL on the GenEval benchmark.

Method	Single object	Two object	Counting	Colors	Position	Attribute binding	Overall ↑
FLUX Labs (2024)	0.99	0.79	0.75	0.78	0.18	0.45	0.66
+ Qwen2.5 3B	0.99	0.84	0.63	0.81	0.35	0.48	0.68
w/ SFT	0.99	0.83	0.64	0.81	0.43	0.44	0.69
w/ RL	0.98	0.83	0.71	0.87	0.41	0.53	0.72
w/ SFT + RL	0.98	0.87	0.77	0.85	0.62	0.49	0.76

Table 4: Quantitative comparison between RePrompt and other method on image generation accuracy and latency with the subset of Geneval. All latency is measured on a single NVIDIA A100 GPU.

Method	Accuracy ↑	Latency (s / img) ↓
FLUX	0.65	20
Show-o Xie et al. (2024)	0.55	3
Idea2Img Yang et al. (2024d) (w/ FLUX)	0.69	140
PARM++ Guo et al. (2025b) (w/ Show-o)	0.72	110
RePrompt (w/ FLUX)	0.76	30

gles with complex reasoning. RL yields larger boosts (+0.04 overall), as it directly optimizes visual correctness. Combining SFT and RL achieves the best results (0.76 overall), with strong improvements in spatial reasoning (0.62) and counting (0.77). These results confirm that SFT offers useful priors, while RL is key for compositional robustness.

Comparison on Accuracy and Latency. Table 4 benchmarks our RePrompt against prior methods in terms of both generation accuracy (GenEval overall) and inference latency. Among base models, FLUX.1 achieves higher accuracy compared to Show-o but suffers from a longer latency (20s vs. 3s per image), reflecting a trade-off between quality and speed. When incorporating advanced prompting techniques, Idea2Img (w/ FLUX.1) improves accuracy to 0.69 but at the cost of a significant latency increase (140s per image), while PARM++ (w/ Show-o) achieves 0.72 accuracy with 110s latency. In contrast, RePrompt achieves the best accuracy while maintaining a much lower latency. This demonstrates the effectiveness of RePrompt in enabling precise visual grounding without compromising efficiency, making it more practical for real-world deployment.

Ablation Study on Trained T2I Models. Figure 4 illustrates the generalization ability of RePrompt across three diverse T2I backbones: FLUX, SD3, and PixArt- Σ . Our method consistently boosts compositional scores on GenEval—improving FLUX from 0.66 to 0.76, SD3 from 0.69 to 0.75, and PixArt- Σ from 0.54 to 0.62—outperforming both baselines and the +Qwen2.5 variants. These results demonstrate RePrompt's robustness and model-agnostic design, working effectively across both large (e.g., SD3) and lightweight (e.g., PixArt- Σ) backbones. Notably, we observe that stronger T2I models used during training lead to better generalization at inference, likely due to richer reasoning patterns induced during learning.

Qualitative Comparison. Figure 5 illustrates the qualitative advantage of RePrompt over existing T2I models. Baseline models often generate images with incorrect spatial relations or hallucinated objects. For example, when prompted with "a fire hydrant with a tennis racket," DALL-E 3 produces unrealistic, stylistic blends where the objects are merged. In contrast, RePrompt accurately grounds the tennis racket "around the base" of the fire hydrant, respecting the intended composition. Similarly, for "a photo of a dog above a cow," our method correctly depicts the dog in the air above the cow with "a high angle shot", aligning with the prompt semantics. These results highlight RePrompt's ability to mitigate typical failures in spatial arrangement and object interaction by generating prompts with explicit compositional cues. More cases are in the Appendix E.

More Ablation Study. We present additional ablation studies on reasoning, training dynamics, and reward functions in the Appendix D to further validate the effectiveness of our reasoning design and reward selection, as well as the stability of the training process.

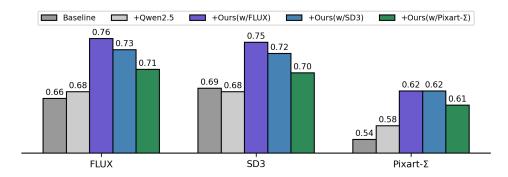


Figure 4: Impact of our method across different base T2I models on the GenEval benchmark. Our method consistently improves the compositional understanding across all base models.

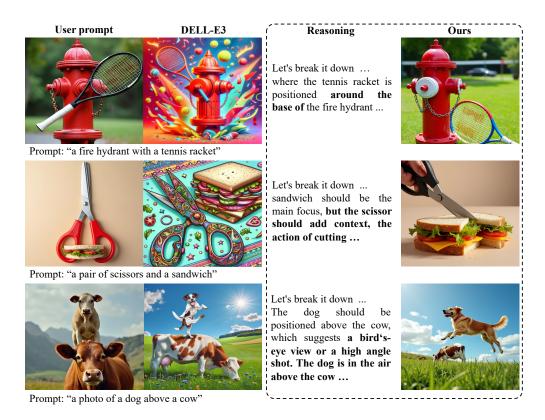


Figure 5: Qualitative results on compositional prompts. Compared to vanilla T2I models, our Re-Prompt improves spatial layout and object relations by generating enhanced prompts with explicit reasoning, leading to more faithful compositions.

