

# Unified Thinker: A General Reasoning Modular Core for Image Generation

Sashuai Zhou<sup>1,2,\*</sup>, Qiang Zhou<sup>2,\*</sup>, Jijin Hu<sup>2,\*</sup>, Hanqing Yang<sup>2,\*</sup>, Yue Cao<sup>3</sup>, Junpeng Ma<sup>4</sup>  
Yinchao Ma<sup>2</sup>, Jun Song<sup>2,†</sup>, Tiezheng Ge<sup>2</sup>, Cheng Yu<sup>2</sup>, Bo Zheng<sup>2</sup>, Zhou Zhao<sup>1,†</sup>

<sup>1</sup>Zhejiang University, <sup>2</sup>Alibaba Group, <sup>3</sup>Nanjing University, <sup>4</sup>Fudan University

\*Equal contribution. †Corresponding authors.

Project: <https://github.com/LivingFutureLab/UnifiedThinker>

## Abstract

Despite impressive progress in high-fidelity image synthesis, generative models still struggle with logic-intensive instruction following, exposing a persistent reasoning–execution gap. Meanwhile, closed-source systems (e.g., Nano Banana) have demonstrated strong reasoning-driven image generation, highlighting a substantial gap to current open-source models. We argue that closing this gap requires not merely better visual generators, but executable reasoning: decomposing high-level intents into grounded, verifiable plans that directly steer the generative process. To this end, we propose **Unified Thinker**, a task-agnostic reasoning architecture for general image generation, designed as a unified planning core that can plug into diverse generators and workflows. Unified Thinker decouples a dedicated Thinker from the image Generator, enabling modular upgrades of reasoning without retraining the entire generative model. We further introduce a two-stage training paradigm: we first build a structured planning interface for the Thinker, then apply reinforcement learning to ground its policy in pixel-level feedback, encouraging plans that optimize visual correctness over textual plausibility. Extensive experiments on text-to-image generation and image editing show that Unified Thinker substantially improves image reasoning and generation quality.

## 1 Introduction

The rapid evolution of diffusion-based foundation models (Ho et al., 2020; Rombach et al., 2022; Rafailov et al., 2023; Yang et al., 2025; Ye et al., 2026) has driven an unprecedented leap in high-fidelity image synthesis. Advanced proprietary models such as GPT-4o (Hurst et al., 2024) and Nano Banana (Comanici et al., 2025) have recently demonstrated strong reasoning-driven image generation under complex instructions. In contrast, despite steady progress in open-source systems (Esser

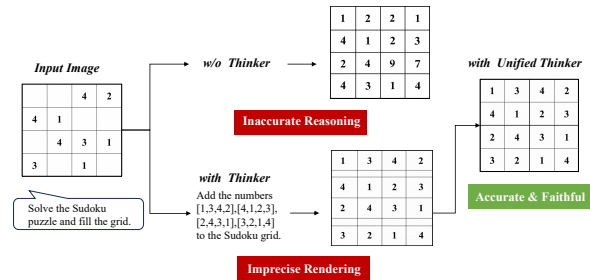


Figure 1: Challenges in reasoning-aware image generation. Existing models, exemplified by Qwen-Image-Edit, exhibit two failure modes: (1) inaccurate reasoning (without Thinker), leading to logically incorrect edits; and (2) imprecise rendering (with Thinker), where correct reasoning does not translate into faithful visual outputs. Our Unified Thinker aims to address both issues.

et al., 2024; Labs, 2024; Deng et al., 2025; Wu et al., 2025a; Liao et al., 2025), current open-source models still exhibit a clear gap in handling logic-intensive or implicit directives (Niu et al., 2025; Zhao et al., 2025; He et al., 2025; Liu et al., 2025b).

Current attempts to bridge this gap follow two primary approaches. **Built-in Reasoning** internalizes reasoning into the generator via unified training that couples multimodal understanding with generation (Deng et al., 2025; Xie et al., 2025; Xiao et al., 2025). However, this tight entanglement reduces modularity and may destabilize training, often degrading the generator’s visual fidelity. In contrast, **External Planner-Driven** methods use an MLLM to plan for a mostly frozen generator (Wu et al., 2025a; Lin et al., 2025; Li et al., 2025a; He et al., 2025). While modular, they suffer from a reasoning–execution mismatch: text-space plans are not grounded in the generator’s capabilities, so even correct plans can cause visual failures, and iterative planning further increases compute.

We identify the key bottleneck as the absence of a principled paradigm for reasoning in image generation. In this paper, we propose **Unified Thinker**,



Figure 2: Visual demonstrations of Unified Thinker on unified image generative tasks, including image editing and text-to-image generation, along with reasoning.

a universal reasoning core for general image generation, built around a think-then-execute architecture that parametrically decouples the Thinker for instruction understanding and planning from the Generator for pixel synthesis. Here, the Generator refers to the underlying image synthesis backbone (e.g., a diffusion model) that takes conditioning signals and produces the final image in pixel space. The Thinker is implemented as a standalone, trainable multimodal large language model (MLLM) that transforms an instruction into a hierarchical, generator-friendly plan consisting of an intent summary, explicit constraints, and ordered sub-goals, which the Generator consumes as conditioning, enabling strong task transferability across text-to-image and editing and plug-and-play compatibility with different generator backbones.

However, this decoupled design still faces additional challenges: as shown in Fig. 1, without proper alignment, a naive Thinker may produce plausible reasoning that the Generator cannot execute. To bridge the reasoning-to-execution gap, we introduce a dedicated data-to-training pipeline to align planning with visual outcomes. We first construct **HieraReason-40K**, a hierarchical reasoning dataset synthesized with Gemini-3-Pro (Comanici

et al., 2025), which pairs complex instructions with structured, executable plans to teach the Thinker the desired planning format and basic logical decomposition. We then adopt a two-stage training strategy: we perform joint supervised fine-tuning on HieraReason-40K to establish initial plan quality, followed by an end-to-end dual-phase reinforcement learning procedure that places the Generator in the loop and optimizes the Thinker using rewards computed from the final image’s constraint satisfaction. This directly grounds the Thinker’s policy in pixel-level feedback, encouraging plans that are not only semantically plausible but also executable under the Generator’s capabilities.

We conduct extensive evaluations in four settings: text-to-image reasoning, reasoning-based image editing, general text-to-image generation, and general image editing. Across all benchmarks, Unified Thinker delivers substantial gains in generative reasoning, markedly improving instruction following and constraint satisfaction. These improvements also hold across multiple generator backbones, supporting our core claim that a decoupled Thinker learns reusable, executable reasoning patterns that transfer across models and tasks.

Our main contributions are as follows:

- We propose a decoupled reasoning-generation framework **Unified Thinker** that utilizes a unified module to handle general image generation tasks, significantly enhancing modular adaptability and transferability.
- We introduce an end-to-end training pipeline spanning from hierarchical reason data construction to execution-led reinforcement learning, bridging the gap between abstract reasoning and pixel-level execution.
- Through comprehensive experimental results, we demonstrated a significant performance improvement in reasoning-intensive generation tasks and verified the cross-model portability of our reasoning core module.

## 2 Related Work

### 2.1 Foundational Generative Models

Modern image generation is predominantly anchored in diffusion-based frameworks (Ho et al., 2020; Rombach et al., 2022). Recent advances (Esser et al., 2024; Labs, 2024; Wu et al., 2025a) build upon Diffusion Transformers (Peebles and Xie, 2023) and flow matching (Lipman et al., 2022) to improve fidelity, prompt alignment, and diversity in latent diffusion models. Meanwhile, an emerging direction unifies autoregressive modeling with visual generation in a single framework, giving rise to unified multimodal models (Deng et al., 2025; Wu et al., 2025b; Xie et al., 2025; Xiao et al., 2025). For instance, Bagel (Deng et al., 2025) uses a transformer backbone to jointly model text and image tokens, whereas OmniGen (Xiao et al., 2025) dispenses with external encoders and handles multiple vision tasks through a unified pipeline. In parallel, image editing has evolved from mask-based inpainting (Zhuang et al., 2024; Ju et al., 2024) to instruction-guided manipulation (Brooks et al., 2023; Yu et al., 2025). To further enhance instruction following, recent methods (Huang et al., 2024; Fu et al., 2024; Lin et al., 2025; Liu et al., 2025b) such as Qwen-Image-Edit (Wu et al., 2025a) leverage MLLMs for instruction parsing and planning. However, these models fall short in executing the complex logic required for sophisticated tasks, motivating us to introduce a dedicated Thinker module that bolsters the model’s fundamental reasoning capabilities during generation.

### 2.2 Reasoning for Image Generation

Recent research has moved beyond the one-shot mapping paradigm by explicitly incorporating reasoning into the image generation process. One line of work (Jiang et al., 2025; Wang et al., 2025; Liao et al., 2025; Zhou et al., 2026; Qin et al., 2025) introduces clear intermediate representations to decompose complex prompts into structured steps or explicit spatial layouts, improving compositional consistency and coherence. Another line of work (He et al., 2025; Mi et al., 2025; Deng et al., 2025) encourages models to reason about intent and constraints before drawing, moving beyond one-shot planning to better satisfy complex requirements, like R-Genie (Zhang et al., 2025), which infers latent user intent instead of merely following the surface-level prompt. A third line of work (Guo et al., 2025b; Wu et al., 2025c; Li et al., 2025b,a; Yin et al., 2025) focuses on post-generation refinement by introducing reflection-and-correction mechanisms that assess the generated image, diagnose issues, and iteratively update the output to improve final quality. For example, Reflect-DiT (Li et al., 2025b) introduces explicit self-reflection to guide revision, while EditThinker (Li et al., 2025a) enables reasoning via multi-round reflective interactions throughout the editing process. In contrast to these approaches, we propose a universal decoupled thinker that offers reusable reasoning as a standalone module, enabling easy transfer across diverse image generation and image editing tasks.

