

# IMPROVING COMPOSITIONAL TEXT-TO-IMAGE GENERATION WITH LARGE VISION-LANGUAGE MODELS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Recent advancements in text-to-image models, particularly diffusion models, have shown significant promise. However, compositional text-to-image models frequently encounter difficulties in generating high-quality images that accurately align with input texts describing multiple objects, variable attributes, and intricate spatial relationships. To address this limitation, we employ large vision-language models (LVLMs) for multi-dimensional assessment of the alignment between generated images and their corresponding input texts. Utilizing this assessment, we fine-tune the diffusion model to enhance its alignment capabilities. During the inference phase, an initial image is produced using the fine-tuned diffusion model. The LVLM is then employed to pinpoint areas of misalignment in the initial image, which are subsequently corrected using the image editing algorithm until no further misalignments are detected by the LVLM. The resultant image is consequently more closely aligned with the input text. Our experimental results validate that the proposed methodology significantly improves text-image alignment in compositional image generation, particularly with respect to object number, attribute binding, spatial relationships, and aesthetic quality.

## 1 INTRODUCTION

Recently, text-to-image models have made great progress, of which diffusion models are remarkable. Compositional text-to-image generation is a more advanced task, which understands and generates images with multiple objects which have variable attributes and complex spatial relationships. Although diffusion models can generate high-quality images, compositional text-to-image generation still struggles to generate images aligned with input texts. These limitations manifest inaccuracies in object number, attribute binding, spatial relationships between objects and aesthetic quality, as shown in Figure 1. These inaccuracies are caused by the compositional complexity of the input text and the cross-attention mechanism in the diffusion model.

Several studies (Agarwal et al., 2023; Chefer et al., 2023) have attempted to resolve issues related to attribute binding by manipulating the attention mechanism within diffusion models. However, these approaches often fall short of comprehensively addressing jointly the interrelated challenges of object number, attribute binding, spatial relationships between objects, and aesthetic quality. Besides, the evaluation of the alignment of the generated image and input text is not fully explored. Several techniques, such as CLIPScore (Hessel et al., 2021) or BLIP (Li et al., 2022), fall short in capturing compositional alignment accurately. While there are methodologies like T2I-CompBench (Huang et al., 2023) that introduce compositional evaluation methods, these are primarily applied to fine-tune diffusion models during the training phase. This approach fails to exploit the full potential of compositional evaluation methods, leaving room for enhancing the overall alignment between generated images and input texts in the inference period.

Parallel to these developments, Large Language Models (LLMs), such as GPT-4, LLama, and Vicuna, have emerged as influential tools in both academia and industry. Building upon LLMs, Large Vision-Language Models (LVLMs) have been developed to integrate visual features with language representations, endowing LLMs with multimodal capabilities. Notable models such as MiniGPT-4 (Zhu et al., 2023), LLama-Adapter (Zhang et al., 2023a), and Bard (Google, 2023) have demonstrated remarkable zero-shot learning, visual perception and reasoning abilities.