5 Conclusion

In this work, we propose a simple yet effective method to enhance compositional understanding in text-to-image (T2I) generation models. By injecting chain-of-thought (CoT) reasoning into the prompt construction pipeline and pairing each CoT with an enhanced prompt, our approach improves the alignment between textual descriptions and generated images. Extensive experiments on the GenEval benchmark demonstrate that our method consistently improves performance across various T2I backbones, including FLUX, SD3, and Pixart- Σ , without requiring additional model retraining. These results highlight the generalizability and plug-and-play nature of our method. We believe our approach offers a new perspective for improving controllability and compositional fidelity in generative models.

ETHICS STATEMENT

We have adhered to the ICLR Code of Ethics throughout the research and submission process. This work does not involve human subjects or sensitive data, and no potentially harmful insights or applications are presented. We have ensured that our dataset usage complies with all relevant licenses and ethical guidelines. The authors declare no conflicts of interest or external sponsorship that could have influenced the research outcomes. We are committed to fairness and reducing bias, and have taken care to avoid discriminatory or unethical practices. Privacy and security considerations have been respected in all stages of this work. We affirm the integrity of our research process and confirm that all ethical standards and institutional requirements have been met.

7 REPRODUCIBILITY STATEMENT

We have taken several steps to ensure the reproducibility of our results. Detailed descriptions of the model architecture, training procedures, and experimental settings are provided in the main text and supplementary materials. All assumptions and derivations underlying our theoretical claims are clearly explained and fully proven in the appendix. The datasets used in our experiments are publicly available. Additionally, we provide anonymized source code as supplementary material to facilitate replication of our experiments. We encourage readers to refer to these materials for a complete understanding of our methodology and experimental setup.

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, et al. Improving image generation with better captions. *Computer Science. https://cdn. openai. com/papers/dall-e-3. pdf*, 2(3):8, 2023.
- Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models with reinforcement learning. *arXiv preprint arXiv:2305.13301*, 2023.
- Tingfeng Cao, Chengyu Wang, Bingyan Liu, Ziheng Wu, Jinhui Zhu, and Jun Huang. Beauti-fulprompt: Towards automatic prompt engineering for text-to-image synthesis. *arXiv* preprint arXiv:2311.06752, 2023.
- Anthony Chen, Jianjin Xu, Wenzhao Zheng, Gaole Dai, Yida Wang, Renrui Zhang, Haofan Wang, and Shanghang Zhang. Training-free regional prompting for diffusion transformers. *arXiv* preprint arXiv:2411.02395, 2024a.
- Junsong Chen, Chongjian Ge, Enze Xie, Yue Wu, Lewei Yao, Xiaozhe Ren, Zhongdao Wang, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart-σ: Weak-to-strong training of diffusion transformer for 4k text-to-image generation. In *European Conference on Computer Vision*, pp. 74–91. Springer, 2024b.
- 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.
- 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.
- Kaituo Feng, Kaixiong Gong, Bohao Li, Zonghao Guo, Yibing Wang, Tianshuo Peng, Benyou Wang, and Xiangyu Yue. Video-r1: Reinforcing video reasoning in mllms. *arXiv preprint arXiv:2503.21776*, 2025.

