

# coDrawAgents: A Multi-Agent Dialogue Framework for Compositional Image Generation

Chunhan Li<sup>1,7</sup> Qifeng Wu<sup>2</sup> Jia-Hui Pan<sup>3</sup> Ka-Hei Hui<sup>3</sup> Jingyu Hu<sup>3</sup> Yuming Jiang<sup>4</sup>  
Bin Sheng<sup>5</sup> Xihui Liu<sup>6</sup> Wenjuan Gong<sup>7</sup> Zhengzhe Liu<sup>1\*</sup>

<sup>1</sup>Lingnan University <sup>2</sup>Carnegie Mellon University <sup>3</sup>The Chinese University of Hong Kong  
<sup>4</sup>Alibaba DAMO Academy <sup>5</sup>Shanghai Jiao Tong University <sup>6</sup>The University of Hong Kong  
<sup>7</sup>China University of Petroleum (East China)

## Abstract

*Text-to-image generation has advanced rapidly, but existing models still struggle with faithfully composing multiple objects and preserving their attributes in complex scenes. We propose **coDrawAgents**, an interactive multi-agent dialogue framework with four specialized agents: Interpreter, Planner, Checker, and Painter that collaborate to improve compositional generation. The Interpreter adaptively decides between a direct text-to-image pathway and a layout-aware multi-agent process. In the layout-aware mode, it parses the prompt into attribute-rich object descriptors, ranks them by semantic salience, and groups objects with the same semantic priority level for joint generation. Guided by the Interpreter, the Planner adopts a divide-and-conquer strategy, incrementally proposing layouts for objects with the same semantic priority level while grounding decisions in the evolving visual context of the canvas. The Checker introduces an explicit error-correction mechanism by validating spatial consistency and attribute alignment, and refining layouts before they are rendered. Finally, the Painter synthesizes the image step by step, incorporating newly planned objects into the canvas to provide richer context for subsequent iterations. Together, these agents address three key challenges: reducing layout complexity, grounding planning in visual context, and enabling explicit error correction. Extensive experiments on benchmarks GenEval and DPG-Bench demonstrate that coDrawAgents substantially improves text-image alignment, spatial accuracy, and attribute binding compared to existing methods.*

## 1. Introduction

Text-to-Image (T2I) generation has emerged as a pivotal area in artificial intelligence, enabling the creation of visual content from textual descriptions [47]. However, cur-

rent T2I models often struggle with user controllability, particularly concerning the reasonable arrangement and relationships of objects within the generated image [22, 66]. These models can exhibit numerical and spatial inaccuracies, leading to challenges in tasks such as faithful layout arrangement [68] and maintaining compositional faithfulness, where the generated image accurately reflects the structure and relationships described in the text [4, 35].

To address the above limitations, existing works have explored the use of Large Language Models (LLMs) and agentic frameworks to assist in the generation of spatial layouts for images. Some approaches [16, 31] leverage the planning capabilities of LLMs to interpret the input text and propose arrangements for the described objects, aiming to improve the structural coherence of the generated images. Furthermore, recent studies explore **agent-based frameworks** that uses multiple specialized LLM agents for text-to-image generation. For instance, MCCD [28] demonstrates the effectiveness of multi-agent collaboration for compositional diffusion. However, both single-agent and existing multi-agent systems lack the interactive, closed-loop reasoning required for reliable compositional generation. Single-agent methods place all parsing, planning, and verification on one model, making early spatial mistakes difficult to detect or revise. Meanwhile, many multi-agent frameworks are essentially fixed pipelines without negotiation or visual grounding, so errors still propagate uncorrected. As a result, neither paradigm provides the mutual checking, iterative refinement, or canvas-aware reasoning needed for stable layout planning.

These limitations are further amplified in scenes containing multiple objects, where stable layout planning requires overcoming several inherent challenges. First, global layout planning incurs quadratic relational complexity among objects, making it difficult for a single planner to capture all dependencies [16, 68]. Second, most approaches predict layouts without access to visual context, forcing the planner

\*Corresponding author: zhengzheliu@ln.edu.hk

to “imagine” the scene in isolation, which often leads to incoherent or unrealistic arrangements. Third, most existing works use diffusion-based models and diffusion pipelines typically commit to a coarse global structure in early denoising steps, with fine details added only later [4, 22]. As a result, errors such as misplaced objects or incorrect attributes, once introduced early, are difficult to correct due to the lack of explicit error-correction mechanisms.

In this work, we propose an interactive multi-agent dialogue framework with four specialized agents: an Interpreter for generative mode selection and text decomposition, a Planner for incremental layout reasoning, a Checker for spatial and semantic verification, and a Painter for visual synthesis. The Interpreter determines whether to invoke the layout-free mode that connect the Painter directly or to activate the layout-aware mode for complex scenes. In the layout-free mode, the interpreter directly call a text-to-image painter for generation. On the contrary, in the layout-aware mode, unlike pipeline-based designs, our four agents engage in dynamic dialogue: it parses the text into attribute enriched object descriptions, ranks them by semantic salience, groups equal-priority objects for joint generation, and schedules iterative plans. Then the Planner incrementally proposes layouts one object (or group) at a time, the Checker validates spatial and semantic consistency against the text and evolving scene, and the Painter synthesizes the image step by step in a training-free and plug-and-play manner, with the evolving canvas providing crucial visual context for subsequent iterations.

Our framework effectively addresses the core challenges of prior methods. First, instead of performing global planning over all objects simultaneously, the Planner, guided by the Interpreter, adopts a divide-and-conquer strategy by reasoning about objects with the same semantic priority level at a time, thereby substantially reducing layout complexity. Second, the Planner leverages the evolving visual context from the Painter’s canvas, ensuring that layout predictions are grounded in the actual scene rather than imagined in isolation. Third, the Checker introduces an explicit error-correction mechanism, validating object placement alignment, and applying necessary adjustments to improve layout faithfulness. Unlike single-agent or pipeline systems, our agents reason collaboratively, allowing planning, checking, and synthesis to inform one another. This closed-loop interaction enables `codrawAgents` to revise errors, yielding more stable and faithful alignment between text and image.

We evaluate `coDrawAgents` on the GenEval and DPG-Bench benchmarks. On GenEval, our framework sets a new state of the art with substantial improvements in compositional fidelity, object relations, and attribute binding over prior approaches. On DPG-Bench, which stresses long-context and multi-object reasoning, `coDrawAgents` consistently outperforms recent baselines, achieving stronger

spatial accuracy and text–image consistency. Qualitative comparisons further highlight these gains: whereas existing methods often misplace objects, miscount, or confuse attributes, `coDrawAgents` generates coherent layouts and high-quality images aligned with the textual descriptions. These results demonstrate the effectiveness of our interactive multi-agent design for compositional image generation tasks. **Codes will be released upon publication.**

In summary, our contributions are threefold:

- We introduce **coDrawAgents**, an interactive multi-agent dialogue framework with four specialized agents: Interpreter, Planner, Checker, and Painter that collaborate dynamically instead of following a fixed pipeline.
- We propose three technical innovations: (i) a divide-and-conquer planning strategy that reduces layout complexity, (ii) grounding layout decisions in the evolving visual context for stronger spatial alignment, and (iii) an explicit error-correction mechanism via the Checker to enhance faithfulness.
- We achieve **state-of-the-art performance** on GenEval and DPG-Bench, showing clear gains in text–image consistency, spatial accuracy, and attribute binding.

## 2. Related Work

**Text-to-Image Generation** The field of text-to-image (T2I) generation has seen rapid progress, initially driven by Vector Quantized GANs (VQGANs) [13] paired with CLIP guidance [45]. The paradigm shifted significantly with the advent of diffusion models, which led to remarkable improvements in image quality and image-prompt alignment. Foundational models such as DALL-E [46], Imagen [48], and Stable Diffusion [47] established the potential of this new approach. Subsequent efforts have focused on scaling, with models including [2, 12, 26, 42] further enhance performance. More recently, there is a growing trend of integrating Multimodal Large Language Models (MLLMs) directly into the generation process to improve prompt comprehension and contextual reasoning [8, 11, 17, 18, 24, 27, 32, 40, 44, 50, 55, 57–59, 63]. However, these monolithic models still often struggle with fine-grained control over object composition and complex spatial relationships.

**Layout-to-Image Generation** To address the challenge of precise object placement, Layout-to-Image (L2I) generation conditions synthesis on explicit spatial information [10, 16, 25, 36, 38, 39, 65, 68–70], typically in the form of bounding boxes or segmentation masks. ControlNet [66] and GLIGEN [29] demonstrated spatial grounding in pre-trained diffusion models, while later works explored LLM-based layout generation [16, 31], training-free constraints [61], and fine-grained regional controls [9]. These models require explicit spatial conditioning as an input, limiting their applicability when given only text input.