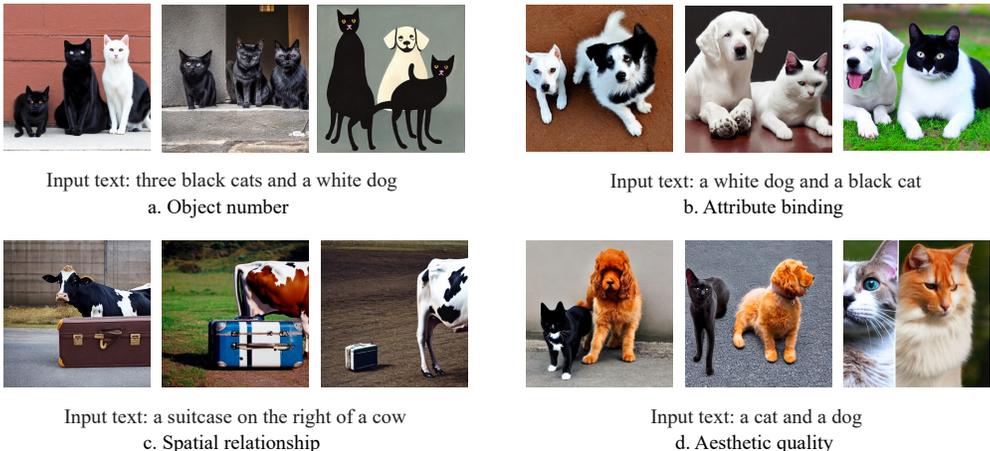


Figure 1: **Illustrating Limitations in Compositional Text-to-Image Generation.** (a) Object Number: The discrepancy between the quantity of objects in the image (e.g., cat and dog) and the input text is evident. (b) Attribute Binding: The attributes of objects depicted do not correspond with the input text; for instance, the cat’s color is black and white, contrasting with the specified black. (c) Spatial Relationship: The arrangement of objects does not conform to the input text, with the suitcase not situated to the right of the cow as described. (d) Aesthetic Quality: The representation of the object is distorted, deviating from conventional aesthetic standards.

Leveraging these advancements, our work incorporates the LVLM to augment the capabilities of compositional text-to-image generation. Specifically, we first utilize LVLMs for evaluation. To capture the compositional feature in alignment, we evaluate the alignment of generated images with input texts mainly in terms of four dimensions: object number, attribute binding, spatial relationship, and aesthetic quality. We employ the LVLM to formulate questions derived from the input text and subsequently input the generated image and questions into the LVLM. The responses obtained serve as LVLM-assessed metrics to evaluate alignment. Following this, we employ Reward Feedback Learning (ReFL) during the training phase to fine-tune diffusion models based on these metrics, thereby enhancing compositional text-to-image generation. To fully utilize the evaluation method, during the inference stage, the LVLM is re-engaged to identify and correct errors in the generated images. An iterative process governed by the LVLM uses image-editing algorithms to eliminate misalignments until the generated image is fully compliant with the input text. Our empirical results substantiate that our approach significantly amplifies the accuracy and fidelity of compositional image generation.

In conclusion, the key contributions of our study include:

- We leverage LVLMs to assess the alignment between generated images and input texts, focusing on object number, attribute binding, spatial relationships, and aesthetic quality within compositional text-to-image models.
- We enhance image-text alignment by fine-tuning the diffusion models through LVLM-based evaluations during the training period.
- We design an LVLM-guided iterative correction process to systematically rectify any misalignments in the generated images during the inference period.

Through these contributions, our work establishes a robust and plug-and-play framework for improving compositional text-to-image diffusion models.

## 2 RELATED WORK

**Text-to-Image Generation.** The aim of text-to-image generation is to produce images based on input textual descriptions. Significant advances in generative models, such as generative adversarial networks (GANs (Goodfellow et al., 2014)), auto-regressive models (Vaswani et al., 2017), and diffusion models (Ho et al., 2020), have paved the way for a plethora of works in this domain. GANs were initially employed for this purpose Reed et al. (2016), and numerous subsequent GAN-based models sought to enhance visual fidelity and caption congruence (Zhang et al., 2017; 2018; Xu et al., 2018; Li et al., 2019; Dong et al., 2017; Zhu et al., 2019; Tao et al., 2020; Ye et al., 2021; Kang et al., 2023; Sauer et al., 2023). However, GANs are not without challenges, particularly in terms of mode-collapse and training instability.

To address these issues, researchers have investigated the use of Transformer-based auto-regressive models for text-to-image generation (Ramesh et al., 2021; Ding et al., 2021; Esser et al., 2021a; Ding et al., 2022; Zhang et al., 2021; Lee et al., 2022; Chang et al., 2023), combined with a discrete VAE (Van Den Oord et al., 2017; Razavi et al., 2019; Esser et al., 2021b) for image tokenization and Transformers (Vaswani et al., 2017) for modeling the joint distribution of textual and image tokens. This often follows a two-stage methodology: initially, a discrete VAE tokenizes the input image, and subsequently, a multi-layer Transformer integrates text and image tokens.