- Dhruba 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.
 - Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025a.
 - Jianshu Guo, Wenhao Chai, Jie Deng, Hsiang-Wei Huang, Tian Ye, Yichen Xu, Jiawei Zhang, Jenq-Neng Hwang, and Gaoang Wang. Versat2i: Improving text-to-image models with versatile reward. *arXiv preprint arXiv:2403.18493*, 2024.
 - Ziyu Guo, Renrui Zhang, Chengzhuo Tong, Zhizheng Zhao, Peng Gao, Hongsheng Li, and Pheng-Ann Heng. Can we generate images with cot? let's verify and reinforce image generation step by step. *arXiv preprint arXiv:2501.13926*, 2025b.
 - Shashank Gupta, Chaitanya Ahuja, Tsung-Yu Lin, Sreya Dutta Roy, Harrie Oosterhuis, Maarten de Rijke, and Satya Narayan Shukla. A simple and effective reinforcement learning method for text-to-image diffusion fine-tuning. *arXiv preprint arXiv:2503.00897*, 2025.
 - Yaru Hao, Zewen Chi, Li Dong, and Furu Wei. Optimizing prompts for text-to-image generation. *Advances in Neural Information Processing Systems*, 36:66923–66939, 2023.
 - Kaiyi Huang, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation. *Advances in Neural Information Processing Systems*, 36:78723–78747, 2023.
 - Qihan Huang, Long Chan, Jinlong Liu, Wanggui He, Hao Jiang, Mingli Song, Jingyuan Chen, Chang Yao, and Jie Song. Boosting mllm reasoning with text-debiased hint-grpo. *arXiv preprint arXiv:2503.23905*, 2025a.
 - Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Yao Hu, and Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *arXiv* preprint arXiv:2503.06749, 2025b.
 - Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv* preprint arXiv:2412.16720, 2024.
 - Dongzhi Jiang, Ziyu Guo, Renrui Zhang, Zhuofan Zong, Hao Li, Le Zhuo, Shilin Yan, Pheng-Ann Heng, and Hongsheng Li. T2i-r1: Reinforcing image generation with collaborative semantic-level and token-level cot. *arXiv preprint arXiv:2505.00703*, 2025.
 - Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Picka-pic: An open dataset of user preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36:36652–36663, 2023.
 - Black Forest Labs. Flux. https://github.com/black-forest-labs/flux, 2024.
 - Kimin Lee, Hao Liu, Moonkyung Ryu, Olivia Watkins, Yuqing Du, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, and Shixiang Shane Gu. Aligning text-to-image models using human feedback. *arXiv preprint arXiv:2302.12192*, 2023.
- Seung Hyun Lee, Yinxiao Li, Junjie Ke, Innfarn Yoo, Han Zhang, Jiahui Yu, Qifei Wang, Fei Deng, Glenn Entis, Junfeng He, et al. Parrot: Pareto-optimal multi-reward reinforcement learning framework for text-to-image generation. In *European Conference on Computer Vision*, pp. 462–478. Springer, 2024.
 - Xinhao Li, Ziang Yan, Desen Meng, Lu Dong, Xiangyu Zeng, Yinan He, Yali Wang, Yu Qiao, Yi Wang, and Limin Wang. Videochat-r1: Enhancing spatio-temporal perception via reinforcement fine-tuning. *arXiv preprint arXiv:2504.06958*, 2025.

- Youwei Liang, Junfeng He, Gang Li, Peizhao Li, Arseniy Klimovskiy, Nicholas Carolan, Jiao Sun,
 Jordi Pont-Tuset, Sarah Young, Feng Yang, et al. Rich human feedback for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
 pp. 19401–19411, 2024.
 - Bingchen Liu, Ehsan Akhgari, Alexander Visheratin, Aleks Kamko, Linmiao Xu, Shivam Shrirao, Chase Lambert, Joao Souza, Suhail Doshi, and Daiqing Li. Playground v3: Improving text-to-image alignment with deep-fusion large language models. *arXiv preprint arXiv:2409.10695*, 2024.
 - Zhengxi Lu, Yuxiang Chai, Yaxuan Guo, Xi Yin, Liang Liu, Hao Wang, Guanjing Xiong, and Hongsheng Li. Ui-r1: Enhancing action prediction of gui agents by reinforcement learning. *arXiv* preprint arXiv:2503.21620, 2025.
 - Nanye Ma, Shangyuan Tong, Haolin Jia, Hexiang Hu, Yu-Chuan Su, Mingda Zhang, Xuan Yang, Yandong Li, Tommi Jaakkola, Xuhui Jia, et al. Inference-time scaling for diffusion models beyond scaling denoising steps. *arXiv preprint arXiv:2501.09732*, 2025a.
 - Xinyu Ma, Ziyang Ding, Zhicong Luo, Chi Chen, Zonghao Guo, Derek F Wong, Xiaoyi Feng, and Maosong Sun. Deepperception: Advancing r1-like cognitive visual perception in mllms for knowledge-intensive visual grounding. *arXiv preprint arXiv:2503.12797*, 2025b.
 - Oscar Mañas, Pietro Astolfi, Melissa Hall, Candace Ross, Jack Urbanek, Adina Williams, Aishwarya Agrawal, Adriana Romero-Soriano, and Michal Drozdzal. Improving text-to-image consistency via automatic prompt optimization. *arXiv* preprint arXiv:2403.17804, 2024.
 - Wenyi Mo, Tianyu Zhang, Yalong Bai, Bing Su, Ji-Rong Wen, and Qing Yang. Dynamic prompt optimizing for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26627–26636, 2024.
 - Ofir Nabati, Guy Tennenholtz, ChihWei Hsu, Moonkyung Ryu, Deepak Ramachandran, Yinlam Chow, Xiang Li, and Craig Boutilier. Personalized and sequential text-to-image generation. *arXiv* preprint arXiv:2412.10419, 2024.
 - Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023.
 - Jie Qin, Jie Wu, Weifeng Chen, Yuxi Ren, Huixia Li, Hefeng Wu, Xuefeng Xiao, Rui Wang, and Shilei Wen. Diffusiongpt: Llm-driven text-to-image generation system. *arXiv* preprint *arXiv*:2401.10061, 2024.
 - 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*, pp. 10684–10695, 2022.
 - 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.
 - Zahraa Al Sahili, Ioannis Patras, and Matthew Purver. Faircot: Enhancing fairness in diffusion models via chain of thought reasoning of multimodal language models. *arXiv preprint arXiv:2406.09070*, 2024.
 - Haozhan Shen, Peng Liu, Jingcheng Li, Chunxin Fang, Yibo Ma, Jiajia Liao, Qiaoli Shen, Zilun Zhang, Kangjia Zhao, Qianqian Zhang, et al. Vlm-r1: A stable and generalizable r1-style large vision-language model. *arXiv preprint arXiv:2504.07615*, 2025.
 - Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*, 2025.

