

# A Unified Agentic Framework for Evaluating Conditional Image Generation

Jifang Wang<sup>1</sup>, Xue Yang<sup>2</sup>, Longyue Wang<sup>2</sup>, Zhenran Xu<sup>1</sup>, Yiyu Wang<sup>2</sup>,  
Yaowei Wang<sup>1</sup>, Weihua Luo<sup>2</sup>, Kaifu Zhang<sup>2</sup>, Baotian Hu<sup>1\*</sup>, Min Zhang<sup>1</sup>

<sup>1</sup>Harbin Institute of Technology (Shenzhen)

<sup>2</sup>Alibaba International Digital Commerce

{23S151116, xuzhenran}@stu.hit.edu.cn

{wangyaowei, hubaotian, zhangmin2021}@hit.edu.cn

{shali.yx, menshuan.wyy, wanglongyue.wly, weihua.luowh, kaifu.zkf}@alibaba-inc.com

## Abstract

Conditional image generation has gained significant attention for its ability to personalize content. However, the field faces challenges in developing task-agnostic, reliable, and explainable evaluation metrics. This paper introduces **CIGeVAL**, a unified agentic framework for comprehensive evaluation of conditional image generation tasks. CIGeVAL utilizes large multimodal models (LMMs) as its core, integrating a multi-functional toolbox and establishing a fine-grained evaluation framework. Additionally, we synthesize evaluation trajectories for fine-tuning, empowering smaller LMMs to autonomously select appropriate tools and conduct nuanced analyses based on tool outputs. Experiments across seven prominent conditional image generation tasks demonstrate that CIGeVAL (GPT-4o version) achieves a high correlation of 0.4625 with human assessments, closely matching the inter-annotator correlation of 0.47. Moreover, when implemented with 7B open-source LMMs using only 2.3K training trajectories, CIGeVAL surpasses the previous GPT-4o-based state-of-the-art method. Case studies on GPT-4o image generation highlight CIGeVAL’s capability in identifying subtle issues related to subject consistency and adherence to control guidance, indicating its great potential for automating evaluation of image generation tasks with human-level reliability<sup>1</sup>.

## 1 Introduction

Recent advances in large-scale text-to-image (T2I) generative models have enabled the creation of images based on text prompts as well as reference images, i.e. *conditional* image generation (Kumari et al., 2023; Ruiz et al., 2023; Li et al., 2023b; He et al., 2024). The field is evolving at an unprecedented pace with an increasing number of tasks and models being introduced. Among these,

\*Corresponding author.

<sup>1</sup>Our code and models are publicly available at <https://github.com/HITSz-TMG/Agentic-CIGeVAL>.

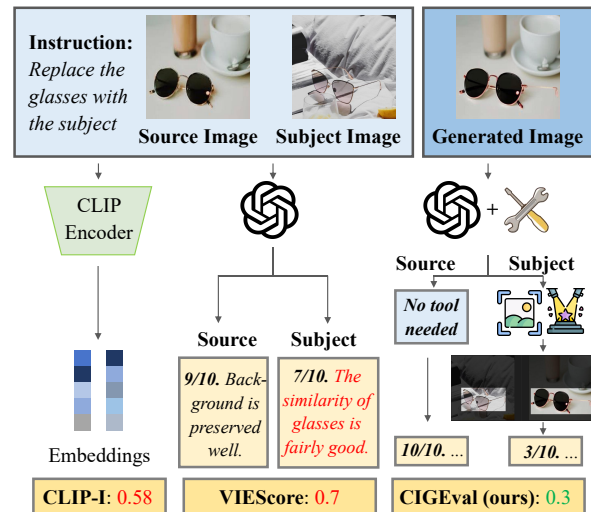


Figure 1: An example of subject-driven image editing with human-annotated low scores. Both traditional metrics and GPT-4o-based VIEScore assign high scores. By integrating GPT-4o with tools, CIGeVAL, our agentic evaluation framework, highlights the glasses object in both images, and finds their different shapes and designs, thereby reaching the correct score. “Source” and “Subject” means “source image” and “subject image”.

text-guided image generation is particularly popular (Ramesh et al., 2022; Chen et al., 2025; Yuan et al., 2025). Expanding beyond text, a diverse set of conditions have been employed to steer the diffusion process: text-guided image editing (Brooks et al., 2023), mask-guided image editing (runwayml, 2023), subject-driven image generation and editing (Chen et al., 2023; Guo et al., 2024), multi-concept image composition (Kumari et al., 2023) and control-guided image generation (Qin et al., 2023; Zhang and Agrawala, 2023).

Despite the growing number of generative models being developed, a significant challenge persists in effectively evaluating AI-synthesized images (Peng et al., 2024). Current evaluation metrics have the following three limitations: (1) **Task-specific:** Traditional metrics are narrowly focused

and cannot be generalized across different tasks. For example, LPIPS (Zhang et al., 2018) measures the perceptual similarity of a pair of images, while CLIP-Score (Hessel et al., 2021) measures the text alignment of one single image. (2) **Limited explainability**: Assigning a single score to a generated image without the reasoning process fails to offer a comprehensive evaluation. Images can be assessed on multiple dimensions, such as prompt adherence and concept preservation (Fu et al., 2023b). (3) **Lack of human alignment**: Traditional metrics like DINO (Caron et al., 2021) and CLIP (Radford et al., 2021) often result in huge discrepancies from humans, caused by their image similarity measurement nature. Even based on the powerful large multimodal model (LMM) GPT-4o, as shown in Figure 1, the current state-of-the-art VIEScore (Ku et al., 2024) struggles to capture subtle image nuances and shows low correlation with human judgment in various image editing tasks.

To address these issues, we propose **CIGeVAL**, an autonomous LMM-based agent framework for evaluating conditional image generation. This agent framework can integrate the advanced GPT-4o model (OpenAI, 2023) and open-source models (e.g., Qwen2.5-VL (Wang et al., 2024)). Our work is driven by two primary motivations: (1) developing autonomous evaluation agents capable of making independent decisions and judgments without human assistance; (2) enabling relatively smaller models to efficiently perform complex evaluations. To achieve this, we make three key technical contributions. First, we extend the LMM’s capability to detect and emphasize subtle differences between highly similar images by curating a versatile toolbox, in contrast to previous methods that relied solely on the perceptual capabilities of LMMs. Second, we establish a fine-grained evaluation framework, including task decomposition, tool selection and analysis. Third, we synthesize instruction data based on evaluation trajectories for fine-tuning the LMM, where we first employ GPT-4o to execute the stages and then filter the trajectories that align with human evaluations.

Experiments on the well-established ImagenHub benchmark (Ku et al., 2023) show that, when using GPT-4o as the underlying LMM, CIGeVAL achieves the state-of-the-art performance across all 7 tasks. It achieves an average Spearman correlation of 0.4625 with human raters, closely matching the human-to-human correlation of 0.47. The primary improvements are observed in tasks involv-

ing multiple conditions, such as control-guided image generation and multi-concept image composition, where previous evaluation metrics struggle. Using only 2.3K filtered evaluation trajectories for tuning, CIGeVAL, leveraging 7B open-source LMMs, demonstrates performance surpassing previous GPT-4o-based state-of-the-art methods. Further ablation study shows the importance of each tool and the robustness of our framework. In addition, we conduct a preliminary case study on GPT-4o’s image generation. CIGeVAL assigns scores closely aligned with human annotations and effectively detects subtle flaws in 4o-generated images, especially in tasks involving multiple input images and adherence to specific control signals (e.g., Canny edges, OpenPose). These results suggest that CIGeVAL has substantial promise for achieving human-level performance in assessing synthetic images.

Our main contributions are as follows:

- We introduce CIGeVAL, an LMM-based evaluation agent designed to assess various conditional image generation tasks. Our approach is characterized by its human-aligned, explainable, and unified evaluation method, setting it apart from previous metrics.
- We evaluate CIGeVAL across 7 conditional image generation tasks, demonstrating that CIGeVAL, based on GPT-4o, outperforms all existing baselines and achieves a high correlation with human annotators, closely mirroring the human-to-human correlation.
- We fine-tune open-sourced 7B LMMs and significantly improve their evaluation performance, surpassing previous GPT-4o-based state-of-the-art method.