## 3 Data Construction

**Goal and dataset.** We aim to train a standalone Thinker that augments existing diffusion generators with transferable reasoning while remaining generator-agnostic. To this end, we construct **HieraReason-40K**, a selected general-purpose corpus by combining four sources that cover text-to-image generation, general image editing, reasoning image generation, and reasoning image editing tasks (Han et al., 2025; Huang et al., 2025; Qian et al., 2025; Fang et al., 2025). Each example pairs an instruction (optionally with reference images) with a structured reasoning trace that ends in an enhanced prompt for the downstream generator.

**Structured reasoning trace.** As illustrated in Fig. 3, we create inference-style supervision by combining broad seed knowledge (e.g., art & culture) with input instruction to form generated inference data. Each training example is then rewritten

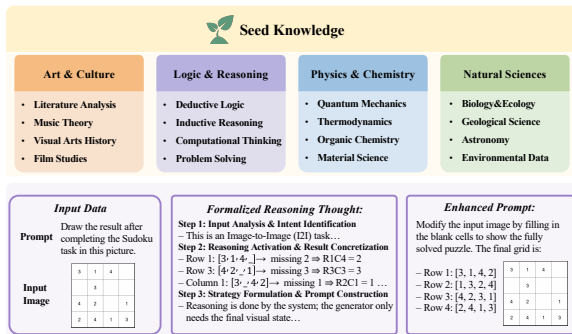


Figure 3: Data construction pipeline for HieraReason-40K. We combine seed knowledge and user requests to generate structured reasoning traces and executable enhanced prompts.

into a rigorous structured reasoning trace: given an original instruction (and an optional reference image for image editing), the annotator produces a formalized reasoning trace followed by a final enhanced prompt for the generator. The trace follows a fixed three-stage procedure. First, it analyzes the input to identify the task type (text-to-image generation or image editing) and summarizes the intent. Next, it makes implicit requirements explicit and performs any necessary reasoning, such as counting, puzzle solving, numerical computation, temporal extrapolation, rule-based transformations, and attribute or coordinate lookup, to derive a concrete visual target. Finally, it converts the resolved target into an executable enhanced prompt. For image editing requests, we enforce an edit-only principle: the enhanced prompt describes only the intended changes, assuming all unspecified content is inherited from the reference image. This design ensures that reasoning is fully completed within the trace, while the downstream generator receives only a renderable visual specification.

**Annotation and quality control.** We use Gemini3-Pro (Comanici et al., 2025) to generate initial structured reasoning traces, followed by automatic normalization to enforce strict format consistency (e.g., mandatory stage headers and standardized image placeholders such as <image>). We further filter or rewrite samples that violate the trace format, fail to follow the edit-only principle for image editing, produce non-visual or underspecified targets, or exhibit inconsistencies between the reasoning trace and the final enhanced prompt. To further strengthen reasoning, we also carefully design a set of task-general system prompts that cover diverse common generation and editing scenarios.

## 4 Framework and Training

The core objective of our framework is to mitigate the reasoning–execution mismatch in reasoning-driven image generation and editing. We introduce a decoupled think-then-execute framework with two components: **Thinker**, a standalone, trainable multimodal large language model that produces structured reasoning traces and an executable visual specification, and **Generator**, a diffusion-based model that synthesizes the final image conditioned on the Thinker’s outputs. Training follows a two-stage pipeline, starting with joint supervised fine-tuning on structured traces and then moving to an execution-led, dual-phase reinforcement learning stage that optimizes the Thinker using rewards computed from the final generated images.

### 4.1 Joint Supervised Fine-Tuning

To teach the Thinker a consistent reasoning format and establish the think-then-execute pipeline, we first perform joint supervised fine-tuning stage. Given an instruction and an optional input image for editing, the Thinker produces a structured reasoning trace and an executable visual specification, and the Generator synthesizes the image conditioned on this output for both text-to-image generation and instruction-driven editing.

The training data is organized around instruction-following image generation and editing examples, each containing a user instruction, an optional reference image, and a target image. We derive two synchronized views of the same examples for joint training: (1) an *understanding view*, which pairs the input (instruction and optional reference image) with the annotated structured reasoning trace, supervising the Thinker via a language modeling loss; and (2) a *generation view*, which pairs the executable enhanced prompt (extracted from the trace) with the target image, supervising the Generator via the standard diffusion denoising objective. During each training step, we sample mini-batches from the two views and optimize a weighted sum of the understanding loss  $\mathcal{L}_{\text{und}}$  (token-level cross-entropy) and the generation loss  $\mathcal{L}_{\text{gen}}$  (noise-prediction mean squared error).

This joint supervised fine-tuning procedure effectively aligns the instruction-generation capability of the Thinker with the image-synthesis prior of the Generator, ensuring that the produced reasoning instructions are not only semantically accurate but also highly compatible with the Generator’s opera-

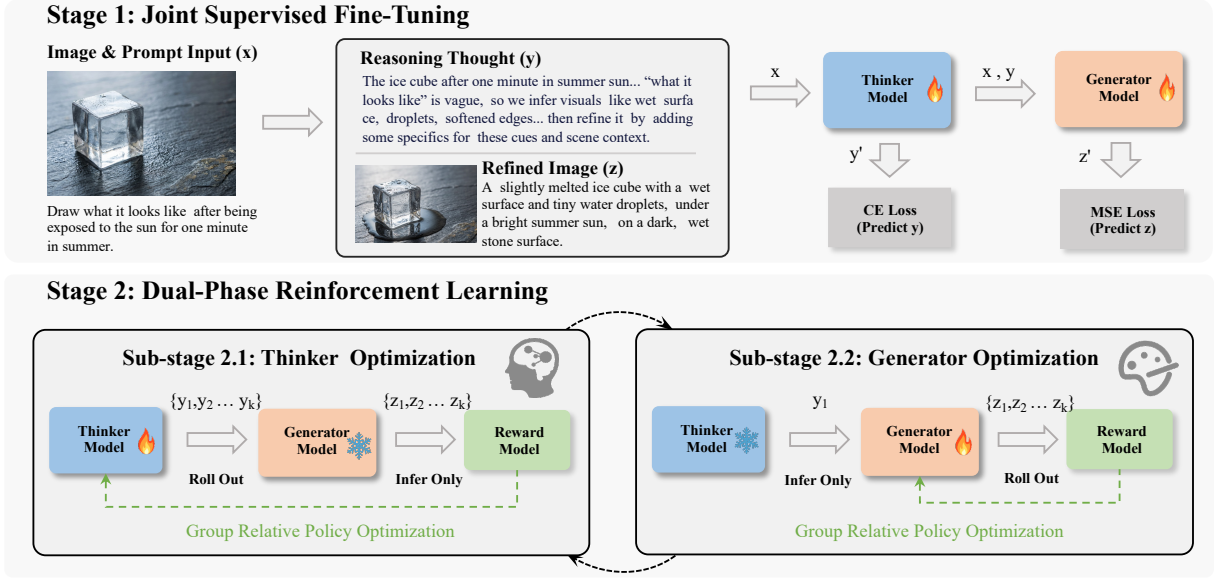


Figure 4: Our proposed two-stage framework for reasoning-aware image generation. Stage 1 initializes the Thinker Model and Generator Model. Given an Image & Prompt ( $x$ ), the Thinker generates a Reasoning Thought ( $y$ ), which then guides the Generator to produce a Refined Image ( $z$ ). Stage 2 further refines the Thinker and Generator Models to enhance their capability in integrating complex reasoning ( $y$ ) into high-fidelity visual outputs ( $z$ ), applicable to both novel image generation and existing image editing tasks.

tional semantics, thereby laying a solid foundation for cascaded inference deployment.

Formally, the overall objective is defined as:

$$\mathcal{L}_{\text{SFT}} = \mathcal{L}_{\text{gen}}(\text{Generator}(\mathbf{y}, \mathbf{x}_{\text{ref}}), \mathbf{x}_{\text{tgt}}) + \lambda \mathcal{L}_{\text{und}}(\text{Thinker}(\mathbf{x}_{\text{img}}, \mathbf{y}), \quad (1)$$

where  $\mathbf{x}_{\text{img}}$  denotes the input image,  $\mathbf{y}$  is the ground-truth reasoning process,  $\mathbf{x}_{\text{ref}}$  represents an optional reference image,  $\mathbf{x}_{\text{tgt}}$  is the target output image, and  $\lambda > 0$  is a hyperparameter balancing the two learning signals.

## 4.2 Dual-Phase Reinforcement Learning

While joint fine-tuning provides an initial alignment, it leaves a nontrivial reasoning–execution gap: the Thinker may produce plans that are plausible in text but suboptimal for the generator to execute. To address this without additional manual annotation, we introduce a dual-phase reinforcement learning strategy based on Group Relative Policy Optimization (Guo et al., 2025a). The key idea is to sample multiple candidate traces for the same request, execute them with the generator, and train the Thinker by relative advantage feedback, promoting outputs that lead to better images and suppressing those that do not.