- Luong Trung, Xinbo Zhang, Zhanming Jie, Peng Sun, Xiaoran Jin, and Hang Li. Reft: Reasoning with reinforced fine-tuning. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 7601–7614, 2024.
- 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*, pp. 8228–8238, 2024.
- Junke Wang, Zhi Tian, Xun Wang, Xinyu Zhang, Weilin Huang, Zuxuan Wu, and Yu-Gang Jiang. Simplear: Pushing the frontier of autoregressive visual generation through pretraining, sft, and rl. *arXiv preprint arXiv:2504.11455*, 2025a.
- Yi Wang, Mushui Liu, Wanggui He, Longxiang Zhang, Ziwei Huang, Guanghao Zhang, Fangxun Shu, Zhong Tao, Dong She, Zhelun Yu, et al. Mint: Multi-modal chain of thought in unified generative models for enhanced image generation. *arXiv preprint arXiv:2503.01298*, 2025b.
- Zhenyu Wang, Aoxue Li, Zhenguo Li, and Xihui Liu. Genartist: Multimodal Ilm as an agent for unified image generation and editing. *Advances in Neural Information Processing Systems*, 37: 128374–128395, 2024.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- Tsung-Han Wu, Long Lian, Joseph E Gonzalez, Boyi Li, and Trevor Darrell. Self-correcting llm-controlled diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6327–6336, 2024.
- Enze Xie, Junsong Chen, Yuyang Zhao, Jincheng Yu, Ligeng Zhu, Yujun Lin, Zhekai Zhang, Muyang Li, Junyu Chen, Han Cai, et al. Sana 1.5: Efficient scaling of training-time and inference-time compute in linear diffusion transformer. *arXiv preprint arXiv:2501.18427*, 2025.
- 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.
- 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.
- Zeyue Xue, Jie Wu, Yu Gao, Fangyuan Kong, Lingting Zhu, Mengzhao Chen, Zhiheng Liu, Wei Liu, Qiushan Guo, Weilin Huang, et al. Dancegrpo: Unleashing grpo on visual generation. *arXiv* preprint arXiv:2505.07818, 2025.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*, 2024a.
- Ling Yang, Zhaochen Yu, Chenlin Meng, Minkai Xu, Stefano Ermon, and Bin Cui. Mastering text-to-image diffusion: Recaptioning, planning, and generating with multimodal llms. In *Forty-first International Conference on Machine Learning*, 2024b.
- Shentao Yang, Tianqi Chen, and Mingyuan Zhou. A dense reward view on aligning text-to-image diffusion with preference. *arXiv preprint arXiv:2402.08265*, 2024c.
- Yi Yang, Xiaoxuan He, Hongkun Pan, Xiyan Jiang, Yan Deng, Xingtao Yang, Haoyu Lu, Dacheng Yin, Fengyun Rao, Minfeng Zhu, et al. R1-onevision: Advancing generalized multimodal reasoning through cross-modal formalization. *arXiv preprint arXiv:2503.10615*, 2025.
- Zhengyuan Yang, Jianfeng Wang, Linjie Li, Kevin Lin, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. Idea2img: Iterative self-refinement with gpt-4v for automatic image design and generation. In *European Conference on Computer Vision*, pp. 167–184. Springer, 2024d.