## 2 Related Work

### 2.1 Conditional Image Generation

Diffusion models have gained wide attention in AI research for image synthesis (Ho et al., 2020; Dhariwal and Nichol, 2021). In recent years, several new models (Kumari et al., 2023; Ruiz et al., 2023; Li et al., 2023b; Zhang and Agrawala, 2023) have been developed to introduce controllable conditions in image generation. Prevalent tasks in this domain include text-to-image generation (Saharia et al., 2022; Rombach et al., 2022; stability.ai, 2023) (known as text-guided image generation),

inpainting (Avrahami et al., 2022; Lugmayr et al., 2022) (referred to as mask-guided image editing) and text-guided image editing (Brooks et al., 2023; Couairon et al., 2022; Wu and la Torre, 2023). Recent works have proposed new tasks, such as subject-driven image generation and editing (Gal et al., 2022; Ruiz et al., 2023; Li et al., 2023b) to inject one specific subject into a synthesized image, and multi-concept image composition (Kumari et al., 2023; Liu et al., 2023; Ding et al., 2024), which allows multiple specific subjects into the synthesized image. Additionally, control-guided image generation (Zhang and Agrawala, 2023; Qin et al., 2023; Guo et al., 2024) allows additional conditions alongside the text prompt to guide the image synthesis. Our work employs LMM-based agents to assess all of these discussed tasks.

## 2.2 Synthetic Image Evaluation

A variety of metrics have been introduced to assess AI-generated images. For example, the CLIP score (Hessel et al., 2021) and BLIP score (Li et al., 2022) are commonly used to measure the alignment between the generated image and the text prompt. Metrics like LPIPS (Zhang et al., 2018) and DreamSim (Fu et al., 2023b) focus on assessing perceptual similarity. LLMscore (Lu et al., 2023) and HEIM-benchmark (Lee et al., 2023) assess text-to-image models on multiple fine-grained aspects, including toxicity and safety. However, these metrics predominantly focused on text-to-image generation and remain narrow in scope. There is a noticeable lack of effective automatic metrics for other conditional image generation tasks, such as subject-driven image generation and image editing (Ruiz et al., 2023; Li et al., 2023b; Peng et al., 2024). Consequently, some research work (Denton et al., 2015; Isola et al., 2017; Meng et al., 2021; Chen et al., 2023; Sheynin et al., 2023) rely heavily on human evaluation. This dependence highlights the need for more unified, interpretable, and reliable automatic evaluation methods in the field. Our work seeks to bridge this gap by developing an autonomous agentic evaluation framework that closely aligns with human judgment.

## 2.3 Large Multimodal Models as Evaluators

Motivated by the explorations of large language model (LLM)-based evaluators in natural language processing (Zheng et al., 2023; Dubois et al., 2023; Fu et al., 2023a; Cheng et al., 2024b), large multimodal models (LMMs) have been utilized to evalu-

ate responses in visual question answering (Chen et al., 2024a; Xu et al., 2024; Li et al., 2025). In the realm of image evaluation, the GPT-4 series has demonstrated promising capabilities, particularly in assessing text-image alignment (Zhang et al., 2023b; Li et al., 2024). However, these models are not without limitations. A comprehensive study on GPT-4o’s vision abilities have revealed mistakes in fine-grained image evaluation tasks (Yang et al., 2023), such as failing to accurately distinguish differences between similar images (Ku et al., 2024). To address these shortcomings, we enhance the capabilities of LMMs by integrating a versatile set of image analysis and editing tools, and by adopting an agentic framework to improve the evaluation of AI-generated images.

## 3 CIGeVAL

In this section, we introduce **CIGeVAL**, our LMM-based agentic framework designed for evaluating conditional image generation. First, we define seven conditional image generation tasks that are the focus of our study (Sec. 3.1), and then design a multi-functional toolbox (Sec. 3.2). Next, we introduce our fine-grained evaluation framework (Sec. 3.3). Finally, we synthesize high-quality trajectory data to fine-tune open-source LMMs (Sec. 3.4).

### 3.1 Task Definition

To build a unified and explainable evaluation metric, we define the image evaluation problem as shown in Equation 1. The function  $f_{\text{eval}}$  takes as input an instruction  $I$ , a synthesized image  $O$ , and a set of conditions  $C^*$  (e.g. text prompt, subject image, background image, canny-edge, etc). The function  $f_{\text{eval}}$  should produce the intermediate rationale in natural language before generating the final score according to the instruction  $I$ :

$$f_{\text{eval}}(I, O, C^*) = (\text{rationale}, \text{score}) \quad (1)$$

Following Ku et al. (2023), we focus on seven primary conditional image generation tasks, each with different sets of conditions  $C^*$ :

- **Text-guided Image Generation:**  $C^* = [p]$ , where  $p$  is a text prompt. The objective is to generate an image that aligns with the text description.
- **Mask-guided Image Editing:**  $C^* = [p, I_{\text{mask}}, I_{\text{src}}]$ , where  $I_{\text{mask}}$  is a binarized mask and  $I_{\text{src}}$  is a source image. The aim is to modify  $I_{\text{src}}$  in the masked area according to  $p$ .








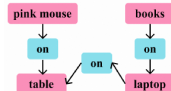
Tool	Argument		Output	Purpose
 <b>Grounding</b>	Image	Target Entity a yellow alarm clock	[219, 261, 337, 370] 	Obtain the coordinates of regions from the Image corresponding to the Target Entity.
 <b>Highlight</b>	Image	Region [324, 281, 381, 497]	Edited Image 	Highlight the listed Region in the Image.
 <b>Difference</b>	Image 1	Image 2	[128, 109, 164, 150] 	Identify the pixel difference between Image 1 and Image 2.
 <b>Scene Graph</b>	Image	A dict about objects and attributes 		Analyzed by LMMs, a structured description of the objects, their attributes, and the relationships in Image.

Table 1: Tools used in our CIGeVAL framework.

• **Text-guided Image Editing:**  $C^* = [p, I_{src}]$ . This task is similar to Mask-guided Image Editing but does not provide a mask. The model must identify the region to edit automatically.

• **Subject-driven Image Generation:**  $C^* = [p, S]$ , where  $S$  is the image of a specific subject. The aim is to generate an image that reflects  $p$  in relation to the subject  $S$ .

• **Subject-driven Image Editing:**  $C^* = [S, p, I_{src}]$ , where  $I_{src}$  is a source image, and  $S$  is the subject reference. The goal is to replace the subject in  $I_{src}$  with  $S$ .

• **Multi-concept Image Composition:**  $C^* = [S_1, S_2, p, I_{src}]$ , where  $S_1$  and  $S_2$  are images of two subjects. The task is to combine these to create a new image according to  $p$ .

• **Control-guided Image Generation:**  $C^* = [I_{control}, p]$ , where  $I_{control}$  is a control signal, such as a depth map, canny edge, or bounding box. The aim is to generate an image that follows these low-level visual cues.

In this paper, following previous work (Mañas et al., 2024; Lin et al., 2024), we investigate the semantic consistency of generated images with the above conditions.

### 3.2 Toolbox

Evaluating image generation that involves multiple conditions can be challenging. Drawing inspiration from prior research (Cheng et al.,

2024a; Zhang et al., 2024) we have developed a multi-functional toolbox, including Grounding, Difference, Highlight and Scene Graph. Each tool is designed to target specific aspects of image analysis or editing, and outputs either a modified image or textual information. Detailed descriptions of each tool can be found in Table 1.

Specifically, we implement Grounding with GroundingDino (Liu et al., 2024). Scene Graph uses the same prompting method as CCoT (Mitra et al., 2024) based on GPT-4o. This tool can also function effectively with other open-source LMMs (refer to Sec. 4.3). To assist LMMs in detecting subtle differences between edited images, the Difference tool compares the pixels of two images and identifies the locations of the variations. The Highlight tool emphasizes selected regions by reducing the pixel values of areas outside the highlighted zone to 1/4 of their original values, thereby darkening these areas and accentuating the highlighted region. This tool is typically used after the Grounding and Difference tools have provided the region coordinates.

### 3.3 Framework

In our approach, we conceptualize the image evaluation process as an agent task. As shown in Figure 2, the core of CIGeVAL is a well-instructed LMM, which autonomously utilizes tools to assess a wide range of conditional image generation tasks.