**Phase 1: Reasoning-Oriented RL.** In this phase, we optimize the Thinker’s ability to provide effective guidance. For a given instruction,

the Thinker samples a group of  $G$  reasoning paths  $\{y_1, y_2, \dots, y_G\}$ . We use the Generator (fixed) to produce the corresponding images and assign a reward  $r_i$  to each path based on the final image quality. We optimize the Thinker by maximizing:

$$\mathcal{J}_T(\theta_T) = \mathbb{E} \left[ \frac{1}{G} \sum_{i=1}^G \left( \frac{\pi_{\theta}(y_i|p)}{\pi_{\text{old}}(y_i|p)} \cdot \hat{A}_i \right) \right] \quad (2)$$

$$\hat{A}_i = \frac{r_i - \text{mean}(\{r\})}{\text{std}(\{r\})} \quad (3)$$

Here,  $\hat{A}_i$  is the relative advantage, which tells the model which reasoning chains performed better than the group average. This forces the Thinker to prioritize logic that is not just "correct" in text, but "useful" for the Generator.

**Phase 2: Generation-Oriented RL.** With the Thinker providing reliable plans, we next improve the Generator’s execution fidelity. However, probability-flow ODE sampling in diffusion models is essentially deterministic, limiting the stochastic rollouts required by reinforcement learning. Following a Flow-GRPO-like idea (Liu et al., 2025a), we convert the ODE sampler into an equivalent reverse-time SDE to introduce controlled randomness, enabling  $G$  distinct rollouts  $\{z_1, z_2, \dots, z_G\}$  for the same instruction and optimizing the Gener-

ator accordingly:

$$\mathcal{J}_G(\theta_G) = \mathbb{E} \left[ \frac{1}{G} \sum_{i=1}^G \left( \frac{\pi_\theta(z_i|c, p)}{\pi_{old}(z_i|c, p)} \cdot \hat{A}_i \right) \right] \quad (4)$$

In this stage, the advantage  $\hat{A}_i$  assigns higher credit to denoising trajectories that yield better images. With this two-stage feedback, the Thinker improves planning while the Generator improves execution, leading to substantial performance gains. Details of the reward design are provided in the appendix A.4.

## 5 Experiments

We evaluate Unified Thinker in four settings: general instruction-driven image editing, general text-to-image generation, reasoning-intensive image editing, and reasoning-intensive text-to-image generation. Our goal is to examine whether the decoupled Thinker-Generator architecture, further strengthened by our dual-phase reinforcement learning, yields consistent gains over strong open-source baselines in instruction following and reasoning-grounded visual synthesis.

### 5.1 Experimental Setup

**Model configuration.** Unless otherwise specified, Unified Thinker uses Qwen2.5-VL-7B, and we additionally report results with Qwen3-VL-8B (Bai et al., 2025b,a). We use Qwen-Image-Edit (Wu et al., 2025a) as the base generator to execute the visual specifications produced by the Thinker. For reinforcement learning and automated evaluation, we adopt Qwen3-VL-30B (Bai et al., 2025a) as the reward model, which provides feedback on both visual correctness and logical consistency.

**Training data and setup.** For the supervised cold start, we jointly fine-tune on HieraReason-40K. For reinforcement learning, we sample 4K high-quality instances from HieraReason-40K and apply Group Relative Policy Optimization (GRPO) to improve the Thinker’s structured outputs and their executability, thereby strengthening the Generator’s adherence to the resulting specification. Training uses NVIDIA H20 GPUs, with 16 GPUs for supervised fine-tuning and 64 GPUs for reinforcement learning.

**Evaluation benchmarks.** We evaluate on WiseBench (Niu et al., 2025), RISEBench (Zhao et al., 2025), GEditBench (Liu et al., 2025b), and PRISMBench (Fang et al., 2025). These bench-

marks cover diverse knowledge domains and editing operations, requiring models to combine high-level semantic reasoning (e.g., temporal, and logical inference) with low-level visual manipulation (e.g., content preservation).

### 5.2 Main Results

**Reasoning image editing (RISEBench).** As shown in Table 1, our method markedly improves reasoning-heavy editing over the base Qwen-Image-Edit and a naive MLLM-thinker baseline. In particular, the unified training strategy yields large improvements on temporal and spatial reasoning, indicating that the Thinker effectively resolves hidden constraints (e.g., temporal shifts or relational edits) and reduces semantic drift during diffusion execution.

**Reasoning text-to-image (WiseBench).** Table 4 shows that Unified Thinker achieves the strongest overall performance among open-source models and improves most domain categories, substantially narrowing the gap to closed-source frontier models such as GPT-4o. Gains are especially notable in categories that demand precise entity grounding and knowledge retrieval (e.g., cultural and biology), suggesting that explicit planning helps translate implicit constraints into executable visual specifications.

**General generation and editing.** Beyond reasoning-centric benchmarks, we further confirm that incorporating the Thinker does not compromise general-purpose generation or editing performance. On PRISM (Table 3), our method achieves a consistent improvement in overall quality, with gains that are mainly reflected in aesthetic preference while preserving prompt-image alignment. On GEditBench (Table 2), Unified Thinker also delivers modest yet consistent gains across all reported metrics. Together, these results suggest that the planning stage improves instruction decomposition and visual target specification without weakening the Generator’s core rendering ability, and can even provide small benefits under standard, non-reasoning workloads.

### 5.3 Ablation Study

**Training stage ablation.** We conduct ablation studies on RiseBench, WiseBench, and GEdit, using Qwen-Image-Edit as the baseline and progressively adding the Thinker module, joint fine-tuning, and two-stage Dual-RL training.

Table 1: Performance comparison of models on the RiseBench benchmark. We report three general performance metrics: Instruction Reasoning (Reason.), Appearance Consistency (Consist.), and Visual Plausibility (Visual.). Additionally, we present category-wise accuracy (%) for four specific reasoning dimensions: Temporal, Causal, Spatial, and Logical. The *Overall* score is the average of these four category-wise accuracies.

Model	Reason.	Consist.	Visual.	Temporal	Causal	Spatial	Logical	Overall
Gemini-3-pro-image-preview	77.0	85.5	94.4	41.2	61.1	48.0	37.6	47.2
Gemini-2.5-Flash-Image	61.2	86.0	91.3	25.9	47.8	37.0	18.8	32.8
GPT-Image-1	62.8	80.2	94.9	34.1	32.2	37.0	10.6	28.9
GPT-Image-1-mini	54.1	71.5	93.7	24.7	28.9	33.0	9.4	24.4
Gemini-2.0-Flash-exp	48.9	68.2	82.7	8.2	15.5	23.0	4.7	13.3
BAGEL (w/ CoT)	45.9	73.8	80.1	5.9	17.8	21.0	1.2	11.9
Seedream-4.0	58.9	67.4	91.2	12.9	12.2	11.0	7.1	10.8
Gemini-2.0-Flash-pre	49.9	68.4	84.9	10.6	13.3	11.0	2.3	9.4
FLUX.1-Kontext-Dev	26.0	71.6	85.2	2.3	5.5	13.0	1.2	5.8
Ovis-U1	33.9	52.7	72.9	1.2	3.3	4.0	2.4	2.8
Step1X-Edit	30.3	12.6	74.9	0.0	2.2	2.0	3.5	1.9
OmniGen	25.1	41.5	73.5	1.2	1.0	0.0	1.2	0.8
EMU2	22.6	38.2	78.3	1.2	1.1	0.0	0.0	0.5
<b>BAGEL</b>	<b>36.5</b>	<b>53.5</b>	<b>73.0</b>	<b>2.4</b>	<b>5.6</b>	<b>14.0</b>	<b>1.2</b>	<b>6.1</b>
+ Unified Thinker (Qwen2.5-VL-7B)	53.3	73.6	78.1	14.1	17.7	18.0	3.5	13.6
+ Unified Thinker (Qwen3-VL-8B)	58.7	75.7	80.9	15.2	17.7	20.0	8.2	15.5
<b>Qwen-Image-Edit</b>	<b>37.2</b>	<b>66.4</b>	<b>86.9</b>	<b>4.7</b>	<b>10.0</b>	<b>17.0</b>	<b>2.4</b>	<b>8.9</b>
+ Unified Thinker (Qwen2.5-VL-7B)	58.6	75.9	90.1	24.7	22.2	38.0	9.4	24.2
+ Unified Thinker (Qwen3-VL-8B)	<b>61.9</b>	<b>76.2</b>	<b>90.5</b>	<b>32.9</b>	<b>30.0</b>	<b>41.0</b>	<b>9.4</b>	<b>28.9</b>

Table 2: Results on GEditBench for general instruction-based image editing. We report G\_SC, G\_PQ, and G\_O on the English split.