- Shih-Ying Yeh, Sang-Hyun Park, Giyeong Oh, Min Song, and Youngjae Yu. Tipo: Text to image with text presampling for prompt optimization. *arXiv preprint arXiv:2411.08127*, 2024.
- En Yu, Kangheng Lin, Liang Zhao, Jisheng Yin, Yana Wei, Yuang Peng, Haoran Wei, Jianjian Sun, Chunrui Han, Zheng Ge, et al. Perception-r1: Pioneering perception policy with reinforcement learning. *arXiv* preprint arXiv:2504.07954, 2025.
- Taeyoung Yun, Dinghuai Zhang, Jinkyoo Park, and Ling Pan. Learning to sample effective and diverse prompts for text-to-image generation. *arXiv preprint arXiv:2502.11477*, 2025.
- Jingyi Zhang, Jiaxing Huang, Huanjin Yao, Shunyu Liu, Xikun Zhang, Shijian Lu, and Dacheng Tao. R1-vl: Learning to reason with multimodal large language models via step-wise group relative policy optimization. *arXiv* preprint arXiv:2503.12937, 2025a.
- Xinchen Zhang, Ling Yang, Guohao Li, Yaqi Cai, Jiake Xie, Yong Tang, Yujiu Yang, Mengdi Wang, and Bin Cui. Iteracomp: Iterative composition-aware feedback learning from model gallery for text-to-image generation. *arXiv preprint arXiv:2410.07171*, 2024.
- Yuyao Zhang, Jinghao Li, and Yu-Wing Tai. Layercraft: Enhancing text-to-image generation with cot reasoning and layered object integration. *arXiv* preprint arXiv:2504.00010, 2025b.
- Hanyang Zhao, Haoxian Chen, Ji Zhang, David D Yao, and Wenpin Tang. Scores as actions: a framework of fine-tuning diffusion models by continuous-time reinforcement learning. *arXiv* preprint arXiv:2409.08400, 2024.

APPENDIX

A BROADER IMPACT

Our work aims to enhance the controllability and fidelity of text-to-image generation by aligning prompt engineering with user intent through self-reflective reasoning. This could have a wide range of beneficial applications, such as improving accessibility tools for the visually impaired, assisting designers in rapid prototyping, or helping educators generate custom visual teaching materials. By reducing the gap between user intentions and generated images, our method empowers non-expert users to communicate complex visual ideas more effectively.

However, like many generative technologies, our system also presents potential risks. Enhanced prompts with richer semantics could inadvertently be exploited to generate more realistic harmful or misleading content, especially in the context of misinformation or deepfakes. To mitigate such risks, we recommend pairing our method with content moderation filters and safety alignment mechanisms in deployment. Additionally, since our method involves reinforcement learning with learned reward models, any bias in the reward signal may propagate into the prompt generation. Future research should explore fairness-aware training objectives and human-in-the-loop evaluation to ensure responsible use.

Overall, we believe that RePrompt contributes to the development of more interpretable, controllable, and user-aligned generative models, which are essential for trustworthy AI applications.

B VARIANCE REDUCTION VIA STRUCTURED REASONING

In this appendix, we provide a full theoretical analysis showing that conditioning prompt generation on an explicit reasoning trace H strictly reduces the variance of the downstream reward estimator. This reduction in variance directly leads to improved sample efficiency for reinforcement learning.

B.1 SETUP AND NOTATION

Let:

- \mathcal{P} be the space of bare prompts P'.
- \mathcal{H} be the space of reasoning traces H.
- $r: \mathcal{P} \to [0, R_{\text{max}}]$ be the reward random variable obtained by sampling $P' \sim \pi(P)$ and generating an image $I = f_{\phi}(P')$.
- $r_H: \mathcal{H} \times \mathcal{P} \to [0, R_{\max}]$ be the reward when first sampling $H \sim \pi_H(P)$, then $P' \sim \pi(P \mid H)$, and finally $I = f_{\phi}(P')$.
- All expectations and variances are taken over the joint sampling of H and P'.

B.2 LAW OF TOTAL VARIANCE

By the law of total variance,

$$\operatorname{Var}[r(P')] = E_H[\operatorname{Var}[r(P') \mid H]] + \operatorname{Var}_H[E[r(P') \mid H]]$$

Since variances are nonnegative, we immediately have:

Theorem B.1 (Variance Reduction).

$$\operatorname{Var}[r(H, P')] = E_H[\operatorname{Var}[r \mid H]] \leq \operatorname{Var}[r(P')].$$

Proof. By definition,

$$\operatorname{Var}[r(H, P')] = E_H[\operatorname{Var}[r \mid H]]$$

and since

$$Var[r(P')] = E_H[Var[r \mid H]] + Var_H[E[r \mid H]],$$

the nonnegativity of $Var_H[E[r \mid H]]$ yields the result.

B.3 SAMPLE COMPLEXITY IMPROVEMENT

Variance directly impacts the number of samples required to estimate the expected reward within a given accuracy. Consider estimating the expected reward $\mu = E[r]$ by drawing N independent samples $\{r_i\}$. By Chebyshev's inequality,

$$\Pr(|\hat{\mu} - \mu| \ge \varepsilon) \le \frac{\operatorname{Var}[r]}{N\varepsilon^2}.$$

Thus, to guarantee $\Pr(|\hat{\mu} - \mu| \ge \varepsilon) \le \delta$, we require

$$N \geq \frac{\operatorname{Var}[r]}{\varepsilon^2 \delta}.$$

Applying Theorem B.1, using reasoning reduces the required sample size:

$$N_{\text{reasoning}} = \frac{\operatorname{Var}[r(H, P')]}{\varepsilon^2 \delta} \le \frac{\operatorname{Var}[r(P')]}{\varepsilon^2 \delta} = N_{\text{bare}}.$$

B.4 DISCUSSION

This analysis formalizes the intuition that introducing an intermediate reasoning trace H conditions the policy on structured, compositional information about the scene, thereby reducing uncertainty (variance) in the reward signal. Empirically, this translates to fewer rollouts needed during GRPO training and more stable gradient estimates—accelerating convergence without additional data or annotations.