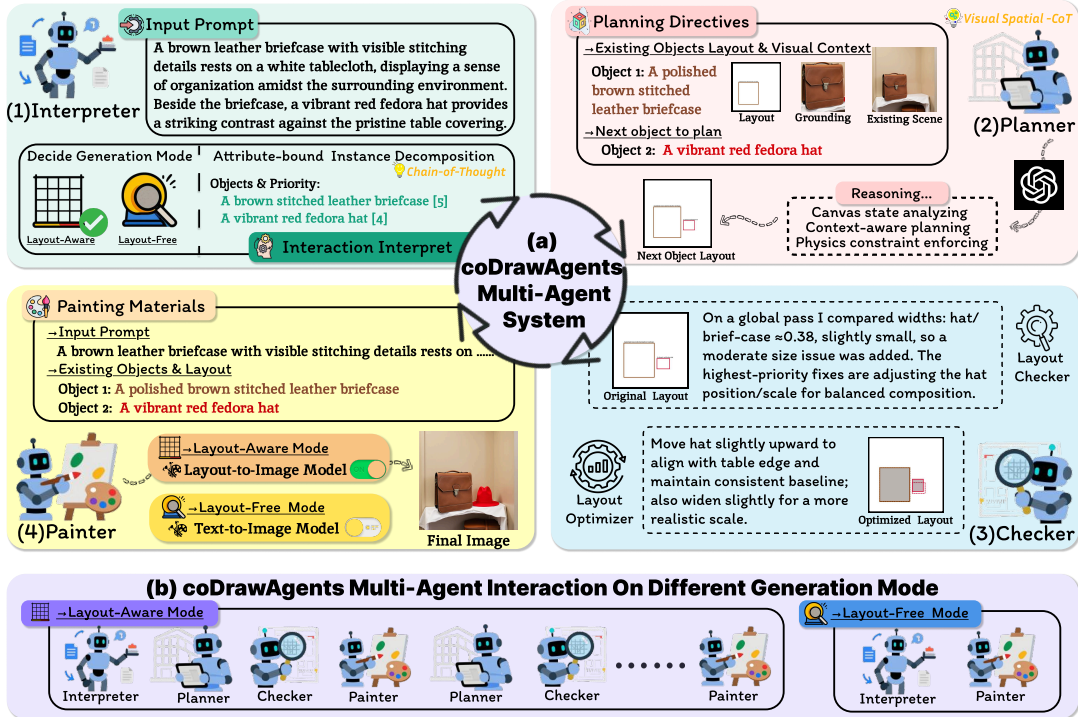


Figure 1. Overview of our proposed coDrawAgents framework. (a) The multi-agent system consists of four specialized agents: (1) Interpreter, (2) Planner, (3) Checker, and (4) Painter. The Interpreter adaptively selects between layout-free and layout-aware modes; in the latter, the Planner incrementally proposes layouts, the Checker validates and refines them, and the Painter synthesizes the evolving canvas. (b) Illustration of the multi-agent interaction process in different generation modes, showing iterative collaboration in layout-aware mode and the direct pathway in layout-free mode.

**Compositional Text-to-Image Generation** Ensuring compositional faithfulness, where the generated images reflect all objects, attributes, and relations in a prompt, remains a key challenge. Early works such as Composable Diffusion [35] and Attend-and-Excite [4] combine concepts or refine attention guidance. Subsequent methods introduced layout reasoning as an intermediate step, for instance LayoutLLM-T2I [43], LLM Blueprint [17], ALR-GAN [51], and LMD [31], which use LLMs or refinement modules to predict layouts that guide diffusion. RPG [62] extends this idea by denoising subregions in parallel, while PlanGen [1] integrates layout planning with synthesis. Most recently, GoT [15] employs a "Generation Chain-of-Thought" to produce a reasoning trace of semantic and spatial relations. While these approaches improve relational reasoning, they typically perform planning without visual feedback, making it difficult to resolve occlusion, depth, or other complex spatial interactions.

**Agent for Image Generation** Recent works have begun to explore agent-based paradigms for image generation, ranging from multi-agent prompt decomposition [28], foreground-conditioned inpainting [52], and self-correcting or interactive editing [37, 56, 60], to more recent directions such as self-improving agents [53], multicultural gen-

eration [3], training-free pipelines [5], and proactive multi-turn dialogue [21]. While these systems demonstrate the potential of agent designs, they are often limited by either fixed sequential pipelines (e.g., T2I-Copilot), planning solely from text without grounding (e.g., MCCD), or relying mainly on task scheduling or user queries without iterative visual feedback (e.g., Talk2Image, Proactive Agents). In contrast, our coDrawAgents framework introduces a closed-loop multi-agent dialogue where the Planner, Checker, and Painter interact continuously with the evolving canvas, achieving stronger interactivity and more faithful compositional generation.

### 3. coDrawAgents System

Given an input text  $T$ , our objective is to generate an image  $I$  that faithfully aligns with the semantic content and spatial arrangement of the text. As shown in Figure 1, we propose coDrawAgents, an interactive multi-agent dialogue framework in which four specialized agents: Interpreter, Planner, Checker, and Painter collaborate for compositional text-to-image generation. We will first introduce the details of agents collaboration and then each agent respectively in the remaining part of this section.

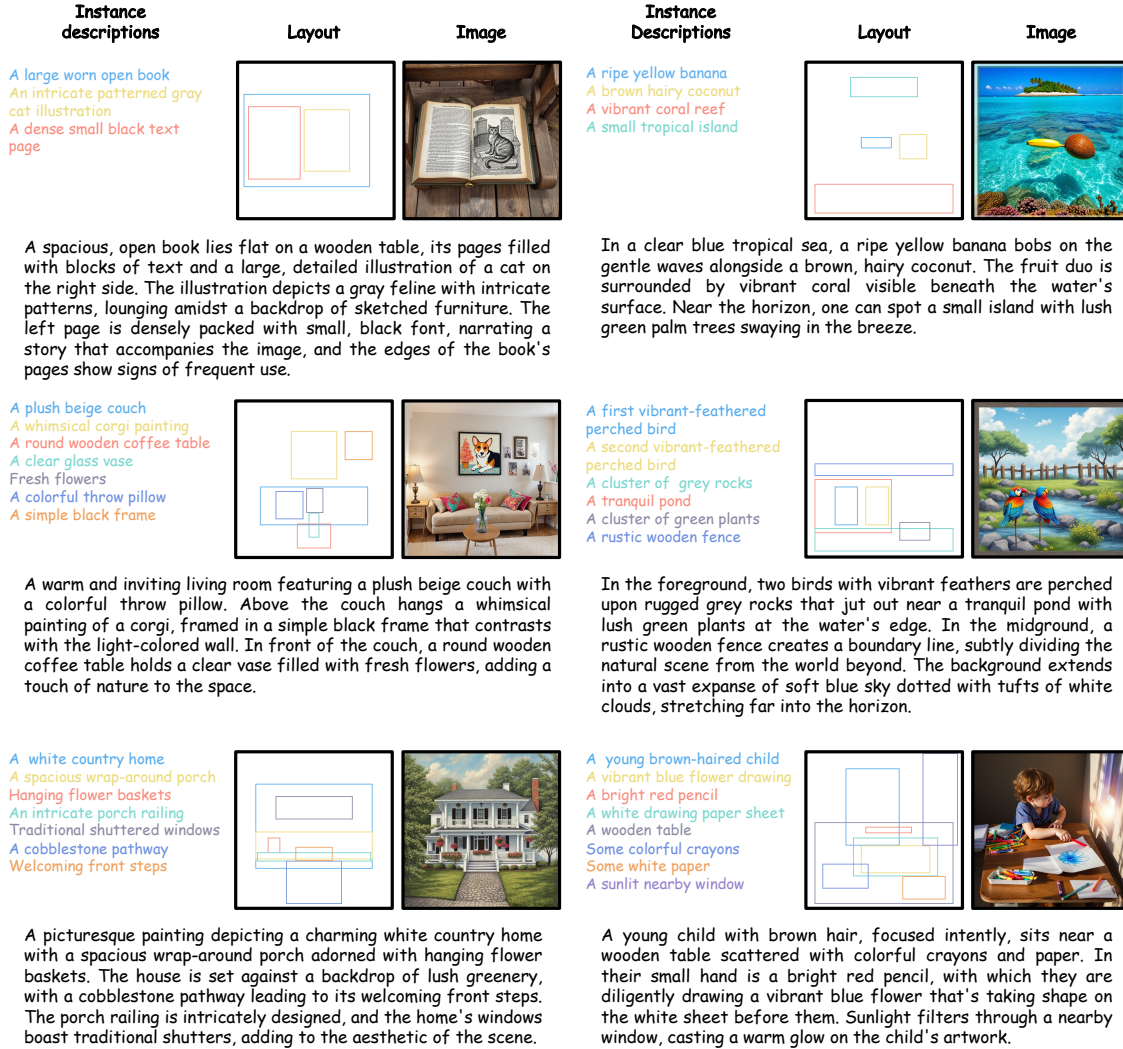


Figure 2. Generative results of our coDrawAgents framework.

### 3.1. Multi-agent Collaboration

To improve the generality of our approach, covering both general text-to-image cases without explicit layouts and more complex cases requiring layout planning, the Interpreter decides whether to enter the *layout-free* or the *layout-aware* mode. In the *layout-free* mode, the Interpreter directly invokes the Painter (a text-to-image model) to generate  $I$  that aligns with  $T$ . In the *layout-aware* mode, the Interpreter first parses  $T$  into attribute-rich object descriptors, ranks them by semantic importance, and groups objects of similar priority for joint generation. The generation then proceeds through a Planner–Checker–Painter loop, iterating once for each semantic priority level. In the  $i^{th}$  iteration, Planner incrementally proposes layouts  $L_i$  for the set of highest-salience objects at a time based on the existing objects grounding and scene visual context rather than the entire scene. The Checker then leverages both the text

and the visual context to validate spatial consistency and objects semantic alignment, followed by layout refinement. The Painter synthesizes the image step by step, incorporating each newly planned object into the evolving canvas, which in turn provides essential visual context for subsequent iterations. After  $N$  layout iterations, the final image  $I$  is produced, closely aligned with the input text  $T$ .