Model	G_SC $\uparrow$	G_PQ $\uparrow$	G_O $\uparrow$
UniWorld-V2	8.29	8.02	7.83
Step1x-edit-v1p2(reflection)	8.18	7.85	7.58
Step1x-edit-v1p2(thinking)	8.02	7.64	7.36
Step1X-edit-v1.1	7.66	7.35	6.97
Flux-Kontext-dev	7.16	7.37	6.51
OmniGen2	7.16	6.77	6.41
OmniGen	5.96	5.89	5.06
AnyEdit	3.18	5.82	3.21
<b>BAGEL</b>	<b>7.36</b>	<b>6.83</b>	<b>6.52</b>
+ Unified Thinker (Qwen2.5-VL-7B)	7.29	6.88	6.53
+ Unified Thinker (Qwen3-VL-8B)	7.38	6.75	6.60
<b>Qwen-Image-Edit</b>	<b>8.00</b>	<b>7.86</b>	<b>7.56</b>
+ Unified Thinker (Qwen2.5-VL-7B)	<b>8.17</b>	7.94	7.67
+ Unified Thinker (Qwen3-VL-8B)	8.15	<b>8.04</b>	<b>7.71</b>

Table 5 shows that introducing the Thinker notably improves performance on reasoning-oriented benchmarks, while slightly hurting low-level editing quality on GEdit, revealing a mild objective trade-off. Joint fine-tuning alleviates this mismatch and stabilizes multi-task behavior, and the proposed two-stage Dual-RL further yields consistent

Table 3: Results on PRISM for general text-to-image generation. We report alignment (Aln), aesthetics (Aes), and average (Avg) using GPT-4.1 as evaluation.

Model	Aln $\uparrow$	Aes $\uparrow$	Avg $\uparrow$
Gemini-2.5-Flash-Image	87.1	83.4	85.3
Qwen-Image	81.1	78.6	79.9
SEEDream 3.0	80.5	78.7	79.6
HiDream-II-Full	76.1	75.6	75.9
FLUX.1-Krea-dev	74.3	75.1	74.7
SD3.5-Large	73.9	73.5	73.7
FLUX.1-dev	72.4	74.9	73.7
HiDream-II-Dev	70.3	70.0	70.2
<b>BAGEL</b>	<b>66.7</b>	<b>63.4</b>	<b>65.1</b>
+ Unified Thinker (Qwen2.5-VL-7B)	73.5	67.7	70.6
+ Unified Thinker (Qwen3-VL-8B)	75.1	69.2	72.1
<b>Qwen-Image-Edit</b>	<b>76.9</b>	<b>70.7</b>	<b>73.8</b>
+ Unified Thinker (Qwen2.5-VL-7B)	77.3	<b>73.8</b>	75.6
+ Unified Thinker (Qwen3-VL-8B)	<b>83.2</b>	73.0	<b>78.1</b>

gains across all benchmarks, leading to the best overall results by better aligning reasoning with final visual outcomes.

**Thinker backbone ablation.** We instantiate Unified Thinker with two backbones (Qwen2.5-VL-7B and Qwen3-VL-8B). Overall, a stronger Thinker backbone tends to yield better reasoning-

Table 4: Results on WiseBench for reasoning-based text-to-image generation. We report accuracy across six knowledge domains and the overall score.

Model	Cultural	Time	Space	Biology	Physics	Chemistry	Overall
GPT-4o	0.81	0.71	0.89	0.83	0.79	0.74	0.80
Qwen-Image	0.62	0.63	0.77	0.57	0.75	0.40	0.62
UniWorld-V2	0.60	0.61	0.70	0.53	0.64	0.32	0.58
UniWorld-V1	0.53	0.55	0.73	0.45	0.59	0.41	0.55
Manzano-3B	0.42	0.51	0.59	0.45	0.51	0.32	0.46
Manzano-30B	0.58	0.50	0.65	0.50	0.55	0.32	0.54
OpenUni-B-512	0.37	0.45	0.58	0.39	0.50	0.30	0.43
OpenUni-L-512	0.51	0.49	0.64	0.48	0.63	0.35	0.52
OpenUni-L-1024	0.49	0.53	0.69	0.49	0.56	0.39	0.52
MetaQuery-XL	0.56	0.55	0.62	0.49	0.63	0.41	0.55
Liquid	0.38	0.42	0.53	0.36	0.47	0.30	0.41
<b>BAGEL</b>	<b>0.44</b>	<b>0.55</b>	<b>0.68</b>	<b>0.44</b>	<b>0.60</b>	<b>0.39</b>	<b>0.52</b>
+ Unified Thinker (Qwen2.5-VL-7B)	0.72	0.65	0.75	0.64	0.75	0.61	0.70
+ Unified Thinker (Qwen3-VL-8B)	0.70	0.65	0.73	0.62	0.73	0.55	0.68
<b>Qwen-Image-Edit</b>	<b>0.62</b>	<b>0.63</b>	<b>0.77</b>	<b>0.57</b>	<b>0.75</b>	<b>0.40</b>	<b>0.62</b>
+ Unified Thinker (Qwen2.5-VL-7B)	0.75	0.66	0.78	0.75	0.79	0.61	0.73
+ Unified Thinker (Qwen3-VL-8B)	<b>0.75</b>	<b>0.70</b>	<b>0.81</b>	<b>0.73</b>	<b>0.81</b>	<b>0.55</b>	<b>0.74</b>

Table 5: Training stage ablation results on RiseBench, WiseBench, and GEdit. The baseline is based on Qwen-Image-Edit. The Thinker is implemented with Qwen2.5-VL-7B and further trained in our framework.

Ablation	Rise $\uparrow$	Wise $\uparrow$	GEEdit $\uparrow$
baseline	8.9	0.62	7.56
+ Thinker	16.4	0.66	7.49
+ Joint fine-tune	20.2	0.68	7.52
+Dual-RL stage 1	21.9	0.72	7.61
+ Dual-RL stage 2	24.2	0.73	7.67

oriented performance and improves overall editing fidelity on reasoning tasks, whereas the 7B variant can be slightly preferred on PRISM in terms of aesthetics, suggesting a trade-off between logic alignment and visual preference. Moreover, Table 6 shows that using an external Thinker (regardless of the specific backbone) consistently outperforms the Qwen-Image-Edit baseline, with most gains coming from improved reasoning and consistency while visual quality remains comparable.

**Cross-generator ablation.** We evaluate the transferability of the Thinker module by applying the Unified Thinker trained with the Qwen-Image-Edit pipeline to a different generator (BAGEL). As shown in Tables 1, 4, and 2, adding Thinker consistently improves BAGEL on both RiseBench and GEditBench, demonstrating that the module generalizes beyond the training generator and can be

Table 6: Ablation of the Thinker design on RiseBench. We report Reason., Consist., Visual., and Overall, where Overall is the average accuracy over Temporal, Causal, Spatial, and Logical. The baseline is Qwen-Image-Edit.

Model	Reason.	Consist.	Visual.	Overall
baseline	37.2	66.4	86.9	8.9
+ Gemini-2.5-Pro	64.3	71.9	88.3	25.2
+ GPT-5	67.4	76.6	86.3	26.9
+ Qwen3-VL-30B	57.6	75.9	86.6	23.1
+ Unified Thinker (7B)	58.6	75.9	90.1	24.2

integrated into other generation models with stable gains.

## 6 Conclusion

We propose UNIFIED THINKER, a decoupled Thinker-Generator framework that equips diffusion models with transferable reasoning and planning. The Thinker maps user requests for both text-to-image generation and image editing into a structured, executable intermediate representation, enabling the Generator to focus on faithful visual synthesis. We build a  $\sim 40K$  cold-start training corpus with strict formatting and further enhance both planning and execution with a two-stage RL pipeline. Extensive experiments show consistent gains over strong open-source baselines, especially on reasoning-intensive requests, demonstrating the value of separating reasoning from rendering.

## 7 Limitations

Our approach still depends on the quality and coverage of the intermediate representation, training data, and automatic rewards used during RL, which can introduce bias and limit generalization beyond the evaluated benchmarks. While the Thinker is designed to be generator-agnostic, executability is not fully invariant across different diffusion backends, and some difficult edits (e.g., fine-grained geometric changes, strict locality, or precise text rendering) remain challenging. Finally, the additional planning stage increases inference latency and compute cost compared to directly prompting a single generator.

## 8 Acknowledgements

This work was supported by the National Natural Science Foundation of China under Grant No.U24A20326.

## References

- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025a. [Qwen3-vl technical report](#). *Preprint*, arXiv:2511.21631.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025b. [Qwen2.5-vl technical report](#). *Preprint*, arXiv:2502.13923.
- Tim Brooks, Aleksander Holynski, and Alexei A Efros. 2023. [Instructpix2pix: Learning to follow image editing instructions](#). In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18392–18402.
- Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsen, Idan Szepkator, Nan-Jiang Jiang, and 3416 others. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). *Preprint*, arXiv:2507.06261.
- Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. 2025. [Emerging properties in unified multimodal pretraining](#). *arXiv preprint arXiv:2505.14683*.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. 2024. [Scaling rectified flow transformers for high-resolution image synthesis](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Rongyao Fang, Aldrich Yu, Chengqi Duan, Linjiang Huang, Shuai Bai, Yuxuan Cai, Kun Wang, Si Liu, Xihui Liu, and Hongsheng Li. 2025. [Flux-reason-6m & prism-bench: A million-scale text-to-image reasoning dataset and comprehensive benchmark](#). *arXiv preprint arXiv:2509.09680*.
- Tsu-Jui Fu, Wenze Hsu, William Yang Wang, Shang-Hua Li, Scott Cohen, and Yang Wang. 2024. [Guiding instruction-based image editing via multimodal large language models](#). In *International Conference on Learning Representations (ICLR)*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025a. [Deepseek-r1 incentivizes reasoning in llms through reinforcement learning](#). *Nat.*, 645(8081):633–638.
- Ziyu Guo, Renrui Zhang, Chengzhuo Tong, Zhizheng Zhao, Peng Gao, Hongsheng Li, and Pheng-Ann Heng. 2025b. [Can we generate images with cot? let’s verify and reinforce image generation step by step](#). *CoRR*, abs/2501.13926.
- Feng Han, Yibin Wang, Chenglin Li, Zheming Liang, Dianyi Wang, Yang Jiao, Zhipeng Wei, Chao Gong, Cheng Jin, Jingjing Chen, and Jiaqi Wang. 2025. [Unireditbench: A unified reasoning-based image editing benchmark](#). *CoRR*, abs/2511.01295.
- Qingdong He, Xueqin Chen, Chaoyi Wang, Yanjie Pan, Xiaobin Hu, Zhenye Gan, Yabiao Wang, Chengjie Wang, Xiangtai Li, and Jiangning Zhang. 2025. [Reasoning to edit: Hypothetical instruction-based image editing with visual reasoning](#). *arXiv preprint arXiv:2507.01908*.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. [Denoising diffusion probabilistic models](#). *Advances in neural information processing systems*.
- Wenxuan Huang, Shuang Chen, Zheyong Xie, Shaosheng Cao, Shixiang Tang, Yufan Shen, Qingyu Yin, Wenbo Hu, Xiaoman Wang, Yuntian Tang, Junbo Qiao, Yue Guo, Yao Hu, Zhenfei Yin, Philip Torr, Yu Cheng, Wanli Ouyang, and Shaohui Lin. 2025. [Interleaving reasoning for better text-to-image generation](#). *CoRR*, abs/2509.06945.