Diffusion models have also been embraced for text-to-image generation (Nichol et al., 2021; Ho et al., 2022; Ramesh et al., 2022; Saharia et al., 2022; Rombach et al., 2022; Xu et al., 2022; Zhang et al., 2023b). For instance, GLIDE (Nichol et al., 2021) innovatively conditions the diffusion model on an input caption, building upon earlier works in the diffusion model sphere (Dhariwal & Nichol, 2021; Ho & Salimans, 2022). Moreover, DALL-E 2 (Ramesh et al., 2022) enhances the GLIDE model by conditioning on a supplemental CLIP image embedding for heightened diversity. Some endeavors, like Stable Diffusion (Rombach et al., 2022), emphasize computational efficiency by first representing input images as low-dimension latent codes. Nevertheless, challenges such as alignment with human preferences and textual input continue to persist.

**Alignment of Text-to-Image Generation Models.** Efforts have been made to align text-to-image generation models with human preferences and aesthetic standards (Hao et al., 2022; Lee et al., 2023; Wu et al., 2023; Xu et al., 2023; Dong et al., 2023a; Fang et al., 2023). For instance, Lee et al. (2023) concentrates on text alignment, using a reward model trained on human-annotated datasets to refine the text-to-image model. Similarly, ImageReward (Xu et al., 2023) offers a universal human preference reward model that encompasses text-image alignment, body problems, aesthetics, toxicity, and biases.

Promptist Hao et al. (2022) introduces prompt adaptation by training a language model to enhance the original prompt. This method leverages both the CLIP model and an aesthetic predictor as reward models, and fine-tunes in a supervised manner within a reinforcement learning framework. Alternatively, given the inefficiencies and instabilities associated with Reinforcement Learning from Human Feedback (RLHF (Ouyang et al., 2022)), Dong et al. (2023a) proposes reward ranked finetuning to better align generative models.

Despite the strides made in aligning text-to-image generation models with human preferences and aesthetic standards, complexities still arise when dealing with intricate prompts delineating multiple objects, diverse attributes, and elaborate spatial relations. Addressing these challenges, our work innovatively integrates large vision-language models (LVLMs) with diffusion models, presenting a refined approach for enhancing text-image alignment, especially in the realm of compositional image generation.

## 3 METHOD

### 3.1 PRELIMINARY

**Latent Diffusion Models.** Recently, diffusion models, exemplified by DALL-E and Midjourney, have gained widespread adoption in the field of text-to-image generation. These generative models aim to create desired data by denoising from a Gaussian distribution  $x_T \sim N(0, 1)$ . Initially, the diffusion models define a forward process by constructing a Markov chain of variables  $x_1, x_2, \dots, x_T$

from the target distribution  $x_0 \sim q(x_0)$  by iteratively adding Gaussian noise based on a predefined schedule  $\beta_t$ :

$$q(x_t|x_{t-1}) = N(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I). \quad (1)$$

Subsequently, the target distribution is transformed to a Gaussian distribution at step  $T$ . The diffusion models are tasked with learning the reverse process to approximate the true posterior  $q(x_{t-1}|x_t)$  by denoising from Gaussian distribution  $x_T \sim \mathcal{N}(0, 1)$ :

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \sigma_t), \quad (2)$$

where  $\mu_\theta$  represents the mean, computed using neural networks. This reverse process generates the desired sample  $x_0$  at last. In contrast, latent diffusion models execute the aforementioned forward and reverse processes in the latent space rather than the pixel space. This adaptation aims to mitigate the computational cost and enhance semantic generation. These models employ the Variational Autoencoder (VAE) to encode to the latent space  $z_0 = f(x_0)$ . They can be used in text-to-image generation by inputting a prompt  $T$  to generate the corresponding image:

$$p_\theta(z_{t-1}|z_t) = \mathcal{N}(z_{t-1}; \mu_\theta(z_t, t, T), \sigma_t). \quad (3)$$