C EXPERIMENT SETTING

Table 5 summarizes the key hyperparameters used in our experiments, including configurations for the GRPO optimization algorithm, the FLUX.1 text-to-image model, and the joint training process. For GRPO, we set the clipping range ε to 0.2 and the KL penalty coefficient β to 0.04 with a group size of 4. The FLUX.1 model operates at a resolution of 512×512 , with 50 diffusion steps and a classifier-free guidance (CFG) scale of 3.5. In joint training, we balance ImageReward and VLM-Reward with equal weights (0.5 each), and constrain prompt lengths between 15 and 77 tokens. Training is conducted using 8 devices with a per-device batch size of 4, a learning rate of 2e-6, gradient accumulation steps of 2, and a total of 3 epochs.

D ABLATION STUDY

Ablation Study on Reasoning. Table 6 presents an ablation study on the GenEval benchmark to assess the impact of reasoning in our reinforcement learning framework. The FLUX baseline, integrating a large language model (+Qwen2.5 3B) brings modest gains across most categories, raising the overall score to 0.68. Applying RL without reasoning achieves a similar overall improvement (0.68), suggesting that reward-driven optimization alone contributes to better alignment. However, incorporating reasoning into the RL loop leads to a more substantial improvement, pushing the overall score to 0.72. Notably, categories that require more complex semantic understanding—such as "Colors" (from 0.83 to 0.87) and "Attribute binding" (from 0.46 to 0.53)—benefit the most. These results demonstrate that step-by-step reasoning helps the model better decompose and interpret textual prompts, thereby enabling more accurate and faithful image generation.

Training Dynamics. Figure 6 presents the reward curve during reinforcement learning of the Re-Prompt. We observe a stable and monotonically increasing trend in the reward, demonstrating that our reward model provides a reliable and effective supervision signal throughout training. The absence of sharp fluctuations or reward collapse suggests that our RL setup maintains stable policy updates. This aligns with the observed downstream improvements, confirming that reward-guided prompt refinement effectively enhances compositional alignment in generated images.

Table 5: Key Parameters for GRPO and T2I Model.

Config	Symbol	Value							
GRPO Config									
Clipping range	ε	0.2							
KL penalty coefficient	β	0.04							
Group size	G	4							
FLUX.1 Config									
Image resolution	$H \times W$	512×512							
Diffusion steps	T	50							
CFG scale	$\lambda_{ m cfg}$	3.5							
Joint Trainin	Joint Training Parameters								
ImageReward weight	α	0.5							
VLM-Reward weight	γ	0.5							
Min prompt's length	L_{\min}	15							
Max prompt's length	$L_{ m max}$	77							
Device number	-	8							

Table 6: Ablation study of reasoning on the GenEval benchmark.

B

lr

2e-6

Per device batch size

Gradient Accumulation

Learning rate

Epoch

Method	Single object	Two object	Counting	Colors	Position	Attribute binding	Overall ↑
FLUX	0.99	0.79	0.75	0.78	0.18	0.45	0.66
+Qwen2.5 3B	0.99	0.84	0.63	0.81	0.35	0.48	0.68
RL w/o reasoning	1.00	0.81	0.68	0.83	0.33	0.46	0.68
RL w/ reasoning	0.98	0.83	0.71	0.87	0.41	0.53	0.72

Ablation Study on Visual Reasoning Rewards. To investigate the effectiveness of different visual reasoning reward signals used in our reinforcement learning-based reprompting strategy, we conduct a detailed ablation study on the GenEval benchmark. As shown in Table 7, our method consistently improves over the FLUX baseline and Qwen2.5 3B variant across all subcategories, demonstrating the efficacy of reward-driven learning in aligning generated images with prompt semantics.

Both ImageReward and VLM-Reward show noticeable gains over the +Qwen2.5 3B baseline, indicating that each reward captures complementary aspects of visual faithfulness. However, their standalone performances are still limited in challenging tasks like Position. Notably, our ensemble reward formulation—which combines both ImageReward and VLM-based feedback—achieves the best overall performance, with consistent improvements compared to the +Qwen2.5 3B baseline. These are precisely the categories that require stronger compositional understanding and spatial reasoning, underscoring the strength of our composite reward design in driving semantically aligned image generation. The ablation confirms that each reward contributes uniquely, and their integration enables more holistic supervision, leading to substantial gains in compositional image-text alignment.