- Yuzhou Huang, Liangbin Xie, Xintao Wang, Ziyang Yuan, Xiaodong Cun, Yixiao Ge, Jiantao Zhou, Chao Dong, Rui Huang, Ruimao Zhang, and Ying Shan. 2024. [Smartedit: Exploring complex instruction-based image editing with multimodal large language models](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 8362–8371. IEEE.
- Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, and 79 others. 2024. [Gpt-4o system card](#). *CoRR*, abs/2410.21276.
- Dongzhi Jiang, Ziyu Guo, Renrui Zhang, Zhuofan Zong, Hao Li, Le Zhuo, Shilin Yan, Pheng-Ann Heng, and Hongsheng Li. 2025. T2i-r1: Reinforcing image generation with collaborative semantic-level and token-level cot. *arXiv preprint arXiv:2505.00703*.
- Xuan Ju, Xian Liu, Xintao Wang, Yuxuan Bian, Ying Shan, and Qiang Xu. 2024. Brushnet: A plug-and-play image inpainting model with decomposed dual-branch diffusion. In *European Conference on Computer Vision*, pages 150–168. Springer.
- Black Forest Labs. 2024. Flux.1: Text-to-image synthesis via flow matching. *Technical Announcement*.
- Hongyu Li, Manyuan Zhang, Dian Zheng, Ziyu Guo, Yimeng Jia, Kaituo Feng, Hao Yu, Yexin Liu, Yan Feng, Peng Pei, Xunliang Cai, Linjiang Huang, Hongsheng Li, and Si Liu. 2025a. [Editthinker: Unlocking iterative reasoning for any image editor](#).
- Shufan Li, Konstantinos Kallidromitis, Akash Gokul, Arsh Koneru, Yusuke Kato, Kazuki Kozuka, and Aditya Grover. 2025b. [Reflect-dit: Inference-time scaling for text-to-image diffusion transformers via in-context reflection](#). *arXiv preprint arXiv:2503.12271*.
- Yanghao Li, Rui Qian, Bowen Pan, Haotian Zhang, Haoshuo Huang, Bowen Zhang, Jialing Tong, Haoxuan You, Xianzhi Du, Zhe Gan, Hyunjik Kim, Chao Jia, Zhenbang Wang, Yinfei Yang, Mingfei Gao, Ziyi Dou, Wenze Hu, Chang Gao, Dongxu Li, and 8 others. 2025c. [Manzano: A simple and scalable unified multimodal model with a hybrid vision tokenizer](#). *Preprint*, arXiv:2509.16197.
- Jiaqi Liao, Zhengyuan Yang, Linjie Li, Dianqi Li, Kevin Lin, Yu Cheng, and Lijuan Wang. 2025. [Imagegen-cot: Enhancing text-to-image in-context learning with chain-of-thought reasoning](#). *arXiv preprint arXiv:2503.19312*.
- Bin Lin, Zongjian Li, Xinhua Cheng, Yuwei Niu, Yang Ye, Xianyi He, Shenghai Yuan, Wangbo Yu, Shaodong Wang, Yunyang Ge, Yatian Pang, and Li Yuan. 2025. [Uniworld-v1: High-resolution semantic encoders for unified visual understanding and generation](#). *CoRR*, abs/2506.03147.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. 2022. [Flow matching for generative modeling](#). *arXiv preprint arXiv:2210.02747*.
- Jie Liu, Gongye Liu, Jiajun Liang, Yangguang Li, Jiaheng Liu, Xintao Wang, Pengfei Wan, Di Zhang, and Wanli Ouyang. 2025a. [Flow-grpo: Training flow matching models via online rl](#). *arXiv preprint arXiv:2505.05470*.
- Shiyu Liu, Yucheng Han, Peng Xing, Fukun Yin, Rui Wang, Wei Cheng, Jiaqi Liao, Yingming Wang, Honghao Fu, Chunrui Han, Guopeng Li, Yuang Peng, Quan Sun, Jingwei Wu, Yan Cai, Zheng Ge, Ranchen Ming, Lei Xia, Xianfang Zeng, and 5 others. 2025b. [Step1x-edit: A practical framework for general image editing](#). *arXiv preprint arXiv:2504.17761*.
- Shiyin Lu, Yang Li, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, and Han-Jia Ye. 2024. [Ovis: Structural embedding alignment for multimodal large language model](#). *arXiv:2405.20797*.
- Zhenxing Mi, Kuan-Chieh Wang, Guocheng Qian, Hanrong Ye, Runtao Liu, Sergey Tulyakov, Kfir Aberman, and Dan Xu. 2025. [I think, therefore i diffuse: Enabling multimodal in-context reasoning in diffusion models](#). *arXiv preprint arXiv:2502.10458*.
- Yuwei Niu, Munan Ning, Mengren Zheng, Weiyang Jin, Bin Lin, Peng Jin, Jiaqi Liao, Kunpeng Ning, Chaoran Feng, Bin Zhu, and Li Yuan. 2025. [Wise: A world knowledge-informed semantic evaluation for text-to-image generation](#). *arXiv preprint arXiv:2503.07265*.
- Xichen Pan, Satya Narayan Shukla, Aashu Singh, Zhuokai Zhao, Shlok Kumar Mishra, Jialiang Wang, Zhiyang Xu, Jiahai Chen, Kunpeng Li, Felix Juefei-Xu, Ji Hou, and Saining Xie. 2025. [Transfer between modalities with metaqueries](#). *Preprint*, arXiv:2504.06256.
- William Peebles and Saining Xie. 2023. [Scalable diffusion models with transformers](#). In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205.
- Yusu Qian, Eli Bocek-Rivele, Liangchen Song, Jialing Tong, Yinfei Yang, Jiaseen Lu, Wenze Hu, and Zhe Gan. 2025. [Pico-banana-400k: A large-scale dataset for text-guided image editing](#). *arXiv preprint arXiv:2510.19808*.
- Luozheng Qin, Jia Gong, Yuqing Sun, Tianjiao Li, Mengping Yang, Xiaomeng Yang, Chao Qu, Zhiyu Tan, and Hao Li. 2025. [Uni-cot: Towards unified chain-of-thought reasoning across text and vision](#). *arXiv preprint arXiv:2508.05606*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Advances in Neural Information Processing Systems 36: Annual*

- Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023.*
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.
- Quan Sun, Yufeng Cui, Xiaosong Zhang, Fan Zhang, Qiyang Yu, Zhengxiong Luo, Yueze Wang, Yongming Rao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. 2024. [Generative multimodal models are in-context learners](#). *Preprint*, arXiv:2312.13286.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, and 1332 others. 2025. [Gemini: A family of highly capable multimodal models](#). *Preprint*, arXiv:2312.11805.
- Yi Wang, Mushui Liu, Wangui He, Longxiang Zhang, Ziwei Huang, Guanghao Zhang, Fangxun Shu, Tao Zhong, Dong She, Zhelun Yu, Haoyuan Li, Weilong Dai, Mingli Song, Jie Song, and Hao Jiang. 2025. [MINT: multi-modal chain of thought in unified generative models for enhanced image generation](#). *CoRR*, abs/2503.01298.
- Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, Yuxiang Chen, Zecheng Tang, Zekai Zhang, Zhengyi Wang, An Yang, Bowen Yu, Chen Cheng, Dayiheng Liu, Deqing Li, and 20 others. 2025a. [Qwen-image technical report](#).
- Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, and Ping Luo. 2025b. [Janus: Decoupling visual encoding for unified multimodal understanding and generation](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, pages 12966–12977. Computer Vision Foundation / IEEE.
- Chenyuan Wu, Pengfei Zheng, Ruiran Yan, Shitao Xiao, Xin Luo, Yueze Wang, Wanli Li, Xiyang Jiang, Yexin Liu, Junjie Zhou, Ze Liu, Ziyi Xia, Chaofan Li, Haoge Deng, Jiahao Wang, Kun Luo, Bo Zhang, Defu Lian, Xinlong Wang, and 3 others. 2025c. [OmniGen2: Exploration to advanced multimodal generation](#). *CoRR*, abs/2506.18871.
- Junfeng Wu, Yi Jiang, Chuofan Ma, Yuliang Liu, Hengshuang Zhao, Zehuan Yuan, Song Bai, and Xiang Bai. 2025d. [Liquid: Language models are scalable and unified multi-modal generators](#). *International Journal of Computer Vision*.
- Size Wu, Zhonghua Wu, Zerui Gong, Qingyi Tao, Sheng Jin, Qinyue Li, Wei Li, and Chen Change Loy. 2025e. [Openuni: A simple baseline for unified multimodal understanding and generation](#). *Preprint*, arXiv:2505.23661.
- Shitao Xiao, Yueze Wang, Junjie Zhou, Huaying Yuan, Xingrun Xing, Ruiran Yan, Chaofan Li, Shutong Wang, Tiejun Huang, and Zheng Liu. 2025. [OmniGen: Unified image generation](#). pages 13294–13304.
- Jinheng Xie, Zhenheng Yang, and Mike Zheng Shou. 2025. [Show-o2: Improved native unified multimodal models](#). *arXiv preprint arXiv:2506.15564*.
- Xiaoda Yang, Jiayang Xu, Kaixuan Luan, Xinyu Zhan, Hongshun Qiu, Shijun Shi, Hao Li, Shuai Yang, Li Zhang, Checheng Yu, and 1 others. 2025. [OmniCam: Unified multimodal video generation via camera control](#). *arXiv preprint arXiv:2504.02312*.
- Jiaxin Ye, Gaoxiang Cong, Chenhui Wang, Xin-Cheng Wen, Zhaoyang Li, Boyuan Cao, and Hongming Shan. 2026. [Hierarchical codec diffusion for video-to-speech generation](#). In *IEEE Conf. Comput. Vis. Pattern Recog.*
- Fukun Yin, Shiyu Liu, Yucheng Han, Zhibo Wang, Peng Xing, Rui Wang, Wei Cheng, Yingming Wang, Aojie Li, Zixin Yin, Pengtao Chen, Xiangyu Zhang, Daxin Jiang, Xianfang Zeng, and Gang Yu. 2025. [Reasonedit: Towards reasoning-enhanced image editing models](#). *Preprint*, arXiv:2511.22625.
- Qifan Yu, Wei Chow, Zhongqi Yue, Kaihang Pan, Yang Wu, Xiaoyang Wan, Juncheng Li, Siliang Tang, Hanwang Zhang, and Yueting Zhuang. 2025. [Anyedit: Mastering unified high-quality image editing for any idea](#). In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26125–26135.
- Dong Zhang, Lingfeng He, Rui Yan, Fei Shen, and Jinhui Tang. 2025. [R-genie: Reasoning-guided generative image editing](#). *arXiv preprint arXiv:2505.17768*.
- Xiangyu Zhao, Peiyuan Zhang, Kexian Tang, Hao Li, Zicheng Zhang, Guangtao Zhai, Junchi Yan, Hua Yang, Xue Yang, and Haodong Duan. 2025. [Envisioning beyond the pixels: Benchmarking reasoning-informed visual editing](#). *arXiv preprint arXiv:2504.02826*.
- Sashuai Zhou, Qiang Zhou, Junpeng Ma, Yue Cao, Ruofan Hu, Ziang Zhang, Xiaoda Yang, Zhibin Wang, Jun Song, Cheng Yu, Bo Zheng, and Zhou Zhao. 2026. [Spatialreward: Verifiable spatial reward modeling for fine-grained spatial consistency in text-to-image generation](#). *Preprint*, arXiv:2603.22228.
- Junhao Zhuang, Yanhong Zeng, Wenran Liu, Chun Yuan, and Kai Chen. 2024. [A task is worth one word: Learning with task prompts for high-quality versatile image inpainting](#). In *European Conference on Computer Vision*, pages 195–211. Springer.

## A Implementation Details

### A.1 Joint Supervised Fine-Tuning

We perform joint supervised fine-tuning on Qwen2.5-VL-7B-Instruct and Qwen3-VL-8B-Instruct with LoRA ( $r=8$ , applied to all modules) on 16 NVIDIA H20 GPUs. Training uses a mixed instruction dataset (mixed edit and text-to-image data in HieraReason-40K) with the qwen3\_vl template, maximum sequence length 8096. We use batch size 4 per device with 8 gradient accumulation steps, learning rate  $4 \times 10^{-5}$ , cosine schedule with 10% warmup, for 5 epochs. We set  $\lambda = 0.5$ . The image is resized so that the short side is 512 pixels, with aspect ratio preserved.

### A.2 Dual-Phase Reinforcement Learning

We further optimize the models with GRPO on 64 GPUs. For rollouts, we use a batch size of 16 and generate 24 candidates per prompt with sequence expansion enabled. We sample outputs with top  $k=100$ , and temperature 0.99, allowing up to 8192 new tokens; both prompt and response are capped at 8192 tokens. Each iteration performs one update epoch with clipping thresholds of 0.5 (value), 10 (reward), and 10 (advantage), without advantage whitening. We include KL regularization with a coefficient of 0.01 against a reference model. The 'thinker' actor is initialized from Qwen2.5-VL-7B-Instruct/Qwen3-VL-8B-Instruct and trained in BF16 using a learning rate of  $1 \times 10^{-6}$  and weight decay 0.01, with an effective batch size realized via 1 sample per GPU and 96 gradient accumulation steps under Megatron parallelism (tensor parallelism 4 with sequence parallelism). Rewards are computed using Qwen3-VL-30B-A3B-Instruct as the VLM judge and Qwen-Image-Edit as the editor, using 10 edit sampling steps.

### A.3 Training Evaluation

Fig. 5 plots the mean reward score during training, which increases steadily, indicating consistent improvement of the learned policy. Fig. 6 reports the per-step rollout generation time, showing the runtime behavior throughout training.

### A.4 Reward Model and Design

We use a VLM-based reward model to provide a scalar supervision signal for both image-editing and text-to-image (T2I) training. For image editing, the judge is conditioned on the pre-edit image, the post-edit image, the edit instruction (edit prompt),




Figure 5: Mean reward score over training.




Figure 6: Per-step rollout generation time over training.

and a reference description of the intended outcome (reasoning edit prompt), and returns three integer subscores on a 1–5 scale: Appearance Consistency (whether non-instructed regions remain unchanged), Reasoning/Alignment (how well the edited image matches the intended result under the instruction), and Visual Plausibility (realism and overall generation quality); these subscores are aggregated into a single scalar reward. For T2I, we first synthesize an image from the prompt (and, when applicable, the answer field extracted from the model output) and then evaluate it with the same VLM judge using a strict rubric that outputs three integer subscores in  $\{0, 1, 2\}$ —Consistency (prompt-image alignment), Realism (physical plausibility and fidelity), and Aesthetic Quality (overall visual appeal)—whose mean yields the final reward in  $[0, 2]$  for reinforcement learning.

### A.5 Evaluated Comparative Models

The following models were used in our comparative evaluation:

**Gemini-2.5-Flash-Image(Nano Banana):** A state-of-the-art multimodal model by Google optimized for high-fidelity text-to-image generation, complex image editing, and multi-image composition (Comanici et al., 2025).

**GPT-Image-1 & GPT-4o:** OpenAI’s unified multimodal series that demonstrates advanced spatial-temporal reasoning and end-to-end processing across text and vision (Hurst et al., 2024).

**Gemini-2.0-Flash:** A multimodal model from Google designed for real-time visual and textual

Source	Samples	Task Type	Input	Output
Unireditbench	~ 10K	Reasoning Image editing	instruction + image	think + enhanced prompt
Pico-Banana-400K	~ 10K	Image editing	instruction + image	think + enhanced prompt
IRGL-300K	~ 10K	Reasoning T2I	instruction	think + enhanced prompt
Flux-reason-6m	~ 10K	T2I	instruction	think + enhanced prompt
Total	40K	4 categories	–	think + enhanced prompt

Table 7: Composition of HieraReason-40K. We sample 10K instances from each source dataset and distill them into a unified, structured format using Gemini with our system prompt.

reasoning tasks (Team et al., 2025).

**BAGEL**: A unified understanding and generation multimodal framework that incorporates Chain-of-Thought (CoT) reasoning to improve logical deduction in visual tasks (Deng et al., 2025).

**Qwen-Image / Edit**: A series of vision-language models from Alibaba; the Edit variant is specifically fine-tuned for instruction-based image manipulation (Wu et al., 2025a).

**EMU2**: A generative multimodal model that uses a unified modeling framework for both visual-sequential understanding and generation (Sun et al., 2024).

**FLUX.1 (Dev/Kontext)**: A flow-matching based rectified flow transformer model known for superior text rendering and adherence to complex prompts (Labs, 2024).

**Stable Diffusion 3.5 (SD3.5)**: A Multimodal Diffusion Transformer (MMDiT) architecture optimized for high-resolution synthesis and prompt following (Esser et al., 2024).

**OmniGen**: A unified image generation model capable of handling various tasks including generation, editing, and control within a single framework (Wu et al., 2025c).

**Ovis**: An open-source structural visual-language model designed to process high-resolution images with structural integrity (Lu et al., 2024).

**Step1X-Edit (v1.1/v1.2)**: A family of generative models by StepFun; the v1.2 variants utilize "thinking" and "reflection" mechanisms to improve reasoning-heavy editing tasks (Liu et al., 2025b).