In the layout-aware mode, the Planner adopts a divide-and-conquer strategy by reasoning about objects of the same semantic priority level at a time, which reduces layout complexity. Guided by the evolving canvas from the Painter, it grounds layout predictions in the actual scene rather than imagining them in isolation. The Checker further validates object placement and attribute alignment and refines layouts. This collaborative loop alleviates the burden of spatial planning and yields images that faithfully reflect the input text, particularly in complex object arrangements.

### 3.2. Interpreter, Planner, Checker, and Painter

In this section, we detail the four specialized agents: interpreter, planner, checker, and painter, each of whom is responsible for a distinct role in our generation process. The details of prompts are displayed in Supplementary Material.

**Interpreter** To accurately represent complex scenes as a structured input for our iterative framework, we introduce an Interpreter agent to process  $T$  for the agent system. The Interpreter first infers the relative importance of objects from  $T$ , and then decides whether to invoke the Painter directly for detail fidelity or to activate the multi-agent dialogue for layout-precise generation.

As illustrated in Figure 1(a), in the layout-aware mode the Interpreter decomposes the prompt  $T$  into structured, semantically rich object descriptors and prepares them for downstream planning. Specifically, we leverage large language models (LLMs) with chain-of-thought (CoT) prompting, guided by task instructions. The process follows three steps: (i) **Identify and decompose** the prompt into distinct semantic units; (ii) **Establish priorities** by ranking objects according to their semantic salience and grouping items with the same semantic-level for joint generation; and (iii) **Enrich attributes and background** through CoT-guided queries, yielding detailed descriptors of objects and their relations. The Interpreter then assigns the highest-priority objects of the current round to the iteration, enabling interactive generation with the Planner, Checker, and Painter in the generation loop.

**Planner** The workflow of Planner is shown in Figure 1 (b). At  $i^{\text{th}}$  iteration, Planner aims to plan the layout  $L_i$  of the objects at the  $i^{\text{th}}$  priority ranked by the Interpreter. Motivated by the multimodal chain-of-thought [67], we propose a stepwise visualization chain-of-thought (VCoT) for layout planning. We employ GPT-5 as MLLM for VCoT.

VCoT takes as input the global text prompt  $T$ , the description of the  $i^{\text{th}}$  priority objects, the layouts generated in the previous  $i - 1$  iterations, and the partial image  $I_{i-1}$  rendered by Painter. It also incorporates object grounding, establishing correspondences between textual entities and image regions in  $I_{i-1}$ , which mitigates the inherent insensitivity of LLMs to spatial coordinates [64] and enables reliable object localization.

We formulate our CoT reasoning as three steps: Canvas state analysis, Context-aware planning, and Physics constrain enforcement. In the ‘‘Canvas state analysis’’ stage, guided by the rich visual context of objects grounding, the image  $I_{i-1}$  and other inputs, Planner meticulously analyzes the spatial layout of existing objects to gain a comprehensive visual understanding the current state of the scene. Afterwards, in the ‘‘Context-aware planning’’ stage, based on the existing canvas state, the MLLM planner leverages its embedded world knowledge to reason about the plausible

interactions between the candidate object  $O_i$  and the existing scene composition  $(O_0, \dots, O_{i-1})$ . Further, to maintain physical plausibility and scene coherence, Planner incorporate a ‘‘Physics constrain enforcement’’ module to prompt the MLLM to take physical and contextual constraints into account, which encourages realistic object placement to reflect real-world interactions and prevents issues like floating objects or improbable contacts. Please refer to Supplementary Material for more details on VCoT.

**Checker** At each iteration  $i$ , the Checker performs a two-stage check–then–refine procedure illustrated in Figure 1(c).

**In the first stage**, it analyzes the current proposal  $L_i$  and conducts checking at the object level (size, scale, boundary coverage) and the global level (relative placement among semantically related objects, inter-object relations, and overall spatial plausibility). Based on these assessments, the Checker updates  $L_i$  accordingly.

**In the second stage**, the Checker reviews all previous layouts  $\{L_1, \dots, L_i\}$  to identify cross-object conflicts such as overlaps, occlusion ordering inconsistencies, or scale drift. Guided by the global prompt, it examines and refines the layouts of all generated objects across iterations, allowing corrections of misplacements or other errors from earlier steps. For each identified issue, the Checker applies targeted, step-by-step fixes and propagates these corrections to subsequent layouts. The refined layout is then passed to the Painter for rendering.

**Painter** The Painter supports two modes. In layout-free mode, it invokes a text-to-image (T2I) model to synthesize the image  $I$  directly from the prompt. In layout-aware mode, it uses a layout-to-image (L2I) model conditioned on current layout. Across iterations  $i$ , the Painter incrementally renders the canvas by integrating each newly confirmed object, providing visual context for subsequent steps. At the final iteration, the Painter renders the final image  $I$ .

The models used by Painter is designed to be plug-and-play, allowing any T2I and L2I model to be seamlessly integrated without additional training. In this paper, we use Flux [26] for T2I model and 3DIS [70] for L2I model. Note that our model is designed to be compatible with other, potentially more advanced L2I models, which could further improve our text-to-image generation performance.

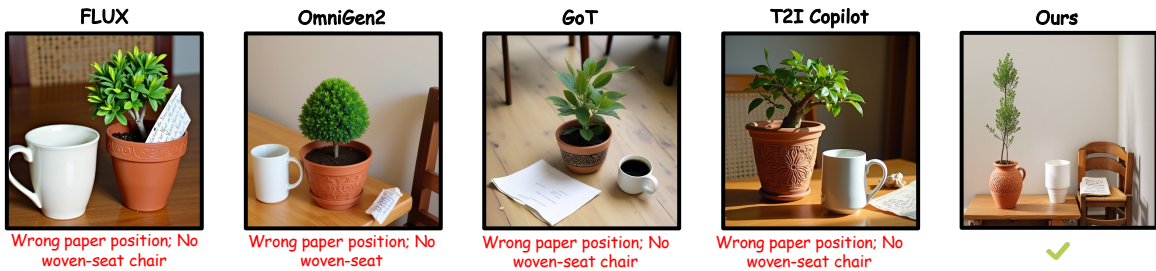
## 4. Experiments

### 4.1. Dataset and Metrics

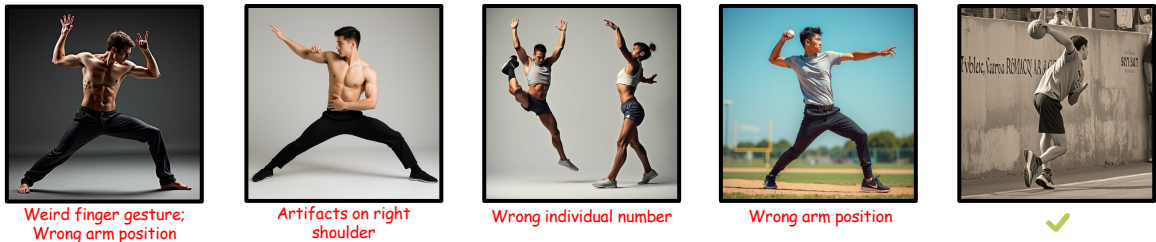
We rigorously evaluated our coDrawAgents framework using two benchmark datasets: GenEval [20] and DPG-Bench [23]. GenEval, a standard for assessing text-to-image generation quality and text-image alignment, provides six metrics including object presence, attribute binding, counting, and spatial relationships. We report the over-

Table 1. Performance comparison on the GenEval [19]. Best results are marked in **bold**. Column names are abbreviated to fit the page.


Model	Single Obj.	Two Obj.	Counting	Colors	Position	Color Attri.	Overall↑
PixArt- $\Sigma$ [6]	0.98	0.50	0.44	0.80	0.08	0.07	0.48
Emu3-Gen [54]	0.98	0.71	0.34	0.81	0.17	0.21	0.54
SDXL [41]	0.98	0.74	0.39	0.85	0.15	0.23	0.55
GoT [15]	0.99	0.69	0.67	0.85	0.34	0.27	0.64
DALL-E 3 [49]	0.96	0.87	0.47	0.83	0.43	0.45	0.67
FLUX.1-dev [26]	0.99	0.81	0.79	0.74	0.20	0.47	0.67
Janus-Pro-1B [8]	0.98	0.82	0.51	0.89	0.65	0.56	0.73
SD3-Medium [14]	0.99	0.94	0.72	0.89	0.33	0.60	0.74
TokenFlow-XL [44]	0.95	0.60	0.41	0.81	0.16	0.24	0.55
UniWorld-V1 [32]	0.99	0.93	0.79	0.89	0.49	0.70	0.80
GPT Image 1 [High] [40]	0.99	0.92	0.85	0.92	0.75	0.61	0.84
<b>coDrawAgents(Ours)</b>	<b>1.00</b>	<b>0.96</b>	<b>0.94</b>	<b>0.97</b>	<b>0.95</b>	<b>0.81</b>	<b>0.94</b>



A small, vibrant green tree sits snugly within a terracotta pot that features intricate patterns etched into its surface. The pot is placed to the left of a simple white ceramic cup with a delicate handle, both resting on a wooden countertop. To the side of the cup is a chair with a woven seat, and the tree in the pot shares this proximity with the chair as well. Perched precariously on the edge of the chair is a crumpled piece of paper, the handwriting upon it partially visible, creating a tableau of everyday items in close association.