Ablation Study on VLM-Reward Models. To evaluate the generality and accessibility of our framework, we replaced GPT-4V with open-source vision-language models (Qwen-VL 72B and 32B) as the reward model. Results show that RePrompt still achieves competitive performance: 0.74 with Qwen-VL 72B and 0.73 with Qwen-VL 32B (vs. 0.76 with GPT-4V), confirming the

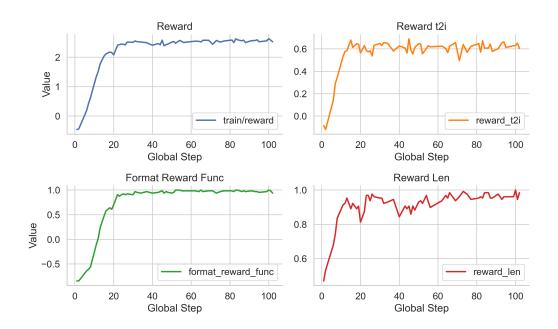


Figure 6: Training curve of our RePrompt during reinforcement learning. The curve shows steady reward improvement, indicating stable training dynamics and effective reward signal.

Table 7: Ablation study of rewards on the GenEval benchmark.

Method	α	γ	Single object	Two object	Counting	Colors	Position	Attribute binding	Overall ↑	
FLUX +Qwen2.5 3B	-	-	0.99 0.99	0.79 0.84	0.75 0.63	0.78 0.81	0.18 0.35	0.45 0.48	0.66 0.68	
RePrompt	1.0 0.7 0.5 0.3 0.0	0.0 0.3 0.5 0.7 1.0	0.99 0.98 0.98 0.99 0.99	0.86 0.86 0.87 0.85 0.85	0.66 0.70 0.77 0.74 0.72	0.89 0.88 0.85 0.84 0.83	0.52 0.56 0.62 0.58 0.53	0.50 0.52 0.49 0.50 0.47	0.74 0.75 0.76 0.75 0.74	

method's robustness without reliance on proprietary tools. We will include full results and adoption with accessible reward models in the final version.

Table 8: Performance comparison of RePrompt with different VLM-Reward models

Method	Single	Two	Counting	Colors	Position	Attribute	Overall ↑
+ GPT-4V	0.98	0.87	0.77	0.85	0.62	0.49	0.76
+ Qwen2.5-VL 72B	0.98	0.86	0.75	0.84	0.58	0.46	0.74
+ Qwen2.5-VL 32B	0.98	0.85	0.74	0.83	0.55	0.45	0.73

Training Dynamics with Different Visual Reasoning Rewards. We analyze the impact of different visual reasoning rewards on training dynamics by plotting training curves under various reward configurations. As shown in Figure 7, training with a single reward model leads to unstable reward fluctuations and degraded sample quality. This instability arises from the limited supervision capacity of a single reward model, which may overfit to specific patterns or neglect important compositional aspects. In contrast, our proposed multi-reward formulation—balancing diverse reasoning signals, enables smoother optimization and more consistent improvements across iterations. These results emphasize the importance of combining complementary reward models to achieve both stability and generalization.

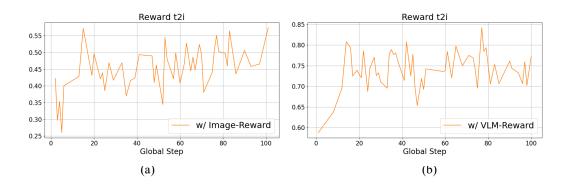


Figure 7: Training curve with different visual reasoning rewards. Using a single reward model leads to instability and suboptimal performance during training. In contrast, our balanced reward design—combining multiple specialized reward signals—yields more stable convergence and higher final reward values.

E MORE CASES

 More Qualitative Comparison. Figure 8, 9, 10 and 11 present qualitative comparisons on different types of compositional prompts, including Position, Two-object, Color, and Attribute Binding. Compared to baseline methods such as DELL-E3, our approach produces images that more faithfully adhere to the spatial, numerical, and attribute-based constraints specified in the prompts. For instance, in Figure 7, our method accurately grounds relative positions (e.g., "above", "left of") by leveraging explicit reasoning, while DELL-E3 often fails to reflect such relations or misplaces objects entirely. Figure 8 highlights our advantage in handling prompts involving multiple objects, where baseline models tend to merge objects or hallucinate irrelevant content. Similarly, in Figure 9, our method is better at preserving specified colors for each object, whereas the baseline sometimes misbinds colors or applies them inconsistently. Finally, as shown in Figure 10, our approach improves attribute binding, ensuring that each attribute is applied to the correct object without confusion. These results demonstrate the effectiveness of integrating reasoning to improve alignment between visual outputs and complex prompt semantics.

F LIMITATIONS

While our method consistently improves compositional quality across standard T2I benchmarks, several limitations warrant future exploration. First, the performance gains on certain fine-grained tasks—such as numeracy and object counting—remain modest, suggesting potential for further enhancement in precise quantitative reasoning. Second, our approach requires fine-tuning with compositional supervision, which may introduce additional computational cost and reliance on structured training signals. However, this design is consistent with common RLHF setups and does not limit practical deployment. Third, the effectiveness of RePrompt depends on the quality of the reward model; while we demonstrate robustness across standard evaluators, improvements in reward fidelity could further amplify performance.