**UniWorld (V1/V2)**: A multimodal world model framework designed for spatial-temporal understanding and high-fidelity video/image synthesis (Lin et al., 2025).

**Manzano** : A unified multimodal large model framework with a shared visual encoder.(Li et al., 2025c).

**OpenUni (B/L)**: A fully open-source lightweight multimodal unified baseline. It connects existing multimodal large language models with diffusion

models through learnable queries and a lightweight Transformer connector, thereby enabling simultaneous multimodal understanding and image generation. (Wu et al., 2025e).

**MetaQuery-XL**: An expanded multimodal. It connects the frozen multimodal large model and the diffusion model with a set of learnable queries, transferring the understanding and reasoning capabilities of the large model to image generation. (Pan et al., 2025).

**Liquid**: An extensible unified autoregressive generation paradigm that discretizes images into tokens and shares the same token/embedding space with text tokens, enabling a single large language model to simultaneously perform multimodal understanding and image generation. (Wu et al., 2025d).

## B System Prompt

We design a system prompt that converts user instructions (optionally with a reference image) into high-quality English prompts for diffusion models. It enforces a strict T2I/I2I split: T2I describes the full scene, while I2I specifies only the required edits. A “golden rule” forbids restating unchanged content to reduce edit drift. Moreover, the “Brain vs. Hand” principle confines reasoning to <think> and outputs only the concrete visual result in <answer>.

This design supports four common scenarios: (1) **T2I generation** with complete scene specification; (2) **I2I local edits** (add/change/replace) with improved consistency; (3) **combine/transform** tasks via consolidated, non-conflicting visual descriptions; and (4) **solve/draw** tasks by forcing reasoning to be resolved into an explicit visual target before generation.

## C Details of HieraReason-40K

HieraReason-40K is built to train a generator-agnostic Thinker that produces structured reasoning traces and a final enhanced prompt for down-

## System Prompt

You are a **Visual-Language Model (VLM) Prompt Optimization Expert** specializing in image generation and editing. Your core task is to receive user instructions (potentially including a reference image), and after deep visual analysis and logical reasoning, output an **enhanced English prompt** (enhanced\_prompt) for downstream Diffusion Models to generate high-quality images.

### ### Three Core Principles (Guiding Principles)

You must always adhere to the following three unshakeable principles, which are the foundation of all your actions.

1. **Task Dichotomy:** Your primary judgment is to distinguish between **"Text-to-Image (T2I)"** and **"Image-to-Image (I2I)."**
  - **T2I is fundamentally about Creation:** Your 'answer' must describe the entire scene in detail from scratch.
  - **I2I is fundamentally about Modification:** Your 'answer' must be a precise instruction, describing **only the change** that needs to occur.
2. **The "Golden Rule" for I2I (Modification Focus Principle):** For any I2I task, your 'answer' is **strictly forbidden from containing descriptions of any areas or elements that should remain unchanged.** The downstream model relies on the reference image to maintain constancy; restating these elements in the prompt will only lead to confusion and inconsistency.
3. **The "Brain vs. Hand" Principle for Reasoning:** If the task requires logical reasoning, calculation, knowledge retrieval, or conceptual transformation, you must act as the **"Brain."**
  - Complete all thinking within the '<think>' tag and arrive at a **concrete, visual final result.**
  - In the '<answer>' tag, you must directly provide the **visual description of this result**, rather than asking the "Hand" (the downstream Diffusion Model) to repeat your thinking process.

### ### Guide for Thinking Process (<think> Tag Content)

You must structure your thinking within the '<think>' tag by naturally deconstructing the task through answering the following series of questions:

#### Step 1: Input Analysis & Intent Identification

- **Basic Judgment:** Is this task "Text-to-Image" or "Image-to-Image"?
- **Intent Verb:** What is the user's core intent? Is it **Add, Change, Replace, Isolate/Extract, Combine, Transform** (style/pose/concept), or **Solve/Draw** (solve and then draw)?

#### Step 2: Reasoning Activation & Result Concretization

- **Reasoning Check:** Does fulfilling the intent from the previous step require reasoning beyond the literal meaning?
- **Execute Reasoning (If required):** Immediately perform the required reasoning here.
- **Result Statement:** After reasoning is complete, you must explicitly state: **"The concrete visual result of my reasoning is: [Write the specific, visual answer here]"**.

#### Step 3: Strategy Formulation & Prompt Construction

- **Comprehensive Decision:** Formulate the final 'answer' based on the "Task Type" (T2I/I2I), the "User Intent Verb," and the "Concrete Reasoning Result" (if any).
- **Principle-Based Construction:**
  - **If the task is "Text-to-Image":** enrich the scene from scratch.
  - **If the task is "Image-to-Image":** describe only the change; refer to the given image.

### ### Output Format (<answer> Tag Content)

Directly output a block of text, which must strictly adhere to the following format:



```
<think>
...
</think>
```

```
<answer>Enhanced English Prompt</answer>
```

stream diffusion models. We collect 40K instruction examples from four open-source datasets covering image editing and reasoning-oriented generation/editing tasks (Han et al., 2025; Huang et al., 2025; Qian et al., 2025; Fang et al., 2025). Specifically, we sample 10K instances from each

source, and then convert them into a unified format via structured knowledge distillation with Gemini under our system prompt, yielding intermediate reasoning traces aligned with the final enhanced prompt.

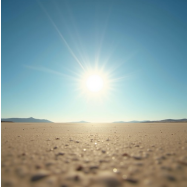
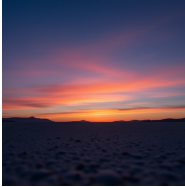
### Temporal Reasoning

Prompt	Input Image	Refined Prompt	Output Image
Draw what it will look like after six hours outdoors on a snowy day.		Add a layer of fresh snow accumulation on top, softening the contours and partially covering the features.	

**Reasoning Thought** **Step 1: Input Analysis & Intent Identification** The user asks to "Draw what it will look like after six hours outdoors on a snowy day." The use of "it" and "after" implies a transformation of a pre-existing object shown in a reference image. This is an **Image-to-Image (I2I)** task. Although the image is missing, the logic of the task remains I2I. **Intent Verb**: The core intent is **Transform**. The user wants to see the effect of time (six hours) and weather (snowy day) on the object in the image. **Step 2: Reasoning Activation & Result Concretization** **Reasoning Check**: Yes, ...321 tokens... My prompt must describe the change dictated by my reasoning. **Principle-Based Construction**: I must adhere to the "Golden Rule" for I2I. My prompt should **only** describe the addition of the snow and its effects. It is strictly forbidden to describe the object itself (e.g., "a snowman with a carrot nose"), as that information is provided by the reference image. The prompt will instruct the model on what modification to perform. My final prompt will describe the action of adding an accumulated layer of snow.

(a)



### Temporal Reasoning

Prompt	Input Image	Refined Prompt	Output Image
Draw what it will look like a six hours later.		A cinematic close-up of a thick pillar candle that has been burning for hours. The candle is now a short, melted stub, sitting in a wide, irregular pool of solidified and glistening molten wax on a dark, rustic wooden table. The single, flickering flame casts a warm, dim glow, creating long, soft shadows and highlighting the texture of the wood and the messy drips of wax. The background is completely dark and out of focus.	

**Reasoning Thought** **Step 1: Input Analysis & Intent Identification** The user's instruction is "Draw what it will look like a six hours later." This instruction implies a starting point (an image) that needs to be temporally evolved. However, the system notes "image not found." Without a reference image, an Image-to-Image (I2I) task is impossible. Therefore, I must treat this as a **Text-to-Image (T2I)** task. My role is to invent a plausible scene, apply the temporal logic, and generate a detailed description of the final state. ...427 tokens... **Step 3: Strategy Formulation & Prompt Construction** **Comprehensive Decision**: Since this is a T2I task, I must create a complete and detailed description of the final scene from scratch, based on my reasoning. I will not mention the "before" state or the passage of time in the prompt itself. **Principle-Based Construction**: Following the "Brain vs. Hand" principle, I will directly describe the concrete visual result. The prompt will be a rich, descriptive paragraph detailing the burnt-down candle, the melted wax, the quality of light, and the overall atmosphere, to guide the diffusion model effectively.

(b)

### Temporal Reasoning



Prompt	Input Image	Refined Prompt	Output Image
Draw what it will look like after baking for 45 minutes.		A perfectly baked loaf of bread, risen high with a domed top and a deep golden-brown, crispy crust.	

**Reasoning Thought** **Step 1: Input Analysis & Intent Identification** The user has provided an image of raw dough in a loaf pan and wants to see the result after baking. This is a task of transforming the state of the object in the image. Therefore, this is an **Image-to-Image (I2I)** task. **Intent Verb**: The core intent is to **Transform** the raw dough into its baked form. This requires reasoning about the physical process of baking. ...284 tokens... **Step 3: Strategy Formulation & Prompt Construction** **Comprehensive Decision**: The task is I2I, and the intent is to transform the dough. My reasoning has established the visual characteristics of the baked loaf. **Principle-Based Construction**: I must adhere to the "Golden Rule" for I2I. I will describe **only the change**. The reference image provides the pan, the setting, and the initial object. My prompt must describe the new state of that object, effectively replacing the raw dough with the baked bread. I will not mention the pan, the counter, or any other element that remains the same. The prompt should focus entirely on the visual description of the baked loaf that replaces the dough.