A dynamic stance captured in a moment of intense action shows an individual with their legs spread apart for balance. Their right arm is drawn back, poised in a throwing position, with their hand just below the level of their head, ready to launch. The left arm is relaxed and lowered, the elbow bent, and the hand gently resting on the stomach area, creating a counterbalance to the tension in the right arm.



In the center of the composition, there is a neatly arranged stack of three vibrant red cubes, each with a smooth, glossy finish that reflects the ambient light. To the right of this stack, there is a deep blue sphere with a matte texture, providing a stark contrast to the geometric sharpness of the cubes. On the left side, two emerald green cones with a slightly textured surface are positioned, their pointed tips directed upwards, creating a symmetrical balance in the arrangement.

Figure 3. Qualitative comparison with existing methods.

Table 2. Performance comparison on the DPG-Bench [23]. Best results are in **bold**.

Model	Global	Entity	Attribute	Relation	Other	Overall↑
Hunyuan-DiT [30]	84.59	80.59	88.01	74.36	86.41	78.87
PixArt- $\Sigma$ [6]	86.89	82.89	88.94	86.59	87.68	80.54
DALL-E 3 [49]	90.97	89.61	88.39	90.58	89.83	83.50
SD3-Medium [14]	87.90	91.01	88.83	80.70	88.68	84.08
FLUX.1-dev [26]	74.35	90.00	88.96	90.87	88.33	83.84
GoT [15]	83.58	82.16	80.07	87.81	65.25	73.53
T2I-Copilot [5]	87.50	81.74	81.07	86.94	48.28	74.34
OmniGen2 [59]	88.81	88.83	90.18	89.37	90.27	83.57
Emu3-Gen [54]	85.21	86.68	86.84	90.22	83.15	80.60
UniWorld-V1 [33]	83.64	88.39	88.44	89.27	87.22	81.38
BLIP3-o 8B [7]	-	-	-	-	-	81.60
<b>coDrawAgents(Ours)</b>	84.78	90.15	87.55	92.92	84.38	<b>85.17</b>

all GenEval Score, along with its sub-scores, to quantify our model’s performance in each of these aspects. For DPG-Bench, which is designed to evaluate a model’s ability to follow lengthy and dense prompts describing multiple objects with diverse attributes and relationships, we follow its established protocol, using an MLLM to adjudicate the generated images based on a series of questions. We report the overall DPG-Bench score, which is the average score across all prompts, along with scores for its main sub-categories.

## 4.2. Efficiency of coDrawAgents

We report agent-usage statistics on DPG-Bench with 1,074 images. As shown in Table 3, the Interpreter, Planner, Checker, and Painter are each invoked only a few times per generation (1.00, 1.52, 1.62, and 1.95 on average, respectively), which is far fewer than the average number of objects present in a scene (2.79). This efficiency arises because the Interpreter groups objects of the same semantic level, enabling multiple objects to be processed within a single round. This design substantially improves efficiency while maintaining strong performance.

## 4.3. Results and Comparisons

As shown in Figure 2, our coDrawAgents framework is capable of handling a wide range of challenging scenarios, including multi-object compositions, diverse visual styles, highly complex layouts with object interactions, and long descriptive prompts. The generated images remain highly consistent with the input descriptions while maintaining high visual quality. Additional results on DPG-Bench and GenEval, as well as our generative results on COCO [34] and customized prompts demonstrating the framework’s ability to generalize to complex scenes and inter-object relationships, are provided in Supplementary Material.

For quantitative evaluation, we compare coDrawAgents with a broad set of recent state-of-the-art models on GenEval and DPG-Bench (Tables 1 and 2). The baseline

models encompass a broad range of state-of-the-art text-to-image generation approaches, including representative T2I models such as DALL-E 3 [49] and FLUX [26], recent multimodal large models like GPT Image 1 [High] [40] and OmniGen2 [59], layout-aware methods such as GoT [15] that perform explicit bounding box reasoning, and multi-agent frameworks like T2I-Copilot [5]. These models collectively represent diverse paradigms in the field, from direct generation to structured reasoning and interactive agent collaboration. The coDrawAgents achieves the best overall results, demonstrating the effectiveness of our interactive multi-agent design.

For qualitative comparisons, we further compare coDrawAgents with representative methods from three categories: general text-to-image generation FLUX [26], vision-language models OmniGen2 [59], one of the most recent open-source compositional image generation frameworks, GoT [15], which explicitly performs one-shot reasoning over all bounding boxes during generation, and the most recent multi-agent text-to-image generation framework T2I-Copilot [5]. As shown in Figure 3, existing approaches often suffer from misplaced objects, incorrect counts, or attribute artifacts, while our method produces coherent layouts and faithful compositions closely aligned with the textual descriptions.

## 4.4. Ablation Studies

We conduct ablation experiments on DPG-Bench to evaluate the contribution of each component in coDrawAgents. As shown in Table 4 and Figure 4, starting from the layout-free baseline, introducing layout-aware planning allows LLM to explicitly generate layout plans and adopt a divide-and-conquer strategy over multiple objects rather than attempting global planning at once, thereby reducing reasoning complexity. Adding visual context enables Planner to leverage partially generated scene as grounding when placing the next set of semantically prioritized objects, which

Table 3. The average number of agent calls and objects in each DPG-bench image.

Agent	Interpreter	Planner	Checker	Painter	Number of Objects
Avg. calls / generation	1.00	1.52	1.62	1.95	2.79

Table 4. Quantitative ablation study.

Model	Global	Entity	Attribute	Relation	Other	Overall $\uparrow$
Layout-free mode	84.50	84.44	86.15	90.87	75.60	77.60
+ Layout-aware mode	79.94	89.32	87.27	92.37	80.65	82.61
+ Visual context	88.89	88.72	89.32	95.95	66.67	84.51
+ Checker (coDrawAgents)	84.78	90.15	87.55	92.92	84.38	<b>85.17</b>

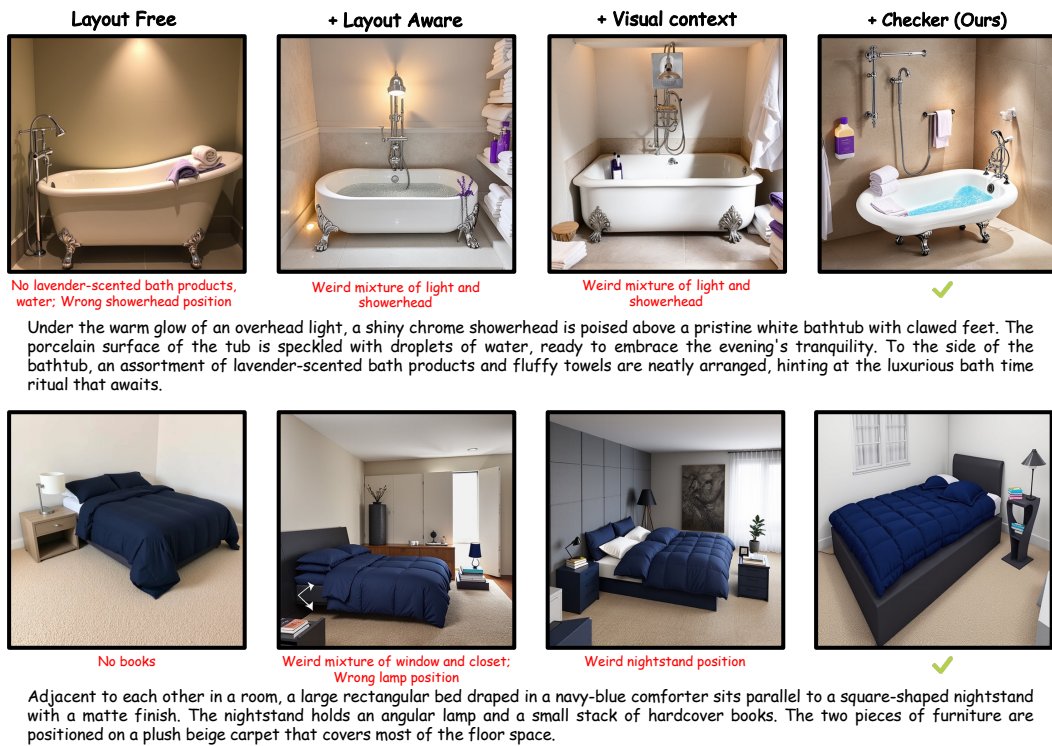


Figure 4. Qualitative ablation results.

enhances spatial coherence. The Checker provides explicit error correction by detecting misplacements and attribute mismatches, further improving entity and attribute faithfulness. Our coDrawAgents model achieves the best overall balance, producing visually coherent scenes with stronger alignment between objects, attributes, and relations.