G LLM USAGE

In this work, large language models (LLMs) were utilized primarily as a general-purpose assistive tool to help with text polishing and improving writing clarity. The core research ideation, experimental design, analysis, and conclusions were developed solely by the authors. The LLM did not contribute to generating any original scientific content or ideas. The authors take full responsibility for the contents of the paper, including all text generated or revised with the assistance of LLMs.



Prompt: "a photo of a cell phone left of a tennis racket'



Prompt: "a photo of a cat below a backpack"



Prompt: "a photo of a bird left of a couch"

Reasoning

Let's break it down ... The table could be floating in mid-air, or it could be suspended somehow. The suitcase should be on the ground beneath the table ...



Let's break it down ...
A close-up image of a modern smartphone placed on the left side of a well-used tennis racket. Both items are set against a simple background...



Let's break it down ...

the cat is in the foreground
and the backpack is in the
background. The backpack
could be hanging or placed
on something, but it's not
specified ...



Let's break it down ...
we could add some
additional details such as the
type of bird, its posture, and
the setting in which it is
found. Let's assume the bird
is a robin, perched on a
branch near the couch ...



Figure 8: Qualitative results on compositional prompts (Position). We compare DELL-E3, the intermediate reasoning process, and our final RePrompt outputs. While DELL-E3 often struggles with spatial relations (e.g., object positions), our reasoning-guided approach enables more accurate and faithful generations that better match the prompt's compositional constraints.

1088

1089

1090

1091

1092

1093

1094

1095

1096 1097 1098

1099

1100

1101

1102

1103

1104

11051106

1107

1108

1109

1110

1111

11121113

1114

1115

1116

1117

1118

1119

1120

1121 1122

11231124

1125

1126

1127

User prompt **DELL-E3** Reasoning Ours Let's break it down ... The bear could be in a playful or curious position, perhaps looking at the hair dryer curiously. The hair dryer could be placed on a surface like a table or a countertop... Prompt: "a photo of a hair drier and a bear" Let's break it down ... The toilet could be a standard porcelain one, while the computer mouse could be a modern wireless placed side by side... Prompt: "a photo of a toilet and a computer mouse' Let's break it down ... The positioning of these objects isn't specified, so they could be placed side by side, one above the other, or in some other arrangement... Prompt: "a photo of a bottle and a refrigerator" Let's break it down ... The toothbrush could be placed on a surface like a bathroom countertop or sink, while the snowboard could be leaning against a wall or laid flat on a wooden floor. ...

Figure 9: Qualitative results on compositional prompts (Two-object). We show comparisons among DELL-E3, our reasoning process, and RePrompt outputs on prompts involving two objects. DELL-E3 often fails to generate both entities accurately, whereas our method uses explicit reasoning to guide the model in generating semantically correct and compositionally faithful images.

Prompt: "a photo of a toothbrush and a snowboard"



Figure 10: Qualitative results on compositional prompts (Color). We present qualitative comparisons on prompts involving specific color attributes. While DELL-E3 tends to ignore or misinterpret color constraints, our approach leverages explicit reasoning to ensure accurate color grounding for each object, resulting in more faithful visual compositions.

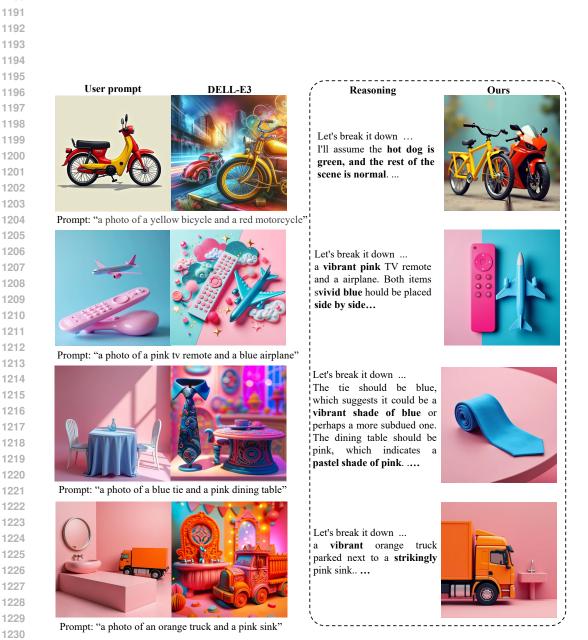


Figure 11: Qualitative results on compositional prompts (Attribute binding). We show examples where prompts specify multiple attributes that must be correctly associated with the corresponding objects. Our method utilizes step-by-step reasoning to disambiguate attribute-object bindings, avoiding attribute swaps or omissions that are common in baseline outputs.