The training loss is derived from the variational lower bound (VLB) loss, which can be simplified as

$$\mathcal{L}(\theta) = E_{t \sim [1, T], z_0, \epsilon_t} [\|\epsilon_t - \epsilon(z_t, t, T)\|_2^2]. \quad (4)$$

Stable Diffusion is based on latent diffusion models, which allows latent diffusion models to be versatile and efficient in generating high-quality, semantically coherent images from textual prompts. This method serves as the baseline for our study.

**ImageReward.** Latent Diffusion Models (LDMs) have shown significant potential as generative models. One of the highlighted challenges in the realm of LDMs is their direct optimization. The ReFL (Xu et al., 2023) methodology, as discussed in prior works, offers a potential solution to this optimization challenge.

LDMs follow a sequential denoising process, which in experiments, expands up to 40 steps. A significant observation from the ReFL approach highlights the behavior of model scores during these denoising steps:

- **Early Stages (Steps 1 to 15):** In this phase, the model scores remain uniformly low for all generated outputs.
- **Intermediate Stages (Steps 15 to 30):** Here, while high-quality generations become evident, it’s still early to conclusively gauge the final quality of all generations based on the present model scores.
- **Late Stages (Steps 30 onwards):** At this juncture, there’s a discernible distinction in generations based on their respective model scores.

These observations suggest that model scores after the 30th denoising step could be potentially reliable indicators for enhancing LDMs, even if they aren’t derived from the final step.

The ReFL algorithm, as elucidated in the referenced literature, aims to harness these model scores as feedback to back-propagate and refine the LDMs. This stands in contrast to traditional methodologies where only the gradient from the final denoising step is retained – an approach found to yield instability.

For effective fine-tuning, there is a balance established between the ReFL loss and the pre-training loss, a strategy to counter rapid overfitting and ensure a more stable fine-tuning process. The corresponding loss functions are illustrated as:

$$\mathcal{L}_{reward} = \lambda \mathbb{E}_{y_i \sim \mathcal{Y}} (\phi(r(y_i, g_\theta(y_i)))) \quad (5)$$

$$\mathcal{L}_{pre} = \mathbb{E}_{(y_i, x_i) \sim \mathcal{D}} (\mathbb{E}_{\mathcal{E}(x_i, y_i, \epsilon \sim \mathcal{N}(0, 1), t)} [\|\epsilon - \epsilon_\theta(z_t, t, \tau_\theta(y_i))\|_2^2]) \quad (6)$$

Here,  $\theta$  denotes the parameters of the LDM, and  $g_\theta(y_i)$  represents the generated image of the LDM using parameters  $\theta$  corresponding to prompt  $y_i$ . Our approach synergistically combines large vision-language models with diffusion models, wherein the LVLM evaluates the generation outcomes, subsequently guiding the ReFL training of the diffusion model. This integration of LVLMs and diffusion models via ReFL underpins a pioneering methodology for enhancing text-image coherence, especially for intricate prompts. This strategy is pivotal to our research and establishes the advanced framework upon which our study is built.

### 3.2 OVERVIEW

In this study, we introduce a comprehensive framework that leverages LVLMs to enhance compositional text-to-image generation. The Figure 2 provided illustrates a schematic representation of our proposed framework, which integrates three core components: LVLM-based Evaluation, Model Fine-tuning, and LVLM-guided Image Editing. Each component strategically utilizes the capabilities of LVLMs to optimize the generation process.