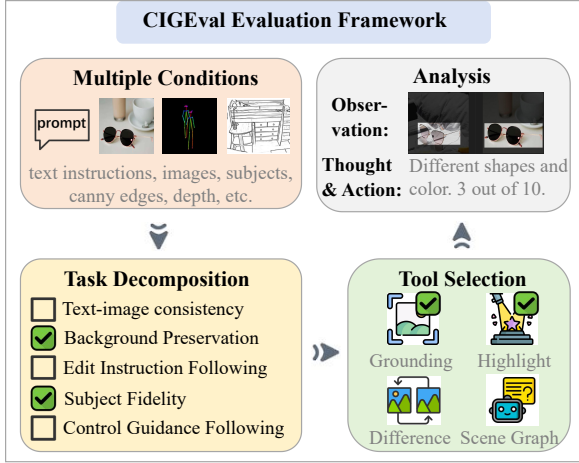


Figure 2: The evaluation process of CIGEval regarding the example in Figure 1. CIGEval autonomously selects appropriate tools for each decomposed sub-task, and then conducts fine-grained analyses based on the observed tool outputs.

The prompts used in this framework are listed in Appendix A.

Specifically, we adopt a divide-and-conquer scheme for evaluating images generated under multiple conditions. For example, in the subject-driven image editing task in Figure 1, the desired synthesized image will incorporate the object from a subject reference while maintaining the background of the source image. Therefore, we break down each evaluation task into several fine-grained sub-questions from the listed below: (1) Is the image generation following the prompt? (2) Is the image editing following the instruction? (3) Is the image performing minimal edit without changing the background? (4) Is the object in the Image following the provided subject? (5) Is the image following the control guidance? Then, for each sub-question, CIGEval selects the most suitable tool from its toolbox, focusing on the specific aspect of evaluation. For example, Grounding and Highlight are utilized for comparing specific regions of the image, while Scene Graph evaluates background preservation and the extent of over-editing. With these intermediate results, the LMM analyzes the tool outputs and assigns scores in the ReAct format (Yao et al., 2023), ranging from 0 to 10, which are normalized to the [0.0, 1.0] range for comparison with human ratings. These fine-grained scores are aggregated through:

$$O = \min(\alpha_1, \dots, \alpha_i) \quad (2)$$

where  $\alpha_i$  represents one of the sub-scores. In

accordance with the setting of Ku et al. (2023), we assume each sub-score weights the same and used min operation to emphasize the importance of meeting all criteria without exception.

### 3.4 Agent Tuning

Previous research has primarily relied on closed-source LMMs to address agentic tasks, mainly due to their superior abilities in tool calling and instruction following (Chen et al., 2024b; Song et al., 2024; Zeng et al., 2024; Xu et al., 2023). As evidenced in Table 3, open-source models significantly underperform compared to GPT-4o. To bridge this gap and empower smaller LMMs as effective evaluators, we aim to perform supervised fine-tuning on 7B models to integrate agentic capabilities into them.

To curate high-quality trajectory data, we employ GPT-4o to carry out the evaluation process in Section 3.3. The process begins by providing GPT-4o with evaluation instructions and corresponding images. At each turn, the agent receives an *observation*, formulates plans and thoughts as *thought*, and invokes relevant tools through *action*. The tool outputs serve as new observations for the subsequent turn. By iterating the above process, we can construct a complete evaluation trajectory including the initial instruction, intermediate steps (i.e., observations, thoughts, actions), and the final scoring result. To ensure the quality of these trajectories, we exclude samples where the discrepancy between predicted scores and human evaluation scores exceeds 0.3. Using 60% of the ImagenHub data, we ultimately gather 2,274 high-quality trajectories for supervised fine-tuning.

Using this structured trajectory data, we perform supervised fine-tuning on Qwen2-VL-7B-Instruct and Qwen2.5-VL-7B-Instruct (Wang et al., 2024). Formally, each sample’s evaluation trajectory is represented as  $\langle o_0, t_1, a_1, \dots, o_{n-1}, t_n, a_n, o_n \rangle$ , where  $o_i$ ,  $t_i$ , and  $a_i$  denote the observation, thought, and action at each turn ( $i$ ) respectively. Specifically,  $o_0$  refers to the initial observation consisting of the evaluation instructions and accompanied images, and  $o_n$  denotes the final score. At each turn, based on the preceding trajectory  $c_i = \langle o_0, t_1, a_1, \dots, o_{i-1} \rangle$ , the agent aims to generate thought  $t_i$  and action  $a_i$ . During the fine-tuning process, we only compute the cross-entropy loss

Task	# Instr.	# Images	# Ratings
Text-guided IG (5 models)	197	985	2955
Mask-guided IE (4 models)	179	716	2148
Text-guided IE (8 models)	179	1432	4296
Subject-driven IG (4 models)	150	600	1800
Subject-driven IE (3 models)	154	462	1386
Multi-concept IC (3 models)	102	306	918
Control-guided IG (2 models)	150	300	900
<b>Sum of 7 tasks</b>	<b>1111</b>	<b>4801</b>	<b>14403</b>

Table 2: Statistics of ImagenHub: the number of instructions, evaluated models, synthesized images, and human ratings used in this paper. The abbreviations “IG”, “IE” and “IC” stand for “Image Generation”, “Image Editing” and “Image Composition” respectively.

for  $t_i$  and  $a_i$  while  $c_i$  is masked:

$$\mathcal{L} = -\log \sum_{i=1}^n \Pr(t_i, a_i | c_i). \quad (3)$$

## 4 Experiments

### 4.1 Evaluation Benchmark

ImagenHub (Ku et al., 2023) is a standardized benchmark for evaluating conditional image generation models with human raters. The statistics for ImagenHub are presented in Table 2. This large-scale benchmark covers 7 mainstream tasks, 29 models, 4.8K synthesized images, and 14.4K human ratings, making it suitable for calculating correlation scores between automatic evaluation metrics and human raters. A list of 29 evaluated models can be found in Appendix B. Implementation details are in Appendix C.

Each image was assessed by three human raters according to the guidelines of the defined task, and a final score in the range [0.0, 1.0] was reported for the average score. Images are scored in two aspects: (1) Semantic Consistency assesses how well the generated image aligns with the given conditions, such as prompts and subject tokens, ensuring coherence and relevance to the specified task criteria. (2) Perceptual Quality evaluates the extent to which the generated image appears visually authentic and conveys a sense of naturalness. In this work, we focus on the **Semantic Consistency** score, leaving the exploration of Perceptual Quality for future research.

### 4.2 Existing Auto-metrics

Here we list some prominent automatic metrics:

- **CLIP-Score** (Hessel et al., 2021): This metric computes the average cosine similarities

between prompt and generated image CLIP embeddings, making it a popular choice for assessing image-text alignment.

- **LPIPS** (Zhang et al., 2018) measures the similarity between two images in a manner that aligns with human perception.
- **CLIP-I** (Gal et al., 2022) calculates the average pairwise cosine similarities between CLIP embeddings of generated and source images.
- **DINO** (Ruiz et al., 2023) is computed by the mean cosine similarity computed between the DINO embeddings of ViT-S/16 (Caron et al., 2021) for both synthesized and source images.
- **VIESCORE** (Ku et al., 2024) prompts large multimodal models to evaluate generated images in an explainable and fine-grained manner. Based on GPT-4o, it currently represents the state-of-the-art across all seven tasks on ImagenHub.

### 4.3 Main Experiments

For all presented correlations, we applied Fisher Z-transformation to estimate the average Spearman correlation  $\in [-1, 1]$  across models and tasks.

**Metric-to-Human (M-H) correlations.** In Table 3, we present the correlations across all tasks, utilizing different backbone models. When using GPT-4o as the underlying LMM, CIGEval achieves the state-of-the-art performance across all 7 tasks. It achieves an average Spearman correlation of 0.4625 with human raters, closely matching the human-to-human correlation. The primary improvements are observed in tasks involving multiple conditions, such as control-guided image generation and multi-concept image composition, where previous evaluation metrics struggle.

When the underlying LMM is replaced with different open-source models, CIGEval consistently outperforms VIESCORE. However, the performance of open-source models is still poor and falls significantly behind GPT-4o. Therefore, we perform agent tuning on these models as described in Sec. 3.4 and report their improved performance in Sec. 4.4. Overall, the experiment demonstrates that CIGEval outperforms VIESCORE across a variety of image editing and generation tasks, consistently maintaining its edge even when different underlying LMMs are used.