## 5. Conclusion

In this paper, we introduced coDrawAgents, an interactive multi-agent dialogue framework for compositional text-to-image generation. The coDrawAgents brings together four specialized agents: the Interpreter that decomposes the text

prompt, the Planner that reasons about object layouts, the Checker that verifies spatial and semantic consistency, and the Painter that renders the final image. By coordinating these roles, coDrawAgents addresses key challenges in complex scene generation, including layout reasoning, grounding in evolving visual context, and explicit error correction. Evaluations on GenEval and DPG-Bench show that coDrawAgents achieves state-of-the-art performance, with notable gains in text-image consistency, spatial accuracy, and attribute binding compared with existing approaches. Limitations are discussed in the Supplementary Material.

## Acknowledgement

This work has been supported by National Natural Science Foundation of China (Category C) fund code 62506149, Lingnan University StartUp Grant fund code: 103684, and Faculty Research Grant fund code:106106 and 106119.

## References

- [1] 360CVGroup. Plangen: Layout planning and generation for compositional text-to-image synthesis. <https://github.com/360CVGroup/PlanGen>, 2024. GitHub repository. 3
- [2] James Betker, Gabriel Geng, Atticus Edwards, Ali Farhadi, Hannaneh Hajishirzi, Jong Wook He, Jing Yu Tan, Jiasen Tao, Sharan Vaswani, Licheng Wang, et al. Improving image generation with better captions. *Computer Science*. <https://cdn.openai.com/papers/dall-e-3.pdf>, 2023. 2
- [3] Parth Bhalerao, Mounika Yalamarty, Brian Trinh, and Oana Ignat. Multi-agent multimodal models for multicultural text to image generation. *arXiv preprint arXiv:2502.15972*, 2025. 3
- [4] Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM TOG*, 2023. 1, 2, 3
- [5] Chieh-Yun Chen, Min Shi, Gong Zhang, and Humphrey Shi. T2i-copilot: A training-free multi-agent text-to-image system for enhanced prompt interpretation and interactive generation. *ICCV*, 2025. 3, 7
- [6] Junsong Chen, Chongjian Ge, Enze Xie, Yue Wu, Lewei Yao, Xiaozhe Ren, Zhongdao Wang, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- $\sigma$ : Weak-to-strong training of diffusion transformer for 4k text-to-image generation. In *ECCV*, pages 74–91. Springer, 2024. 6, 7
- [7] Jiuhai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi Zhou, Saining Xie, Silvio Savarese, et al. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*, 2025. 7
- [8] 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. 2, 6
- [9] Jiaxin Cheng, Zixu Zhao, Tong He, Tianjun Xiao, Zheng Zhang, and Yicong Zhou. Rethinking the training and evaluation of rich-context layout-to-image generation. 2024. 2
- [10] Omer Dahary, Or Patashnik, Kfir Aberman, and Daniel Cohen-Or. Be yourself: Bounded attention for multi-subject text-to-image generation. In *ECCV*, pages 432–448. Springer, 2024. 2
- [11] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. Cogview: Mastering text-to-image generation via transformers. In *NeurIPS*, 2021. 2
- [12] Le H. Duong and et al. Onediffusion: One transformer diffusion for image, video, and 3d generation. In *CVPR*, 2025. 2
- [13] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, 2021. 2
- [14] 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 *ICML*, 2024. 6, 7
- [15] Rongyao Fang, Chengqi Duan, Kun Wang, Linjiang Huang, Hao Li, Shilin Yan, Hao Tian, Xingyu Zeng, Rui Zhao, Jifeng Dai, et al. Got: Unleashing reasoning capability of multimodal large language model for visual generation and editing. *arXiv preprint arXiv:2503.10639*, 2025. 3, 6, 7
- [16] Weixi Feng, Wanrong Zhu, Tsu-jui Fu, Varun Jampani, Arjun Akula, Xuehai He, Sugato Basu, Xin Eric Wang, and William Yang Wang. Layoutgpt: Compositional visual planning and generation with large language models. *NeurIPS*, 36:18225–18250, 2023. 1, 2
- [17] Hanan Gani, Shariq Farooq Bhat, Muzammal Naseer, Salman Khan, and Peter Wonka. LLM blueprint: Enabling text-to-image generation with complex and detailed prompts. In *ICLR*, 2024. 2, 3
- [18] Yu Gao, Lixue Gong, Qiushan Guo, Xiaoxia Hou, Zhichao Lai, Fanshi Li, Liang Li, Xiaochen Lian, Chao Liao, Liyang Liu, et al. Seedream 3.0 technical report. *arXiv preprint arXiv:2504.11346*, 2025. 2
- [19] Dhruva Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *NeurIPS*, 36:52132–52152, 2023. 6
- [20] Dhruva Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. 2023. 5
- [21] Meera Hahn, Wenjun Zeng, Nithish Kannan, Rich Galt, Kartikeya Badola, Been Kim, and Zi Wang. Proactive agents for multi-turn text-to-image generation under uncertainty. *ICML*, 2025. 3
- [22] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 1, 2
- [23] Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024. 5, 7
- [24] Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024. 2
- [25] Chengyou Jia, Minnan Luo, Zhuohang Dang, Guang Dai, Xiaojun Chang, Mengmeng Wang, and Jingdong Wang. Ssmg: Spatial-semantic map guided diffusion model for free-form layout-to-image generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2480–2488, 2024. 2

- [26] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 2, 5, 6, 7
- [27] 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 *ECCV*, 2024. 2
- [28] Mingcheng Li, Xiaolu Hou, Ziyang Liu, Dingkan Yang, Ziyun Qian, Jiawei Chen, Jinjie Wei, Yue Jiang, Qingyao Xu, and Lihua Zhang. Mccd: Multi-agent collaboration-based compositional diffusion for complex text-to-image generation. In *CVPR*, 2025. 1, 3
- [29] Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. Gligen: Open-set grounded text-to-image generation. In *CVPR*, 2023. 2
- [30] Zhimin Li, Jianwei Zhang, Qin Lin, Jiangfeng Xiong, Yanxin Long, Xincheng Deng, Yingfang Zhang, Xingchao Liu, Minbin Huang, Zedong Xiao, et al. Hunyuan-dit: A powerful multi-resolution diffusion transformer with fine-grained chinese understanding. *arXiv preprint arXiv:2405.08748*, 2024. 7
- [31] Long Lian, Boyi Li, Adam Yala, and Trevor Darrell. Llm-grounded diffusion: Enhancing prompt understanding of text-to-image diffusion models with large language models. *arXiv preprint arXiv:2305.13655*, 2023. 1, 2, 3
- [32] Bin Lin, Zongjian Li, Xinhua Cheng, Yuwei Niu, Yang Ye, Xianyi He, Shenghai Yuan, Wangbo Yu, Shaodong Wang, Yunyang Ge, et al. Uniworld: High-resolution semantic encoders for unified visual understanding and generation. *arXiv preprint arXiv:2506.03147*, 2025. 2, 6
- [33] Bin Lin, Zongjian Li, Xinhua Cheng, Yuwei Niu, Yang Ye, Xianyi He, Shenghai Yuan, Wangbo Yu, Shaodong Wang, Yunyang Ge, et al. Uniworld: High-resolution semantic encoders for unified visual understanding and generation. *arXiv preprint arXiv:2506.03147*, 2025. 7
- [34] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 7
- [35] Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B Tenenbaum. Compositional visual generation with composable diffusion models. In *ECCV*, 2022. 1, 3
- [36] Zhengyao Lv, Yuxiang Wei, Wangmeng Zuo, and Kwan-Yee K Wong. Place: Adaptive layout-semantic fusion for semantic image synthesis. In *CVPR*, pages 9264–9274, 2024. 2
- [37] Shichao Ma, Yunhe Guo, Jiahao Su, Qihe Huang, Zhengyang Zhou, and Yang Wang. Talk2image: A multi-agent system for multi-turn image generation and editing. *arXiv preprint arXiv:2508.06916*, 2025. 3
- [38] Yuhang Ma, Shanyuan Liu, Ao Ma, Xiaoyu Wu, Dawei Leng, and Yuhui Yin. Hico: Hierarchical controllable diffusion model for layout-to-image generation. *NeurIPS*, 37: 128886–128910, 2024. 2
- [39] Wolf Nuyts, Ruben Cartuyvels, and Marie-Francine Moens. Explicitly representing syntax improves sentence-to-layout prediction of unexpected situations. *TACL*, 2024. 2
- [40] OpenAI. Gpt-image-1, 2025. Available at <https://openai.com/index/introducing-4o-image-generation/>. 2, 6, 7
- [41] 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. 6
- [42] 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*, 2024. 2
- [43] Leigang Qu, Shengqiong Wu, Hao Fei, Liqiang Nie, and Tat-Seng Chua. Layoutllm-t2i: Eliciting layout guidance from llm for text-to-image generation. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 643–654, 2023. 3
- [44] Liao Qu, Huichao Zhang, Yiheng Liu, Xu Wang, Yi Jiang, Yiming Gao, Hu Ye, Daniel K. Du, Zehuan Yuan, and Xinglong Wu. Tokenflow: Unified image tokenizer for multi-modal understanding and generation. In *CVPR*, 2025. 2, 6
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. 2021. 2
- [46] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. 2021. 2
- [47] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 1, 2
- [48] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Seyed Ghasem Kamyar, S Sara Ghasemipour, Burcu Karagol Ayan, S Sajjad Mahdavi, et al. Photorealistic text-to-image diffusion models with deep language understanding. 2022. 2
- [49] Zhan Shi, Xu Zhou, Xipeng Qiu, and Xiaodan Zhu. Improving image captioning with better use of captions. *arXiv preprint arXiv:2006.11807*, 2020. 6, 7
- [50] Jiao Sun, Deqing Fu, Yushi Hu, Su Wang, Royi Rassin, Da-Cheng Juan, Dana Alon, Charles Herrmann, Sjoerd Van Steenkiste, Ranjay Krishna, et al. Dreamsync: Aligning text-to-image generation with image understanding feedback. *arXiv preprint arXiv:2311.17946*, 2023. 2
- [51] Hongchen Tan, Baocai Yin, Kun Wei, Xiuping Liu, and Xin Li. Alr-gan: Adaptive layout refinement for text-to-image synthesis. *TMM*, 25:8620–8631, 2023. 3
- [52] Xie Tianyidan, Rui Ma, Qian Wang, Xiaoqian Ye, Feixuan Liu, Ying Tai, Zhenyu Zhang, Lanjun Wang, and Zili Yi. Anywhere: A multi-agent framework for user-guided, reliable, and diverse foreground-conditioned image generation. In *AAAI*, 2025. 3
- [53] Xingchen Wan, Han Zhou, Ruoxi Sun, Hootan Nakhost, Ke Jiang, Rajarishi Sinha, and Sercan Ö Arık. Maestro: Self-improving text-to-image generation via agent orchestration. *arXiv preprint arXiv:2509.10704*, 2025. 3