Initially, we deploy the LVLM to assess the alignment between generated images and input texts. The LLMs analyze the input texts and formulate questions aimed at capturing compositional features inherent in the texts. The generated images, along with these questions, are then fed into the LVLM, which produces answers serving as evaluative metrics for alignment assessment. Subsequent to the LVLM-based evaluation, we employ Reward Feedback Learning (ReFL) to fine-tune the diffusion models. This process aims to optimize the models based on the evaluative metrics derived from LVLM evaluation during the training period, thereby enhancing alignment between the images generated and the input texts. To fully harness the potential of the LVLM, we incorporate it during the inference stage as well. Specifically, the LVLM is used to identify any misalignment between the generated images and the input texts. Upon detection of misalignments, the LVLM is used to guide the correction process with image-editing algorithms until the generated images are fully aligned with the input texts, ensuring no misalignment remains.

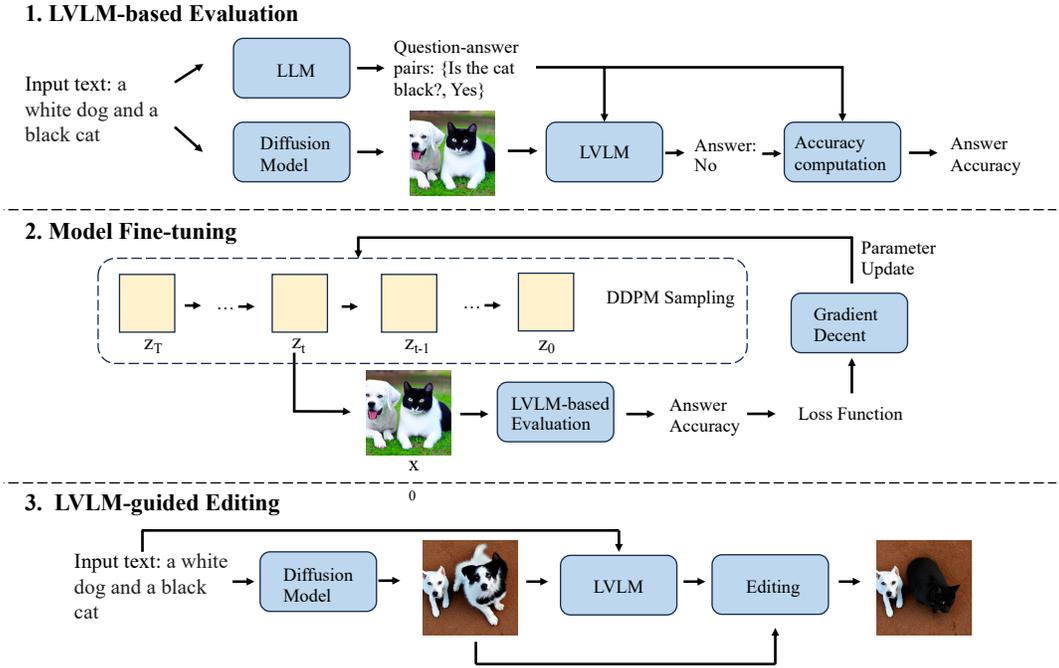


Figure 2: **Overview of the Proposed Methodology.** Our methodology is structured around three core components: (1) LVLM-based Evaluation: Drawing inspiration from TIFA, we initially employ LLM to formulate question-answer pairs grounded in the input text. Subsequently, the LVLM is utilized to procure answers by processing the formulated questions alongside the image. A comparative analysis of answers derived from both image and text is then undertaken to calculate the answer accuracy, serving as our evaluative metric. (2) Model Fine-tuning: The LVLM-based evaluation metric is incorporated as a weight within the diffusion loss function, facilitating the fine-tuning of the diffusion model. The objective is to guide the diffusion model’s focus towards enhancing answer accuracy. (3) LVLM-guided Editing: In the inference phase, the LVLM is deployed to identify misalignments between image and text. Subsequent to this identification, image-editing algorithms are applied iteratively to rectify the image until no alignment is detected.