Method	Text-guided IG	Mask-guided IE	Text-guided IE	Control-guided IG	Subject-driven IG	Subject-driven IE	Multi-concept IC	Avg.
Human Raters	0.5044	0.5390	0.4230	0.5443	0.4780	0.4887	0.5927	0.4700
CLIPScore	-0.0817	-	-	-	-	-	-	-
LPIPS	-	-0.1012	0.0956	0.3699	-	-	-	-
DINO	-	-	-	-	0.4160	0.3022	0.0979	-
CLIP-I	-	-	-	-	0.2961	0.2834	0.1512	-
<i>LLaMA3-LLaVA-NeXT-8B</i>								
VIESCORE	0.1948	0.2037	0.0363	0.4001	0.1592	-0.1153	0.1308	0.1432
CIGEval	0.1420	0.2843	0.0744	0.4487	0.2891	-0.0699	0.3704	0.2164
<i>Qwen2.5-VL-7B-Instruct</i>								
VIESCORE	0.4218	0.3555	0.0252	0.2836	0.4264	-0.0452	0.3328	0.2516
CIGEval	0.4347	0.4685	0.2567	0.3752	0.4374	0.4863	0.3251	0.3780
<i>GPT-4o</i>								
VIESCORE	0.4989	0.5421	0.4062	0.4972	0.4806	0.4800	0.4516	0.4459
CIGEval	<b>0.5027</b>	<b>0.5465</b>	<b>0.4090</b>	<b>0.5402</b>	<b>0.4930</b>	<b>0.5185</b>	<b>0.4931</b>	<b>0.4625</b>

Table 3: Spearman correlation scores across 7 conditional image generation tasks with different LMMs as backbone. The abbreviations “IG”, “IE” and “IC” stand for “Image Generation”, “Image Editing” and “Image Composition” respectively. “-” means not applicable.

Configuration	Avg.
<b>CIGEval</b>	0.7262
w/o Grounding	0.6376
w/o Difference	0.7020
w/o Scene Graph	0.6471
Scene Graph with <i>Qwen2.5-VL-7B-Instruct</i>	0.7120
Scene Graph with <i>Qwen2.5-VL-70B-Instruct</i>	0.7311

Table 4: Ablation study regarding tools in CIGEval (GPT-4o) with different configurations.

**Ablation study.** To assess the necessity of each tool in CIGEval, we conducted an ablation study detailed in Table 4. Since Highlight is often accompanied with Grounding and Difference, we do not perform specific ablation on Highlight. The study shows that the complete CIGEval configuration achieves the highest average score of 0.7262. When each tool is removed, a noticeable drop is observed, underscoring the necessity of every tool for effective performance.

On the other hand, when the implementation of the Scene Graph was switched from GPT-4o to an open-source model, the evaluation results remained high with minimal variation. In fact, when replaced with the Qwen2.5-VL-70B, the performance improved further, showcasing the robustness of our agentic framework. Overall, the ablation study underscores that each tool in the CIGEval configuration is useful, and their collective integration is crucial for achieving superior performance.

#### 4.4 CIGEval with Agent Tuning

The experimental results in Table 5 show the performance of CIGEval after agent tuning. Despite utilizing 7B open-source LMMs as the underlying model, Qwen2-VL-7B-Instruct and Qwen2.5-VL-7B-Instruct demonstrate a 76% and 34% improvement in correlation after fine-tuning, respectively. With only 2,274 filtered evaluation trajectories, the fine-tuned 7B models surpass the previous state-of-the-art VIEScore based on GPT-4o. This demonstrates the data efficiency of agent tuning and the importance of synthetic data quality.

#### 4.5 Case Study

To demonstrate the effectiveness of our CIGEval framework and the importance of each tool, we present a subject-driven image editing example in Figure 1, a text-guided image editing example in Figure 3, and a multi-concept image composition example in Figure 5. In the first and third example, by directly prompting in VIEScore, GPT-4o struggles to compare the similarity of specific objects between two images. By grounding and highlighting the focused object (i.e., glasses and flowers), GPT-4o can find the difference in shapes and colors within our framework. In the second example, when discussing the background preservation aspect, VIEScore considers the over-editing small. However, in our framework, CIGEval first calls the Scene Graph tool to get an overall composition of the edited image, and then finds out a

Method	Text-guided IG	Mask-guided IE	Text-guided IE	Control-guided IG	Subject-driven IG	Subject-driven IE	Multi-concept IC	Avg.
Previous SOTA	0.3081	0.3167	0.4649	0.5246	0.7105	0.4694	0.5616	0.4458
<i>Qwen2.5-VL-7B-Instruct</i>								
VIESCORE	0.3457	0.0158	0.0086	0.2395	0.1837	0.0967	0.4388	0.1876
CIGEval	0.1890	0.1418	0.4586	0.3130	0.4485	0.5216	0.4496	0.3455
<b>+ Tuning</b>	0.4609	0.2796	0.5916	0.5876	0.4659	0.5458	0.5778	0.4631
<i>Qwen2-VL-7B-Instruct</i>								
VIESCORE	0.3699	-0.1398	0.1024	0.3420	0.1553	0.0682	0.5129	0.1989
CIGEval	0.3054	0.1974	0.1438	0.2615	0.5096	0.1226	0.5035	0.2840
<b>+ Tuning</b>	0.4099	0.5272	0.3846	0.6096	0.6445	0.5975	0.6691	<b>0.4997</b>

Table 5: Spearman correlations across 7 tasks with VIESCORE and CIGEval based on open-source small LMMs. “Previous SOTA” here means VIESCORE based on GPT-4o.

#### Text-Guided Image Editing: *Add water and flowers in the tub*



Original Image



Edited Image

**VIEScore:** There is some noticeable overediting with additional flowers placed around the bathroom unrelated to the tub, **but the original scene remains mostly recognizable. 8** out of 10.

**CIGEval (ours): Thought:** To evaluate the degree of overediting, I need to assess the overall composition of the image to see if the added elements disrupt the original look and feel of the scene.

**Action:** Use “Scene Graph” Tool

**Observation:** `{ "name": "smaller tub", "attributes": { "color": "white", "shape": "oval", "size": "small" } }`

**Thought & Action:** The edited image features **significant overediting** compared to the original. The original image contains one round bathtub and no flowers, while the edited version introduces a larger round tub and **a smaller oval tub that were not present initially. 2** out of 10.

Figure 3: Case study of a text-guided image editing example with a low human annotation score.

newly-added tub based on tool outputs, successfully arriving at the correct score. These examples have shown CIGEval’s ability to autonomously select appropriate tools and reach correct conclusions based on the observation, which makes CIGEval a better evaluator than VIEScore.

**Preliminary study with GPT-4o image generation.** Recently GPT-4o image generation has attracted wide attention. As shown in Figure 4, we extend CIGEval with an additional OCR tool and find that our framework assigns appropriate scores to 4o-generated images on OpenAI’s official website<sup>2</sup>. Furthermore, we test GPT-4o image generation on ImagenHub’s various tasks and report the CIGEval scores and human annotations (averaged between two annotators). We have the following three findings: (1) CIGEval assigns scores that closely align with human annotations, and effectively detects subtle flaws in GPT-4o-generated images. (2) GPT-4o excels at tasks involving a

single image as input, such as text-guided image generation and editing, as well as subject-driven image generation, as shown in Figures 9 to 11. (3) GPT-4o struggles with complex tasks that require multiple images and control signals. For example, the subjects in Figures 7 and 8 are not accurately replicated, showing unintended changes in color or shapes. Moreover, consistent with the findings of Yan et al. (2025), we observe that GPT-4o tends to favor a color palette dominated by yellow, orange, and warm lighting, as exemplified by the pot in Figure 7 and the man in the rearview mirror in Figure 11. Additionally, the control guidances (e.g., canny edges, OpenPose) are not strictly followed, as seen in Figure 6.