- [54] Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiyang Yu, et al. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*, 2024. [6](#), [7](#)
- [55] Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiyang Yu, et al. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*, 2024. [2](#)
- [56] Zhenyu Wang, Aoxue Li, Zhenguo Li, and Xihui Liu. Genartist: Multimodal llm as an agent for unified image generation and editing. In *NeurIPS*, 2024. [3](#)
- [57] Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified multimodal understanding and generation. *arXiv preprint arXiv:2410.13848*, 2024. [2](#)
- [58] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025.
- [59] Chenyuan Wu, Pengfei Zheng, Ruiran Yan, Shitao Xiao, Xin Luo, Yueze Wang, Wanli Li, Xiyan Jiang, Yexin Liu, Junjie Zhou, et al. Omnigen2: Exploration to advanced multimodal generation. *arXiv preprint arXiv:2506.18871*, 2025. [2](#), [7](#)
- [60] Tsung-Han Wu, Long Lian, Joseph E Gonzalez, Boyi Li, and Trevor Darrell. Self-correcting llm-controlled diffusion models. In *CVPR*, 2024. [3](#)
- [61] Jinheng Xie, Yuexiang Li, Yawen Huang, Haozhe Liu, Wentian Zhang, Yefeng Zheng, and Mike Zheng Shou. Boxdiff: Text-to-image synthesis with training-free box-constrained diffusion. In *ICCV*, 2023. [2](#)
- [62] 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. 2024. [3](#)
- [63] Ling Yang, Zhilong Zhang, Zhaochen Yu, Jingwei Liu, Minkai Xu, Stefano Ermon, and Bin Cui. Cross-modal contextualized diffusion models for text-guided visual generation and editing. In *ICLR*, 2024. [2](#)
- [64] Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. In *ICLR*, 2024. [5](#)
- [65] Hui Zhang and et al. Creatilayout: Layout-conditioned text-to-image generation with diffusion models. In *ICCV (ICCV)*, 2025. [2](#)
- [66] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. [1](#), [2](#)
- [67] Zhuosheng Zhang, Aston Zhang, Mu Li, hai zhao, George Karypis, and Alex Smola. Multimodal chain-of-thought reasoning in language models. *TMLR*, 2024. [5](#)
- [68] Guangcong Zheng, Xianpan Zhou, Xuewei Li, Zhongang Qi, Ying Shan, and Xi Li. Layoutdiffusion: Controllable diffusion model for layout-to-image generation. In *CVPR*, 2023. [1](#), [2](#)
- [69] Dewei Zhou, Ji Xie, Zongxin Yang, and Yi Yang. 3dis: Depth-driven decoupled instance synthesis for text-to-image generation. *arXiv preprint arXiv:2410.12669*, 2024.
- [70] Dewei Zhou, Ji Xie, Zongxin Yang, and Yi Yang. 3dis-flux: simple and efficient multi-instance generation with dit rendering. *arXiv preprint arXiv:2501.05131*, 2025. [2](#), [5](#)

# coDrawAgents: A Multi-Agent Dialogue Framework for Compositional Image Generation

## Supplementary Material

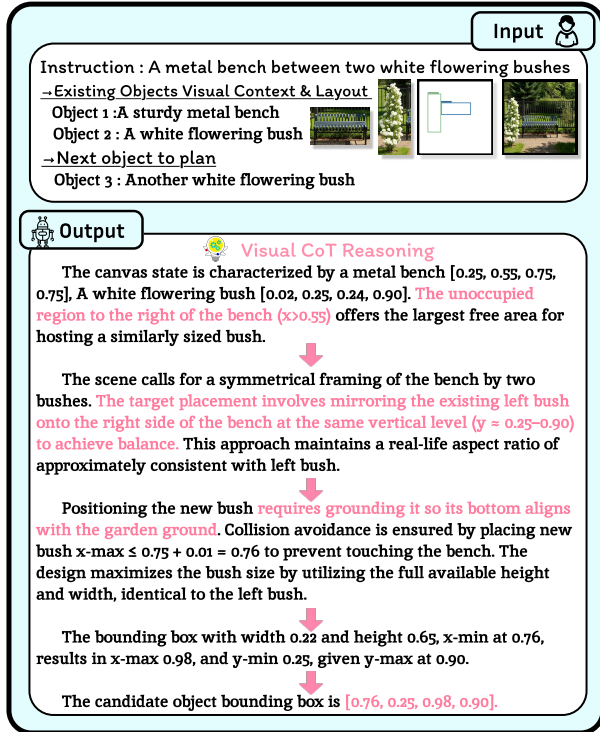


Figure 5. A reasoning example of the VS-COT module in our framework

## 6. Details of Visual CoT (VCoT) Reasoning

Figure 5 illustrates the detailed reasoning steps of our Visual Spatial CoT (VCoT) module. Given the current canvas state and instruction, VCoT performs the following steps:

- Scene parsing.** The module first parses the existing layout, identifying objects and their bounding boxes. For example, a metal bench is located at [0.25, 0.55, 0.75, 0.75], and a white flowering bush is positioned at [0.02, 0.25, 0.34, 0.90].
- Free-space identification.** VCoT then analyzes the unoccupied regions of the canvas. In this case, the right side of the bench ( $x \approx 0.25-0.95$ ) is identified as the feasible area for placing another bush.
- Symmetry reasoning.** To maintain balanced composition, the target placement is chosen to mirror the existing bush on the left, aligning along the same vertical level ( $y \approx 0.90$ ).
- Grounding and constraint enforcement.** The candi-

date bounding box is grounded to the garden floor (bottom aligned), ensuring physical plausibility. Collision constraints are checked to avoid overlap with the bench by setting the bush’s  $x\text{-max} \leq 0.76$ .

- Bounding box refinement.** The box dimensions are adjusted to maximize use of available space while keeping consistency with the left bush. The final bounding box is given as [0.76, 0.25, 0.98, 0.90].

Through these step-by-step spatial reasoning processes, VCoT generates placements that are physically valid, compositionally balanced, and text-aligned.

## 7. Details of Prompts

Figure 6 presents the detailed design of our multi-agent prompting strategy.

## 8. Details of Multi-agent collaboration

At each step, 3DIS first predicts a depth map and then renders the corresponding image using the depth-to-image model. During multi-agent involvement, we iteratively accumulate the depth information from the previous iteration, estimate a refined depth map for the next iteration, and synthesize an updated image based on this new depth. At each iteration, the intermediate image serves only as planning context for the subsequent step, enabling the system to refine local regions while keeping the global structure stable.

Figure 7 illustrates the details of our multi-agent collaboration framework. Given an input prompt, the Interpreter determines the generation mode and decomposes the prompt into instance descriptions. The Planner then proposes an initial layout, which is further checked and refined by the Checkers. Using the refined layouts, the Painter performs iterative planning and synthesis, where each intermediate painting—along with its grounding information—serves as visual context for the next iteration. Finally, the Painter produces the final image.