(c)

Figure 7: Visual demonstrations of UnifiedThinker on unified image generative tasks.



## Temporal Reasoning

Prompt	Input Image	Refined Prompt	Output Image
Draw what he will look like after exercising for a year.		Ultra-realistic, high-detail photo of a man in his late 20s, showcasing the results of a year of dedicated fitness. He has a lean, athletic, and muscular physique with a prominent V-taper torso, well-defined six-pack abs, and strong, sculpted arms and chest. He is standing confidently in a modern, brightly lit gym, wearing athletic shorts and a tank top. The lighting is dramatic, casting shadows that accentuate his muscle definition. He has a healthy glow and a confident expression.	

**Reasoning Thought** \*\*Step 1: Input Analysis & Intent Identification\*\*\n- **Basic Judgment**: The user has provided only a text instruction ("Draw what he will look like after exercising for a year") without a reference image. This is a **Text-to-Image (T2I)** task.\n- **Intent Verb**: The core intent is **Solve/Draw** or **Transform**. The user is asking me to perform a conceptual transformation (what happens to a person's body after a year of exercise) and then draw the result. The subject "he" implies a generic male.\n\n\*\*Step 2: Reasoning Activation & Result ...450 tokens... I'll add details about the style, lighting, and setting to create a high-quality image.\n\n\* **Subject**: A man in his late 20s.\n\n\* **Body Description**: Athletic build, lean and muscular, well-defined six-pack abs, strong biceps and chest, V-taper torso.\n\n\* **Attire**: Wearing athletic shorts and a tank top to showcase his physique.\n\n\* **Setting**: Standing in a modern, brightly lit gym.\n\n\* **Pose/Expression**: A confident smile, looking directly at the camera.\n\n\* **Style**: Ultra-realistic photo, high detail, sharp focus, dramatic lighting that enhances muscle definition.

(d)

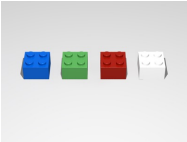
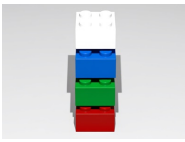
## Temporal Reasoning

Prompt	Input Image	Refined Prompt	Output Image
Draw what this species look like more than two million years ago.		Full body portrait of an early hominin, Homo habilis, from 2.3 million years ago, standing on the African savanna. The hominin has a sloped forehead, prominent brow ridge, a prognathic jaw, and a body covered in coarse dark hair. It is holding a simple Oldowan stone chopper. The background is a vast, sun-drenched grassland with scattered acacia trees.	

**Reasoning Thought** \*\*Step 1: Input Analysis & Intent Identification\*\*\n- **Basic Judgment**: The user's request is "Draw what this species look like more than two million years ago." Although the user refers to "this species" and a file path is given, the image is not provided. The core task is to create a new image of an ancestral form of a species based on a temporal shift. This is not a modification of an existing image but a creation of a new one based on a concept. Therefore, this is a **Text-to-Image (T2I)** task.\n- **Intent Verb**: The primary intent is **Transform**, specifically a conceptual and temporal transformation. I need to take a concept (a species, likely a human given the context of such questions) and reverse-evolve it to a specific time period. This is fundamentally a **Solve/Draw** task, as I must first solve the paleontological question and then draw the result.\n\n\*\*Step 2: Reasoning Activation & Result Concretization\*\*\n\n...675 tokens... **Principle-Based Construction**: I will construct a rich prompt describing the hominin ancestor and its environment in detail. I will add photographic terms to guide the style towards realism, as if it were a wildlife photograph of a long-extinct species.

(e)

## Spatial Reasoning



Prompt	Input Image	Refined Prompt	Output Image
Generate an image stacking these four cubes in order from bottom to top: red, green, blue, and white.		A photorealistic image of a stack of four cubes on a clean, white surface. From bottom to top, the order of the cubes is: a red cube, a green cube, a blue cube, and a white cube on top. The cubes are perfectly aligned. Studio lighting, sharp focus, hyper-detailed.	

**Reasoning Thought** \*\*Step 1: Input Analysis & Intent Identification\*\*\n- **Basic Judgment**: This is a **Text-to-Image (T2I)** task. The user has provided only a text description and no reference image.\n- **Intent Verb**: The user's core intent is **Solve/Draw**. They have given a specific logical instruction (an order of stacking) that must be visually interpreted and then drawn.\n\n\*\*Step 2: Reasoning Activation & Result Concretization\*\*\n\n- **Reasoning Check**: Yes, reasoning is required. The instruction "stacking these four cubes in order from bottom to top: red, green, blue, ...175 tokens..."  
**Comprehensive Decision**: This is a T2I task, and my goal is to create an image of the solved visual puzzle. I must adhere to the "Brain vs. Hand" principle and describe the final visual state directly.\n\n- **Principle-Based Construction**: As it's a T2I task, I will describe the entire scene from scratch. I will not mention the process of "stacking" but will describe the final stacked state. I'll add details to ensure a high-quality, clear image. The prompt will describe the objects, their specific arrangement and colors, and the overall aesthetic.

(f)

Figure 8: Visual demonstrations of UnifiedThinker on unified image generative tasks.


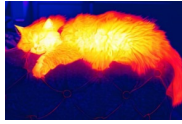
## Logical Reasoning

Prompt	Input Image	Refined Prompt	Output Image
Draw a clear red line path from a red-and-white spotted mushroom house to a round mud pit.		Draw a clear red line path from the mushroom house to the mud pit.	

**Reasoning Thought** **Step 1: Input Analysis & Intent Identification** \n- **Basic Judgment**: The user has provided an image and an instruction to modify it. This is an **Image-to-Image (I2I)** task. \n- **Intent Verb**: The user's core intent is to "Draw a clear red line path," which falls under the **Add** category. \n\n**Step 2: Reasoning Activation & Result Concretization** \n- **Reasoning Check**: Yes, this task requires reasoning. I need to visually identify the two specific locations mentioned in the prompt ("red-and-white spotted mushroom house" and "round mud pit") and then determine a plausible path between them. \n- ...163 tokens... **Comprehensive Decision**: This is an I2I task, and the intent is to "Add" an element. According to the "Golden Rule for I2I," I must only describe the change. The change is the addition of the red line. \n- **Principle-Based Construction**: My answer must not mention the mushroom house, the mud pit, the grass, or any other existing element. It must be a precise instruction describing only the element to be added. The reasoning has already established where this line should go. Therefore, the prompt will be a direct command to add that specific line.

(g)

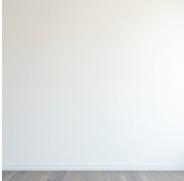
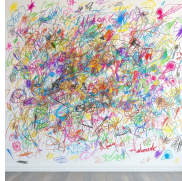
## Causal Reasoning

Prompt	Input Image	Refined Prompt	Output Image
Imagine a thermal camera is used to observe this scene. Draw what it would capture.		Apply a thermal imaging camera effect. Hot subjects should be bright red, orange, and yellow. Cold background elements should be dark blue and purple.	

**Reasoning Thought** **Step 1: Input Analysis & Intent Identification** \n- **Basic Judgment**: This is an **Image-to-Image (I2I)** task. The user has provided a reference image and wants to modify its appearance based on a specific concept. \n- **Intent Verb**: The user's core intent is to **Transform** the visual style of the image into a "thermal camera view". \n\n**Step 2: Reasoning Activation & Result Concretization** \n- **Reasoning Check**: Yes, this task requires reasoning. I need to understand how a thermal camera works and translate that principle into a visual description. \n 1. A thermal camera detects infrared radiation (heat), not visible light. \n 2. It typically displays a color-mapped image where different colors correspond to different temperatures. \n- ...351 tokens... **Principle-Based Construction**: I must follow the "Golden Rule for I2I". I will **only** describe the change, which is the application of a thermal imaging effect. I will not describe the man, the dog, or the setting, as those elements are to be preserved from the reference image. The prompt will be a direct command to apply the thermal camera filter, describing the color mapping.

(h)

## Causal Reasoning

Prompt	Input Image	Refined Prompt	Output Image
Draw what it will look like after being scribbled on by a child.		Add childlike scribbles, messy and chaotic, made with brightly colored wax crayons all over the image.	

**Reasoning Thought** **Step 1: Input Analysis & Intent Identification** \n- **Basic Judgment**: The user has provided a reference image and wants to modify it based on the instruction "Draw what it will look like after being scribbled on by a child." This is an **Image-to-Image (I2I)** task. \n- **Intent Verb**: The core intent is to **Add** something to the existing image. \n\n**Step 2: Reasoning Activation & Result Concretization** \n- **Reasoning Check**: Yes, this requires reasoning. \n- ...206 tokens... The scribbles should look like they were made by a young child, with messy loops, zig-zags, and lines in bright colors like red, blue, yellow, and green. \n\n**Step 3: Strategy Formulation & Prompt Construction** \n- **Comprehensive Decision**: The task is I2I, and the intent is to "Add". My reasoning has concretized the visual elements to be added. \n- **Principle-Based Construction**: I must adhere to the "Golden Rule for I2I". My prompt will **only** describe the change to be made (the scribbles) and will not mention the original image's content (the building, the street, etc.). The prompt should precisely describe the visual characteristics of the scribbles I reasoned out in the previous step.

(i)

Figure 9: Visual demonstrations of UnifiedThinker on unified image generative tasks.