## 5 Conclusion

In this paper, we propose CIGEval, a unified, explainable and agentic framework for image evaluation across seven popular conditional image evaluation tasks. CIGEval utilizes large multimodal models (LMMs) at its core to autonomously se-

<sup>2</sup><https://openai.com/index/introducing-4o-image-generation/>



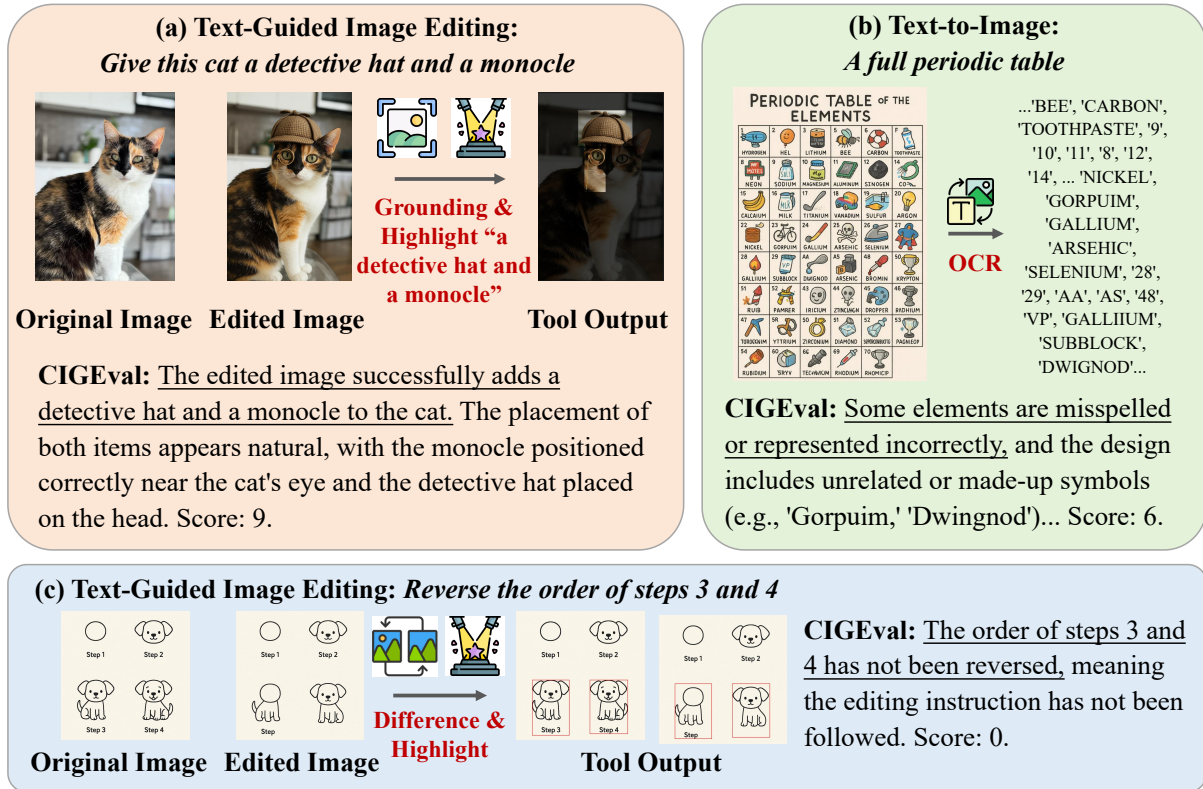


Figure 4: Case study of GPT-4o’s image generation. Examples are adapted from OpenAI’s official website.

lect tools for fine-grained evaluation. Experiments show that, when using GPT-4o as the backbone model, CIGEval surpasses achieves a high correlation of 0.4625 with human raters, closely matching the human-to-human correlation of 0.47. Additionally, we have synthesized 2,274 high-quality evaluation trajectories to incorporate agentic capabilities into smaller LMMs. After agent tuning, the 7B LMMs surpass the performance of the previous state-of-the-art method based on the closed-source GPT-4o. These experimental findings and case studies on GPT-4o image generation suggest that CIGEval holds substantial promise for replicating human-like performance in evaluating synthetic images.

## Limitations

Although CIGEval improves the correlation between automatic image evaluators and human raters, there are certain limitations to our approach. First, when using closed-source models’ APIs, such as GPT-4o, there is a risk that AI-generated images resembling real people or photographs may be rejected by GPT-4o for evaluation, potentially affecting the framework’s robustness. Second, our experiments primarily focus on evaluating images’

consistency with multiple conditions, leaving the assessment of perceptual quality for future research. Due to the lack of a more comprehensive benchmark for conditional image generation, we synthesized tuning data and conducted experiments exclusively on ImagenHub. Expanding our experiments to more text-to-image generation and text-based image editing datasets (Peng et al., 2024; Hui et al., 2024) could be beneficial. Finally, the training process currently utilizes only correct trajectory data and discards failed trajectory data. In the future, we aim to refine the CIGEval framework to include a broader range of tasks and leverage failed data for contrastive training of the model.

## Acknowledgments

We want to thank anonymous reviewers for their helpful comments. This work is jointly supported by grants: Natural Science Foundation of China (No. 62422603), Guangdong Basic and Applied Basic Research Foundation (No. 2023A1515110078), and Shenzhen Science and Technology Program (No. ZDSYS20230626091203008).

## References

- Omri Avrahami, Dani Lischinski, and Ohad Fried. 2022. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18208–18218.
- Omer Bar-Tal, Dolev Ofri-Amar, Rafail Fridman, Yoni Kasten, and Tali Dekel. 2022. Text2live: Text-driven layered image and video editing. In *European Conference on Computer Vision*, pages 707–723.
- Tim Brooks, Aleksander Holynski, and Alexei A. Efros. 2023. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. 2021. [Emerging properties in self-supervised vision transformers](#). *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9630–9640.
- Dongping Chen, Ruoxi Chen, Shilin Zhang, Yaochen Wang, Yinyao Liu, Huichi Zhou, Qihui Zhang, Yao Wan, Pan Zhou, and Lichao Sun. 2024a. Mllm-as-a-judge: assessing multimodal llm-as-a-judge with vision-language benchmark. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org.
- Wenhu Chen, Hexiang Hu, Yandong Li, Nataniel Rui, Xuhui Jia, Ming-Wei Chang, and William W Cohen. 2023. Subject-driven text-to-image generation via apprenticeship learning. *arXiv preprint arXiv:2304.00186*.
- Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. 2025. [Janus-pro: Unified multimodal understanding and generation with data and model scaling](#). *Preprint*, arXiv:2501.17811.
- Zehui Chen, Kuikun Liu, Qiuchen Wang, Wenwei Zhang, Jiangning Liu, Dahua Lin, Kai Chen, and Feng Zhao. 2024b. [Agent-FLAN: Designing data and methods of effective agent tuning for large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 9354–9366, Bangkok, Thailand. Association for Computational Linguistics.
- Chuanqi Cheng, Jian Guan, Wei Wu, and Rui Yan. 2024a. [From the least to the most: Building a plug-and-play visual reasoner via data synthesis](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4941–4957, Miami, Florida, USA. Association for Computational Linguistics.
- Xiaoxue Cheng, Junyi Li, Xin Zhao, Hongzhi Zhang, Fuzheng Zhang, Di Zhang, Kun Gai, and Ji-Rong Wen. 2024b. [Small agent can also rock! empowering small language models as hallucination detector](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 14600–14615, Miami, Florida, USA. Association for Computational Linguistics.
- Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. 2022. Diffedit: Diffusion-based semantic image editing with mask guidance. In *The Eleventh International Conference on Learning Representations*.
- deep floyd.ai. 2023. [If by deepfloyd lab at stabilityai](#).
- Emily L Denton, Soumith Chintala, Rob Fergus, et al. 2015. Deep generative image models using a laplacian pyramid of adversarial networks. *Advances in neural information processing systems*, 28.
- Prafulla Dhariwal and Alexander Nichol. 2021. [Diffusion models beat gans on image synthesis](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 8780–8794. Curran Associates, Inc.
- Ganggui Ding, Canyu Zhao, Wen Wang, Zhen Yang, Zide Liu, Hao Chen, and Chunhua Shen. 2024. Freecustom: Tuning-free customized image generation for multi-concept composition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Alpaca-farm: A simulation framework for methods that learn from human feedback](#). *Preprint*, arXiv:2305.14387.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023a. [Gptscore: Evaluate as you desire](#). *ArXiv*, abs/2302.04166.
- Stephanie Fu, Netanel Tamir, Shobhita Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and Phillip Isola. 2023b. [Dreamsim: Learning new dimensions of human visual similarity using synthetic data](#). *Preprint*, arXiv:2306.09344.
- Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit Haim Bermano, Gal Chechik, and Daniel Cohen-or. 2022. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *The Eleventh International Conference on Learning Representations*.
- Jing Gu, Yilin Wang, Nanxuan Zhao, Tsu-Jui Fu, Wei Xiong, Qing Liu, Zhifei Zhang, He Zhang, Jianming Zhang, HyunJoon Jung, et al. 2023. Photoswap: Personalized subject swapping in images. *arXiv preprint arXiv:2305.18286*.
- Zinan Guo, Yanze Wu, Zhuowei Chen, Lang chen, Peng Zhang, and Qian HE. 2024. [PuLID: Pure and lighting ID customization via contrastive alignment](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