## 9. Additional results on DPG Bench

Figure 9, 10 and 11 presents six representative cases from DPG-Bench, covering diverse everyday and imaginative scenarios. The prompts feature multiple objects, long textual descriptions, and complex interactions, such as indoor arrangements (books, couches, kitchen scenes), dynamic activities (a surfer riding waves, a man and dog playing with a frisbee), and even global landmarks (Sydney Opera House, Eiffel Tower, Mount Everest). Across

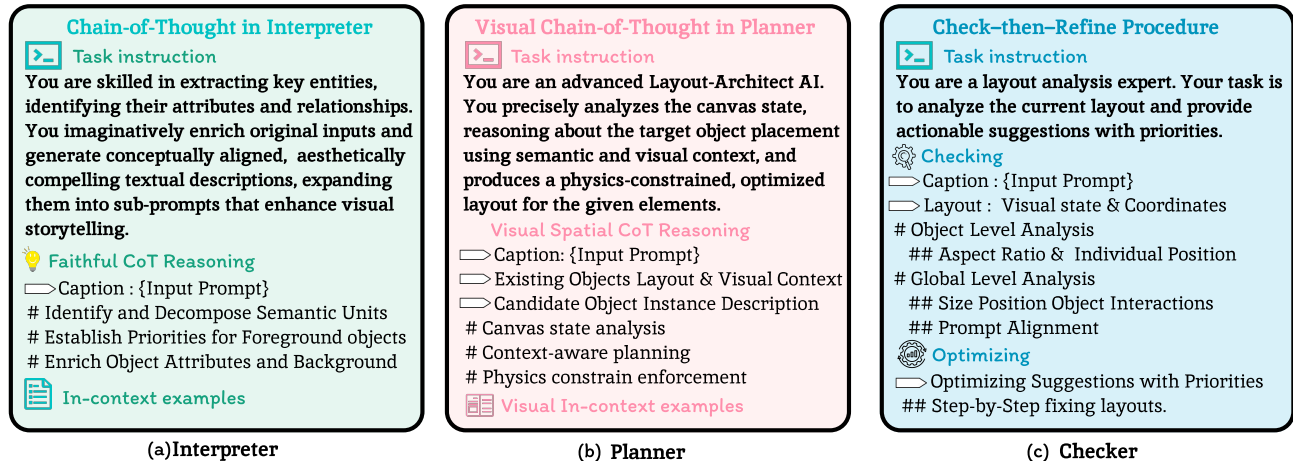


Figure 6. Prompting structure for agents in our framework.

all these cases, coDrawAgents produces images that remain highly faithful to the text, accurately capturing object positions, counts, and attributes while preserving visual coherence and style. These results further confirm that our multi-agent framework scales effectively to the challenging long-context, multi-object compositions posed by DPG-Bench.

Notably, the inclusion of layout-free examples, where instance decomposition and bounding boxes are not available, demonstrates the flexibility of our approach, which can adaptively select between layout-free and layout-aware modes to accommodate varying text prompt requirements.

## 10. Additional results on GenEval

Figures 12 present our generative results on GenEval. coDrawAgents accurately captures object attributes and spatial relationships across diverse scenarios. The framework successfully generates realistic and coherent images, such as a dog positioned to the right of a teddy bear, as well as indoor scenes involving a brown dining table and a white sofa frame.

Notably, coDrawAgents exhibits strong generalization capabilities in handling unconventional or semantically implausible compositions, such as a majestic brown horse placed alongside a leather couch or a computer keyboard, which are rare or unrealistic in real-world contexts. These cases highlight the model’s robustness in following uncommon or imaginative prompts.

## 11. Limitation

While coDrawAgents Dialogue demonstrates significant progress in compositional text-to-image generation, it still has several limitations:

First, the multi-agent system, while beneficial for quality and compositional accuracy, introduces a computational overhead. The framework requires more processing time

compared to single-pass methods due to the multi-agent calls. However, empirical analysis 4.2 shows coDrawAgents Dialogue still achieves competitive inference efficiency, outperforming many existing methods despite its iterative nature. Further optimization of the multi-agent loop remains a key area for future work.

Second, the performance of our Painter is inherently dependent on the underlying T2I and L2I models. This dependency means that limitations of the base models, such as imperfect attribute rendering or biased visual priors, may propagate into coDrawAgents, e.g., “a radish with black skin”. Conversely, it also indicates that coDrawAgents will naturally benefit from future advances in text-to-image and layout-to-image generation.

Third, the Planner and Checker rely on multimodal LLMs for layout reasoning and error detection, making the system susceptible to LLM-specific issues such as hallucination and overconfidence in incorrect layouts. These limitations may lead to invalid object placements or missed corrections, especially in highly compositional or ambiguous prompts. Conversely, it also indicates that coDrawAgents will naturally benefit from future advances in more reliable and grounded LLMs with reduced hallucination tendencies.

Finally, as with most iterative frameworks, coDrawAgents may be affected by error accumulation across iterations. For instance, small placement inaccuracies in early steps can propagate if not fully corrected by the Checker. Nevertheless, our design explicitly mitigates this risk by introducing verification and refinement mechanism, and we observe that the overall error accumulation is significantly lower than in single-pass generation pipelines, as shown in Tables 1 and 2.

**Input Prompt**

A quaint Parisian bistro table, with an ornate metal base, sits on a cobbled street, its surface hosting a classic French kettle with an elegant, sweeping curvilinear profile and a glossy finish that catches the sunlight. Next to it lies a soft, felt French beret, in a deep shade of navy blue, adding a touch of artistic flair to the setting. The backdrop is a bustling Paris afternoon, with the silhouette of the Eiffel Tower looming in the distance, framed by the vibrant green leaves of trees lining the avenue.