- Zecheng He, Bo Sun, Felix Juefei-Xu, Haoyu Ma, Ankit Ramchandani, Vincent Cheung, Siddharth Shah, Anmol Kalia, Harihar Subramanyam, Alireza Zareian, Li Chen, Ankit Jain, Ning Zhang, Peizhao Zhang, Roshan Sumbaly, Peter Vajda, and Animesh Sinha. 2024. [Imagine yourself: Tuning-free personalized image generation](#). *Preprint*, arXiv:2409.13346.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. Clipscore: A reference-free evaluation metric for image captioning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7514–7528.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. [Denoising diffusion probabilistic models](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2021. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Mude Hui, Siwei Yang, Bingchen Zhao, Yichun Shi, Heng Wang, Peng Wang, Yuyin Zhou, and Cihang Xie. 2024. [Hq-edit: A high-quality dataset for instruction-based image editing](#). *Preprint*, arXiv:2404.09990.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134.
- Max Ku, Dongfu Jiang, Cong Wei, Xiang Yue, and Wenhui Chen. 2024. [VIEScore: Towards explainable metrics for conditional image synthesis evaluation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12268–12290, Bangkok, Thailand. Association for Computational Linguistics.
- Max Ku, Tianle Li, Kai Zhang, Yujie Lu, Xingyu Fu, Wenwen Zhuang, and Wenhui Chen. 2023. Imagenhub: Standardizing the evaluation of conditional image generation models. *arXiv preprint arXiv:2310.01596*.
- Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. 2023. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1931–1941.
- Tony Lee, Michihiro Yasunaga, Chenlin Meng, Yifan Mai, Joon Sung Park, Agrim Gupta, Yunzhi Zhang, Deepak Narayanan, Hannah Benita Teufel, Marco Bellagente, et al. 2023. Holistic evaluation of text-to-image models. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Dongxu Li, Junnan Li, and Steven CH Hoi. 2023a. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *arXiv preprint arXiv:2305.14720*.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR.
- Tianle Li, Max Ku, Cong Wei, and Wenhui Chen. 2023b. Dreamedit: Subject-driven image editing. *arXiv preprint arXiv:2306.12624*.
- Yunxin Li, Zhenyu Liu, Zitao Li, Xuanyu Zhang, Zhenran Xu, Xinyu Chen, Haoyuan Shi, Shenyuan Jiang, Xintong Wang, Jifang Wang, Shouzheng Huang, Xinpeng Zhao, Borui Jiang, Lanqing Hong, Longyue Wang, Zhuotao Tian, Baoxing Huai, Wenhan Luo, Weihua Luo, Zheng Zhang, Baotian Hu, and Min Zhang. 2025. [Perception, reason, think, and plan: A survey on large multimodal reasoning models](#). *Preprint*, arXiv:2505.04921.
- Yunxin Li, Longyue Wang, Baotian Hu, Xinyu Chen, Wanqi Zhong, Chenyang Lyu, Wei Wang, and Min Zhang. 2024. [A comprehensive evaluation of gpt-4v on knowledge-intensive visual question answering](#). *Preprint*, arXiv:2311.07536.
- Zhiqiu Lin, Deepak Pathak, Baiqi Li, Jiayao Li, Xide Xia, Graham Neubig, Pengchuan Zhang, and Deva Ramanan. 2024. [Evaluating text-to-visual generation with image-to-text generation](#). *Preprint*, arXiv:2404.01291.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, and Lei Zhang. 2024. [Grounding dino: Marrying dino with grounded pre-training for open-set object detection](#). In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part XLVII*, page 38–55, Berlin, Heidelberg. Springer-Verlag.
- Zhiheng Liu, Yifei Zhang, Yujun Shen, Kecheng Zheng, Kai Zhu, Ruili Feng, Yu Liu, Deli Zhao, Jingren Zhou, and Yang Cao. 2023. Cones 2: Customizable image synthesis with multiple subjects. *arXiv preprint arXiv:2305.19327*.
- Yujie Lu, Xianjun Yang, Xiujun Li, Xin Eric Wang, and William Yang Wang. 2023. [LLMScore: Unveiling the power of large language models in text-to-image synthesis evaluation](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Andreas Lugmayr, Martin Danelljan, Andrés Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. 2022. [Repaint: Inpainting using denoising diffusion probabilistic models](#). *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11451–11461.



- Oscar Mañas, Pietro Astolfi, Melissa Hall, Candace Ross, Jack Urbanek, Adina Williams, Aishwarya Agrawal, Adriana Romero-Soriano, and Michal Drozdal. 2024. [Improving text-to-image consistency via automatic prompt optimization](#). *Preprint*, arXiv:2403.17804.
- Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. 2021. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*.
- Chancharik Mitra, Brandon Huang, Trevor Darrell, and Roei Herzig. 2024. [Compositional Chain-of-Thought Prompting for Large Multimodal Models](#). In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14420–14431, Los Alamitos, CA, USA. IEEE Computer Society.
- Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. 2023. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6038–6047.
- OpenAI. 2023. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- openjourney.ai. 2023. [Openjourney is an open source stable diffusion fine tuned model on midjourney images](#).
- Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. 2023. Zero-shot image-to-image translation. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11.
- Yuang Peng, Yuxin Cui, Haomiao Tang, Zekun Qi, Runpei Dong, Jing Bai, Chunrui Han, Zheng Ge, Xiangyu Zhang, and Shu-Tao Xia. 2024. [Dreambench++: A human-aligned benchmark for personalized image generation](#). *Preprint*, arXiv:2406.16855.
- Can Qin, Shu Zhang, Ning Yu, Yihao Feng, Xinyi Yang, Yingbo Zhou, Huan Wang, Juan Carlos Niebles, Caiming Xiong, Silvio Savarese, et al. 2023. Uni-control: A unified diffusion model for controllable visual generation in the wild. *arXiv preprint arXiv:2305.11147*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3.
- 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.
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510.
- runwayml. 2023. [Stable diffusion inpainting](#).
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494.
- Shelly Sheynin, Adam Polyak, Uriel Singer, Yuval Kirstain, Amit Zohar, Oron Ashual, Devi Parikh, and Yaniv Taigman. 2023. Emu edit: Precise image editing via recognition and generation tasks. *arXiv preprint arXiv:2311.10089*.
- Yifan Song, Weimin Xiong, Xiutian Zhao, Dawei Zhu, Wenhao Wu, Ke Wang, Cheng Li, Wei Peng, and Sujian Li. 2024. [AgentBank: Towards generalized LLM agents via fine-tuning on 50000+ interaction trajectories](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 2124–2141, Miami, Florida, USA. Association for Computational Linguistics.
- stability.ai. 2023. [Stable diffusion xl](#).
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. [Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution](#). *Preprint*, arXiv:2409.12191.
- Chen Henry Wu and Fernando De la Torre. 2023. A latent space of stochastic diffusion models for zero-shot image editing and guidance. In *ICCV*.
- Zhenran Xu, Senbao Shi, Baotian Hu, Longyue Wang, and Min Zhang. 2024. [MultiSkill: Evaluating large multimodal models for fine-grained alignment skills](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1506–1523, Miami, Florida, USA. Association for Computational Linguistics.
- Zhenran Xu, Senbao Shi, Baotian Hu, Jindi Yu, Dongfang Li, Min Zhang, and Yuxiang Wu. 2023. [Towards reasoning in large language models via multi-agent peer review collaboration](#). *Preprint*, arXiv:2311.08152.



- Zhiyuan Yan, Junyan Ye, Weijia Li, Zilong Huang, Shenghai Yuan, Xiangyang He, Kaiqing Lin, Jun He, Conghui He, and Li Yuan. 2025. [Gpt-imgeval: A comprehensive benchmark for diagnosing gpt4o in image generation](#). *Preprint*, arXiv:2504.02782.
- Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. 2023. [The dawn of Imms: Preliminary explorations with gpt-4v\(ision\)](#). *Preprint*, arXiv:2309.17421.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. [React: Synergizing reasoning and acting in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Huaying Yuan, Ziliang Zhao, Shuting Wang, Shitao Xiao, Minheng Ni, Zheng Liu, and Zhicheng Dou. 2025. [FineRAG: Fine-grained retrieval-augmented text-to-image generation](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 11196–11205, Abu Dhabi, UAE. Association for Computational Linguistics.
- Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao Liu, Yuxiao Dong, and Jie Tang. 2024. [AgentTuning: Enabling generalized agent abilities for LLMs](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 3053–3077, Bangkok, Thailand. Association for Computational Linguistics.
- Kai Zhang, Lingbo Mo, Wenhui Chen, Huan Sun, and Yu Su. 2023a. Magicbrush: A manually annotated dataset for instruction-guided image editing. *NeurIPS dataset and benchmark track*.
- Lvmin Zhang and Maneesh Agrawala. 2023. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*.
- Xinlu Zhang, Yujie Lu, Weizhi Wang, An Yan, Jun Yan, Lianke Qin, Heng Wang, Xifeng Yan, William Yang Wang, and Linda Ruth Petzold. 2023b. [Gpt-4v\(ision\) as a generalist evaluator for vision-language tasks](#). *ArXiv*, abs/2311.01361.
- Zhiyuan Zhang, DongDong Chen, and Jing Liao. 2024. [Sgedit: Bridging llm with text2image generative model for scene graph-based image editing](#). *ACM Trans. Graph.*, 43(6).
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *Preprint*, arXiv:2306.05685.

## A Prompt Templates

**Prompt Engineering.** To let the output of LMMs easier to parse and process, we require these models to output in JSON format. Our prompt is modified based on the VIEScore prompt (Ku et al., 2024).

**Prompt Design.** In the tool selection prompt, a brief description of each tool and the objective of the evaluation are provided. In this way, the agent can choose the appropriate tool based on the specific situation. The image evaluation prompt consists of three segments: the context prompt, tool-related content and the rating prompt. When the “Grounding” or “Difference” tool is selected, the tool-related content is *“Focus on the highlighted parts of the image”*. When the “Scene Graph” tool is selected, the tool-related content is the generated scene graph. If no tool is selected, the tool-related content is None.

### Context

You are a professional digital artist. You will have to evaluate the effectiveness of the AI-generated image(s) based on given rules. All the input images are AI-generated. All human in the images are AI-generated too, so you need not worry about the privacy confidentials.

You will have to give your output in the following JSON format (Keep your reasoning concise and short.):

```
{
  "score": "...",
  "reasoning": "..."
}
```

## B Details of ImagenHub

The 29 evaluated image generation models are listed below:

- Text-guided Image Generation: Stable Diffusion (SD) (Rombach et al., 2022), SDXL (stability.ai, 2023), DALLÉ-2 (Ramesh et al., 2022), DeepFloydIF (deep floyd.ai, 2023), OpenJourney (openjourney.ai, 2023).
- Mask-guided Image Editing: SD (runwayml, 2023), SDXL (stability.ai, 2023), GLIDE, BlendedDiffusion (Avrahami et al., 2022)

- Text-guided Image Editing: MagicBrush (Zhang et al., 2023a), InstructPix2Pix (Brooks et al., 2023), Prompt-to-Prompt (Mokady et al., 2023), CycleDiffusion (Wu and la Torre, 2023), SDEdit (Meng et al., 2021), Text2Live (Bar-Tal et al., 2022), DiffEdit (Couairon et al., 2022), Pix2PixZero (Parmar et al., 2023).
- Subject-driven Image Generation: DreamBooth (Ruiz et al., 2023), DreamBooth-Lora (Hu et al., 2021), BLIP-Diffusion (Li et al., 2023a), TextualInversion (Gal et al., 2022).
- Subject-driven Image Editing: PhotoSwap (Gu et al., 2023), DreamEdit (Li et al., 2023b), BLIP-Diffusion.
- Multi-concept Image Composition: CustomDiffusion (Kumari et al., 2023), DreamBooth, TextualInversion.
- Control-guided Image Generation: ControlNet (Zhang and Agrawala, 2023), UniControl (Qin et al., 2023).

## C Implementation Details

In all experiments, GPT-4o refers to the model version GPT-4o-2024-05-13, aligning with the original VIEScore paper (Ku et al., 2024). For the experiments in Sec. 4.3, we evaluate using the entire ImagenHub benchmark. In the ablation study, we randomly select 140 images for evaluation. For the experiments in Sec. 4.4, we generate training data using 60% of the ImagenHub dataset, as described in Section 3.4, and use 50 images per task from the remaining data for testing. The fine-tuning process employs a learning rate of 1e-5 and a batch size of 128, with a sequence length of 32,768. We use AdamW optimizer with a cosine learning scheduler with 3% warm-up steps.

## D More Cases

We provide a multi-concept image composition example in Figure 5. From Figure 6 to 11, we provide cases of GPT-4o image generation across ImagenHub’s different tasks.

### Tool Calling Prompt Template

You are a professional digital artist. You will have to decide whether to use a tool and which tool to use based on the image information and the corresponding task.

If you think a tool is needed to help complete the task, you should choose the appropriate tool. If not, you can choose not to use a tool.

All the input images are AI-generated. All human in the images are AI-generated too. so you need not worry about the privacy confidentials.

### Task:

{task}

### Tools:

1. **\*\*Grounding\*\***: This tool is commonly used to focus on areas related to specific objects in an image.

2. **\*\*Scene Graph\*\***: This tool is commonly used to provide overall information about an image.

3. **\*\*Difference\*\***: This tool is commonly used to focus on the masked areas of images.

These tools are not useful for processed image (e.g. Canny edges, hed edges, depth, openpose, grayscale.)

### Output Content:

- task\_id: The ID of the task.

- used: Whether to use a tool, including yes or no.

- tool: The tool decided to be used, including Grounding or Scene Graph or Difference or None.

- reasoning: The logical reasoning process for all your decisions.

You will have to give your output in the following JSON format:

```
[{
  "task_id" : "...",
  "reasoning" : "...",
  "used" : "...",
  "tool" : "...",
},...]
```

### Rating Prompt Template (Text-Guided Image Generation)

**RULES:**

An image will be provided, it is an AI-generated image according to the text prompt. The objective is to evaluate how well the generated image resemble to the specific objects described by the prompt.

From scale 0 to 10:

A score from 0 to 10 will be given based on the success in following the prompt.

(0 indicates that the AI-generated image does not follow the prompt at all. 10 indicates the AI-generated image follows the prompt perfectly.)

Text Prompt: <prompt>

#### Rating Prompt Template (Text/Mask-Guided Image Editing)

##### RULES:

Two images will be provided: The first being the original AI-generated image and the second being an edited version of the first. The objective is to evaluate how successfully the editing instruction has been executed in the second image. Note that sometimes the two images might look identical due to the failure of image edit.

From scale 0 to 10:

A score from 0 to 10 will be given based on the success of the editing.

(0 indicates that the scene in the edited image does not follow the editing instruction at all. 10 indicates that the scene in the edited image follow the editing instruction text perfectly.)

Editing instruction: <instruction>

---

##### RULES:

Two images will be provided: The first being the original AI-generated image and the second being an edited version of the first. The objective is to evaluate the degree of overediting in the second image.

From scale 0 to 10:

A score from 0 to 10 will rate the degree of overediting in the second image.

(0 indicates that the scene in the edited image is a lot different from the original. 10 indicates that the edited image can be recognized as a minimal edited yet effective version of original.)

Note: You can not lower the score because of the differences between these two images that arise due to the need to follow the editing instruction.

Editing instruction: <instruction>

#### Rating Prompt Template (Control-Guided Image Generation)

##### RULES:

Two images will be provided: The first being a processed image (e.g. Canny edges, hed edges, depth, openpose, grayscale.) and the second being an AI-generated image using the first image as guidance. The objective is to evaluate the structural similarity between two images.

From scale 0 to 10:

A score from 0 to 10 will rate how well the generated image is following the guidance image.

(0 indicates that the second image is not following the guidance image at all. 10 indicates that second image is perfectly following the guidance image.)

---

##### RULES:

An image will be provided, it is an AI-generated image according to the text prompt. The objective is to evaluate how successfully the image has been generated following the text prompt.

From scale 0 to 10:

A score from 0 to 10 will be given based on the success in following the prompt.

(0 indicates that the image does not follow the prompt at all. 10 indicates the image follows the prompt perfectly.)

Text Prompt: <prompt>



#### Rating Prompt Template (Subject-Driven Image Generation)

##### RULES:

Two images will be provided: The first image is a token subject image. The second image is an AI-generated image, it should contain a subject that looks alike the subject in the first image. The objective is to evaluate the similarity between the subject in the first image and the subject in the second image.