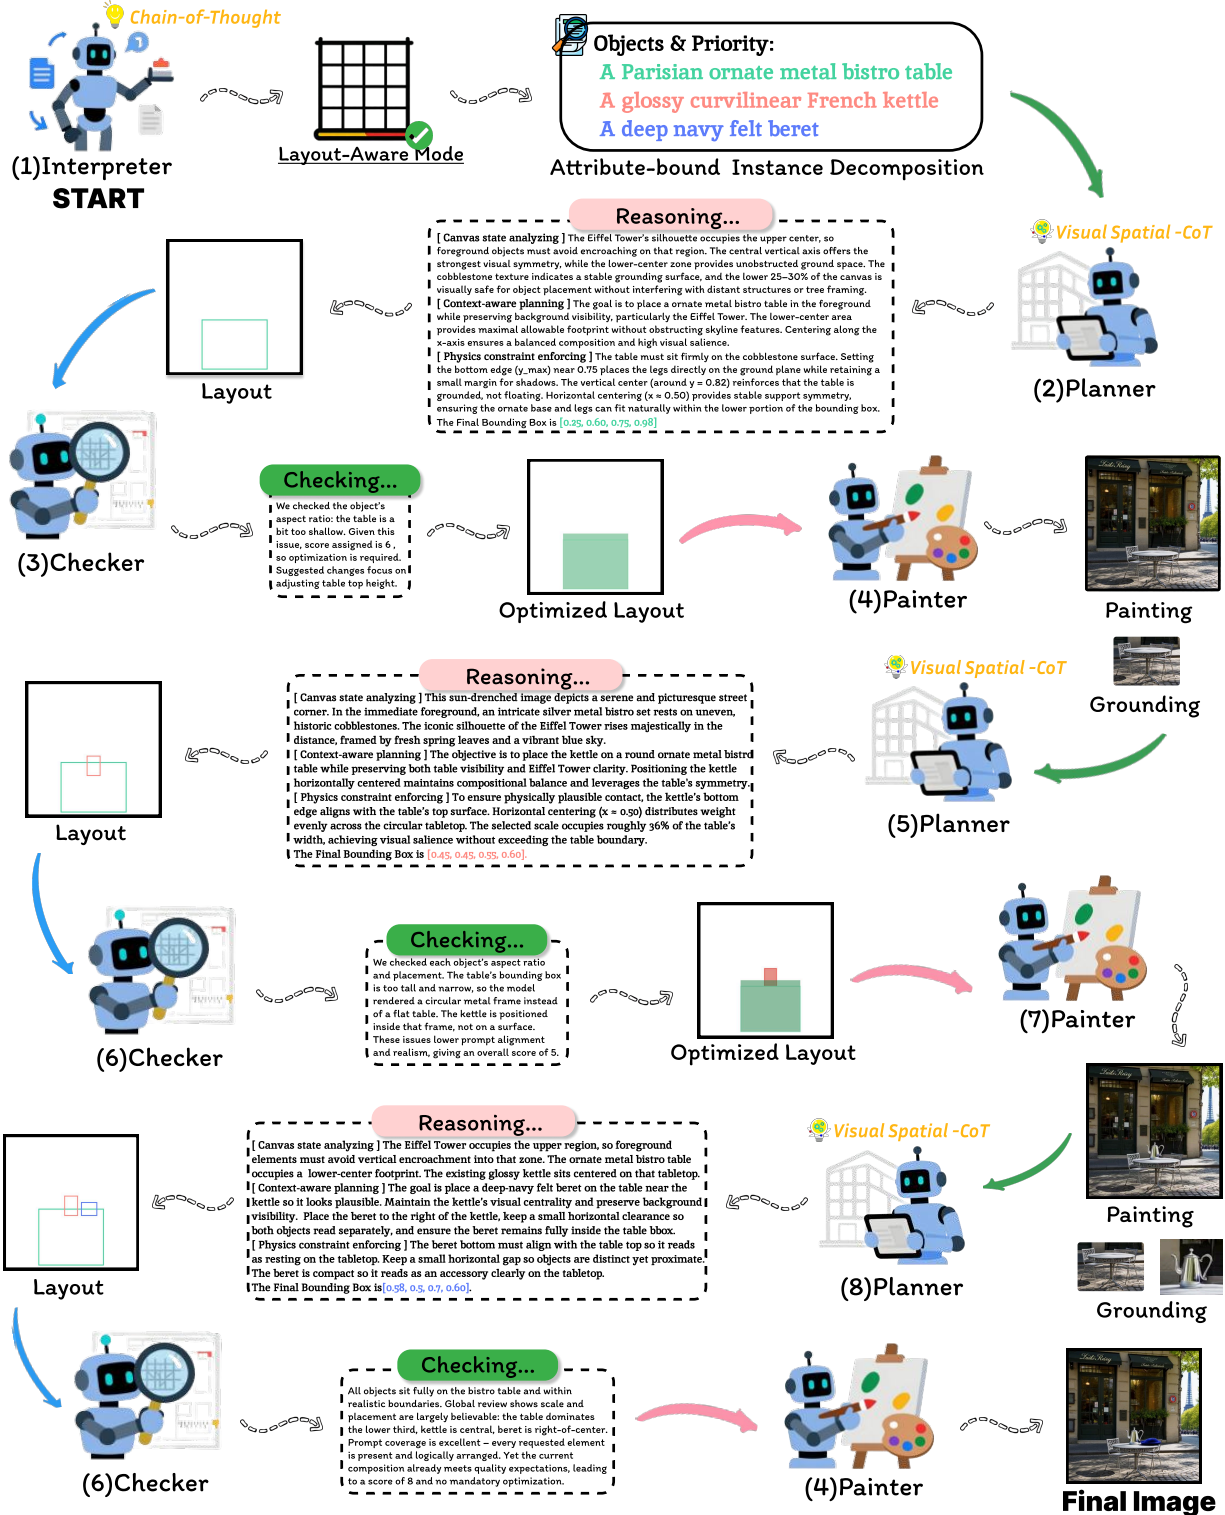



Figure 7. Multi-agent collaboration.

## 12. Qualitative Results



View of kitchen and study area, kitchen with tiled countertop and floor, oven and overhead microwave, spice rack, kettle, and in study area, wood floor, shelves, window, desk with desk lamp, and a padded, rolling office chair.



a whimsical illustration of a small, white baby daikon radish with rosy cheeks and green shoots atop its head, donning a pink, frilly tutu. It is walking a brown, fluffy dog on a red leash, which looks up at the radish with a playful expression. The background features a simple, pastel-colored path that winds through a grassy field.



a large, round orange pumpkin carved with a smiling face, sitting on a wooden table. Inside the hollowed-out pumpkin, a small flickering candle casts a warm glow through the cut-out eyes and mouth. The pumpkin is surrounded by a scattering of fallen autumn leaves.



An imaginative scene where the iconic Sydney Opera House, with its white sail-like shells, sits prominently on the left. To the right, the Eiffel Tower, constructed of intricate iron lattice work, towers over the landscape. Behind both landmarks, the majestic Mount Everest looms, its snow-capped peak piercing the sky.



An abstract oil painting that depicts a chaotic blend of vibrant colors and swirling patterns, giving the impression of a vast, disorienting landscape. The canvas is filled with bold strokes of reds, blues, and yellows that seem to clash and compete for space, symbolizing the complexity and confusion of navigating through life. Amidst the turmoil, a small, indistinct figure appears to be wandering, searching for direction in the overwhelming expanse.



A colorful collection of four cartoon-styled calendars, each uniquely illustrating the essence of a different season. The spring calendar bursts with shades of green and pink, featuring blooming flowers and sprouting leaves. The summer calendar glows with vibrant sun motifs and vivid blue skies. Autumn is represented with warm oranges and browns, showcasing falling leaves and harvest themes. The winter calendar is adorned with soft whites and blues, depicting snowy scenes and cozy fireside images. Each calendar is distinct, yet they all share a whimsical charm that captures the spirit of their respective seasons.



A picturesque outdoor scene featuring a ceramic vase prominently placed to the left of a lush, green lawn. The vase, with its smooth texture and intricate patterns, stands in the foreground, with the expansive, clear blue sky stretching overhead. Beyond the vase, a wooden bench can be seen, slightly obscured by the vase's presence. To the right, a dense, leafy bush rises up against the sky, situated just above a paved street that runs adjacent to the bush.



A playful collection of 2x2 emoji icons, each resembling a vibrant macaron with a distinct facial expression. The top left macaron is a sunny yellow with a beaming smile, while the top right is a fiery red with furrowed brows and an angry scowl. Below them, the bottom left is a bright blue with wide, surprised eyes, and the bottom right is a soft lavender with a tearful, sobbing face. Each of the macaron emojis is whimsically topped with a miniature brown cowboy hat, adding a touch of whimsy to their appearance.



a modern storefront with large glass windows and a bold sign above the entrance that reads 'openai' in sleek, white lettering. The facade is painted in a muted gray, complementing the contemporary design. Inside, through the transparent windows, one can see rows of neatly arranged products and a few customers browsing.



a brightly colored storefront with large, bold letters spelling out 'AwesomePurchase' above the entrance. The shop's window displays are neatly arranged with an array of products, and a small, potted plant sits to the left of the door. The facade of the building is a clean, modern white, contrasting with the vibrant signage.




A collection of nine vibrant VIP sign icons, indicating membership or exclusive status. Each label varies in color, providing a rainbow spectrum from red to violet, and they are designed with a sleek, glossy texture. The symbols feature a simple, bold font that stands out against the solid background, rendering them ideal for vector illustrations where distinction is key.




In the distance, towering black mountains with their peaks blanketed in thick layers of snow stand majestically. Against this dramatic backdrop, a flock of black birds is captured in their dynamic mid-flight, crisscrossing the scene with elegance and energy. Above them, the sky is a tapestry of deep grays clashing with the remnants of serene blue, creating a striking contrast that defines the horizon.


Figure 8. Qualitative results.



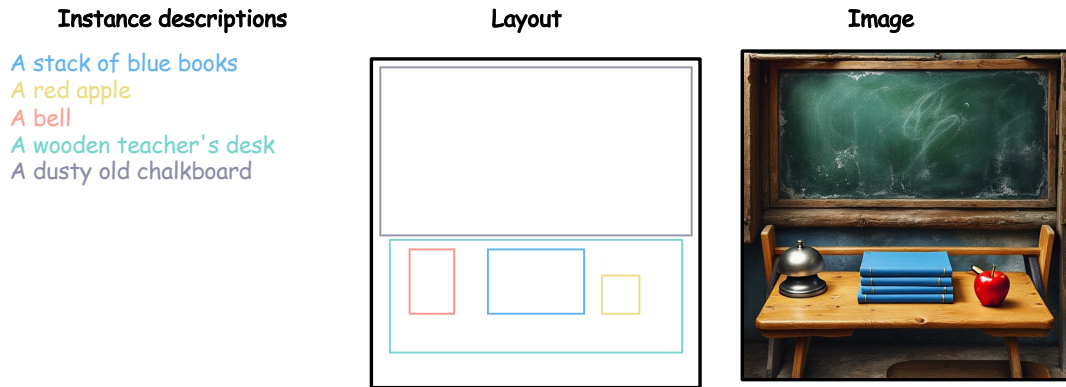
A man, his arm across the woman next to him, stands in a blandly colored kitchen area, in front of a black-rimmed window, next to a counter with a microwave, plates, with and without food, and wine bottles.




a photo of a brown dining table and a white suitcase




A sectional sofa in a front room with a bookshelf and mirror with two floor lamps and two vases of flowers on either side.




A dusty old chalkboard fills the background of this image, which features a wooden teacher's desk that has blue books, a red apple, and a bell placed on top.




A beach scene captures a man, clad in blue and white striped swim shorts, standing barefoot on the warm, golden sand. To his side, a playful black and white dog, with its gaze fixed on an object in the sky, waits in anticipation. Suspended in the air above them is a spinning white frisbee, creating a dynamic moment of play and excitement just off the coast, where the gentle waves lap at the shore.




A photo of someone's living room complete with a bookshelf full of DVDs, two leather chairs, a flat screen TV, fireplace, and a overly large decorative clock.



an individual balancing on a bright yellow surfboard, riding the crest of an ocean wave. parallel to the shore, a series of tall buildings stand in close proximity to one another, creating a dense urban skyline. the closest building has a reflective glass facade, while the one alongside it features beige brickwork.



An old-fashioned kitchen setting with a cast-iron kettle and a ceramic teapot sitting atop a rough-hewn, wooden table that bears the marks and patina of age. The kettle's metallic surface has a dull gleam, reflecting the warm ambient light, while the teapot, adorned with a floral pattern, adds a touch of nostalgia to the setting. In the background, there is a window with curtains partially drawn, allowing for a soft natural light to fill the room. Nearby, a woven basket filled with dried flowers accentuates the rustic charm of the cozy interior.




Inside the microwave sits a clear glass bowl, filled to the brim with scoops of colorful ice cream with visible flecks of vanilla beans. The microwave's interior light casts a warm glow on the ice cream, which threatens to melt if the door were to remain closed for long. It's an odd place for a cold dessert that's usually served at a chilly temperature to avoid its creamy contents from turning into a soupy mess. The microwave is positioned on a countertop, surrounded by assorted kitchen gadgets and a spice rack full of various seasonings.

**Instance descriptions**

The first hot dog  
The second hot dog  
The third hot dog

**Layout**




**Image**




a photo of three hot dogs

A majestic brown horse  
A sleek computer keyboard



a photo of a horse and a computer keyboard

A large brown bear  
A wooden dining table



a photo of a dining table and a bear