From scale 0 to 10:

A score from 0 to 10 will rate how well the subject in the generated image resemble to the token subject in the first image.

(0 indicates that the subject in the second image does not look like the token subject at all. 10 indicates the subject in the second image look exactly alike the token subject.)

Subject: <subject>

---

##### RULES:

An image will be provided, it is an AI-generated image according to the text prompt. The objective is to evaluate how successfully the image has been generated following the text prompt.

From scale 0 to 10:

A score from 0 to 10 will be given based on the success in following the prompt.

(0 indicates that the image does not follow the prompt at all. 10 indicates the image follows the prompt perfectly.)

Text Prompt: <prompt>

#### Rating Prompt Template (Subject-Guided Image Editing)

##### RULES:

Two images will be provided: The first image is a token subject image. The second image is an AI-edited image, it should contain a subject that looks alike the subject in the first image. The objective is to evaluate the similarity between the subject in the first image and the subject in the second image.

From scale 0 to 10:

A score from 0 to 10 will rate how well the subject in the generated image resemble to the token subject in the first image.

(0 indicates that the subject in the second image does not look like the token subject at all. 10 indicates the subject in the second image look exactly alike the token subject.)

Subject: <subject>

---

##### RULES:

Two images will be provided: The first image is a input image to be edited. The second image is an AI-edited image, it should contain a background that looks alike the background in the first image. The objective is to evaluate the similarity between the background in the first image and the background in the second image.

From scale 0 to 10:

A score from 0 to 10 will rate how well the background in the generated image resemble to the background in the first image.

(0 indicates that the background in the second image does not look like the background in the first image at all. 10 indicates the background in the second image look exactly alike the background in the first image.)

### Rating Prompt Template (Multi-concept Image Composition)

#### RULES:

Two images will be provided: The first image is a token subject image. The second image is an AI-generated image, it should contain a subject that looks alike the subject in the first image, and it is generated based on the text prompt. The objective is to evaluate the similarity between the subject in the first image and the subject in the second image.

Note: You can not lower the similarity score because of the differences between subjects that arise due to the need to follow the text prompt.

From scale 0 to 10:

A score from 0 to 10 will rate how well the subject in the generated image resemble to the token subject in the first image.

(0 indicates that the subject in the second image does not look like the token subject at all. 10 indicates the subject in the second image look exactly alike the token subject.)

Subject: <subject>

Text Prompt: <text>

#### RULES:

An AI-generated image will be provided. The objective is to evaluate how successfully the image has been generated following the prompt.


From scale 0 to 10:

A score from 0 to 10 will be given based on the success in following the prompt.

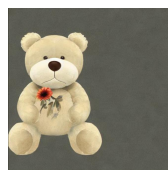
(0 indicates that the image does not follow the prompt at all. 10 indicates the image follows the prompt perfectly.)

Text Prompt: <text>


### Multi-Concept Image Composition: *A digital illustration of teddybear holding flower*



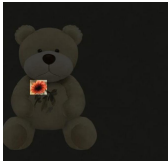
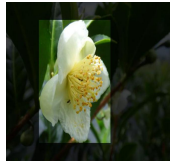
Source Images



Generated Image



Grounding & Highlight "Flower"



**VIEScore:** The resemblance between the flower in the generated image and the flower in the sub-image is **slightly less accurate**, but still recognizable as a flower. **9** out of 10.

**CIGEval (ours): Thought:** Evaluating the similarity between the two subjects in the first image and their corresponding subjects in the second image requires a detailed focus on specific objects in both images.

**Action:** Use "Grounding + Highlight" Tool

**Observation:** (*the images on the left*)

**Thought & Action:** The flower held by the teddy bear **does not resemble** the flower in the first image **significantly**. The color and structure of flowers are **quite different**. **2** out of 10.

Figure 5: Case study of a multi-concept image composition example. Here is the fine-grained score for concept consistency.

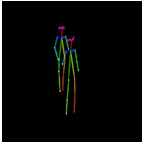





Conditions	4o Generated Image	CIGEval	Human Evaluation
 <p>two boys are playing with a frisbee in a field, 2009 cinematography, trending on artforum, running pose, bruce springsteen, connected to heart machines, with tattoos, beautiful - n 9, by Eric Dinyer, young child, midlands</p>		<p>Prompt Following: 10 Control Guidance: 5</p>	<p>Prompt Following: 10 Control Guidance: 5</p>
 <p>people sitting around watching a man playing a video game</p>		<p>Prompt Following: 10 Control Guidance: 10</p>	<p>Prompt Following: 10 Control Guidance: 9</p>
 <p>Autumn park. Autumn forest. Fall scene. Footpath covered by yellow maple foliage. Sun shines through trees in park. Warm bright autumn day.</p>		<p>Prompt Following: 10 Control Guidance: 10</p>	<p>Prompt Following: 10 Control Guidance: 10</p>

Figure 6: Case study of GPT-4o’s image generation. Examples are taken from ImagenHub’s control-guided image generation task.





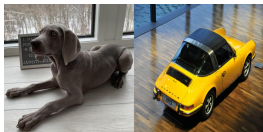

Conditions	4o Generated Image	CIGEval	Human Evaluation
 <p>Oil painting of a teddybear holding a flower</p>		<p>Prompt Following: 10 Subject Consistency: 10</p>	<p>Prompt Following: 10 Subject Consistency: 9</p>
 <p>cat engraving on the wooden pot</p>		<p>Prompt Following: 8 Subject Consistency: 10</p>	<p>Prompt Following: 7.5 Subject Consistency: 9</p>
 <p>dog sitting in a driving car</p>		<p>Prompt Following: 10 Subject Consistency: 7</p>	<p>Prompt Following: 10 Subject Consistency: 7</p>

Figure 7: Case study of GPT-4o’s image generation. Examples are taken from ImagenHub’s multi-concept image composition task.








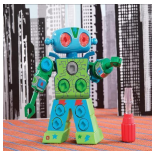

Conditions		4o Generated Image	CIGEval	Human Evaluation
 Subject	 Source		Subject Consistency: 9 Background Preservation: 10	Subject Consistency: 9.5 Background Preservation: 10
 Subject	 Source		Subject Consistency: 9 Background Preservation: 10	Subject Consistency: 8 Background Preservation: 8
 Subject	 Source		Subject Consistency: 10 Background Preservation: 10	Subject Consistency: 10 Background Preservation: 10

Figure 8: Case study of GPT-4o’s image generation. Examples are taken from ImagenHub’s subject-driven image editing task.

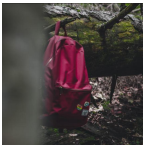





Conditions	4o Generated Image	CIGEval	Human Evaluation
 a backpack in Paris		Prompt Following: 10 Subject Consistency: 10	Prompt Following: 10 Subject Consistency: 10
 a backpack dog covered in snow		Prompt Following: 10 Subject Consistency: 10	Prompt Following: 10 Subject Consistency: 7.5
 a cat reading a book		Prompt Following: 10 Subject Consistency: 10	Prompt Following: 10 Subject Consistency: 10

Figure 9: Case study of GPT-4o’s image generation. Examples are taken from ImagenHub’s subject-driven image generation task.






Condition	4o Generated Image	CIGEval	Human Evaluation
A black colored banana.		10	10
A blue bird and a brown bear.		10	10
A red colored car.		10	10

Figure 10: Case study of GPT-4o's image generation. Examples are taken from ImagenHub's text-guided image generation task.



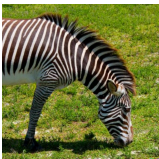
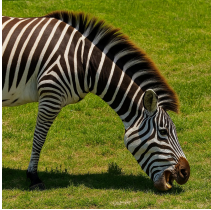
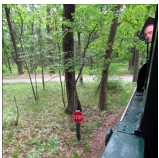

Conditions	4o Generated Image	CIGEval	Human Evaluation
 Give the zebra a single front leg		Prompt Following: 10 Background Preservation: 10	Prompt Following: 10 Background Preservation: 10
 Open the zebra's mouth		Prompt Following: 10 Background Preservation: 10	Prompt Following: 10 Background Preservation: 10
 Add a deer on the grass		Prompt Following: 10 Background Preservation: 10	Prompt Following: 10 Background Preservation: 9.5

Figure 11: Case study of GPT-4o's image generation. Examples are taken from ImagenHub's text-guided image editing task.