### 3.3 LVLM-BASED EVALUATION

To evaluate the alignment of generated images in compositional text-to-image synthesis, we integrate Large Vision-Language Models (LVLMs) into the Text-to-Image Faithfulness evaluation with Question Answering (TIFA (Hu et al., 2023)) framework. This framework assesses various elements such as objects, shapes, materials, attributes, and spatial relationships in the images by formulating specific prompts. The evaluation process unfolds as follows:

**Text Analysis.** The initial step involves analyzing the input text  $T$ . In adherence to TIFA guidelines, we employ LLMs to generate question-answer pairs  $\{Q_i, A_i\}_{i=1}^N$  derived from the input text  $T$ , thereby utilizing the zero-shot, in-context, and reasoning capabilities of LLMs. These generated pairs encapsulate the compositional information present in the text, encompassing aspects like object number, attribute binding and spatial relationship.

**LVLM-based Question Answering:** Based on the TIFA framework, the LVLM is used to answer the formulated questions. Upon generating question-answer pairs  $Q_i, A_i$  from the given text  $T$ , both the questions  $Q_i$  and the generated image  $I$  are input into the LVLM. The LVLM, through reasoning on the generated image  $I$ , produces the answers  $\tilde{A}_i$ .

**Accuracy Computation:** For each pair of text and image  $T, I$ , we compare the answers derived from text  $Q_i$  and the image  $\tilde{Q}_i$ . The accuracy is computed as:

$$\text{ACC}(T, I) = \sum_{i=1}^n \mathbb{1}[Q_i = \tilde{Q}_i]. \quad (7)$$

By doing so, the LVLM facilitates a nuanced and detailed evaluation in terms of the answer accuracy of how well the generated images align with the input text, examining the fidelity of the representation across object number, attribute binding and spatial relationship.

### 3.4 MODEL FINE-TUNING

To refine the alignment between text and image within diffusion models, we adopt a strategy known as Reward Feedback Learning (ReFL (Xu et al., 2023)). ReFL is employed to fine-tune diffusion models based on the LVLM-based evaluation during the training phase. The rewards in ReFL are used to backpropagate and update the diffusion parameters after a predetermined range of steps, because the latter steps yield clearer images, conducive for use in the LVLM, thereby enhancing the stability of the training process. During our model fine-tuning, we first sample plenty of text-image pairs from the diffusion model, subsequently employing the previously mentioned answer accuracy to assess each pair. However, the LVLM-based evaluation is non-differential. Different from ReFL, the answer accuracy serves as the weight of the loss function for fine-tuning the diffusion model. Higher answer accuracy indicates improved alignment between image and text, thus the optimization process prioritizes these instances. The loss function is formulated as:

$$\mathcal{L}'(\theta) = E_{(T, I)}[\text{ACC}(T, I) \cdot \|\epsilon - \epsilon(z_t, t, T)\|_2^2], \quad (8)$$

where  $(T, I)$  represents a sample of text-image pairs generated from the diffusion model, and  $\text{ACC}(T, I)$  is the answer accuracy derived from LVLM-based evaluation. The algorithm is depicted in Algorithm 1.

After the fine-tuning on sampled image-text pairs, the diffusion model is optimized with a focus on enhancing the answer accuracy in subsequently generated text-image pairs.

### 3.5 LVLM-GUIDED EDITING

To maximize the utility of LVLM-based evaluation, we incorporate it not only during the training period but also in the inference period. We note that despite the significant improvement in the alignment between text and image through fine-tuning, discrepancies can still arise during inference. To address this, we iteratively correct any misalignment in the initially generated image during the inference phase. For example, if the LVLM identifies a color discrepancy between an object in the generated image and the corresponding input text, an editing algorithm is activated to adjust the object’s color to align with the text description. In this process, we first generate an initial

**Algorithm 1** Training Procedure

---

**Input:** Fine-tuning text  $\{T_j\}_{j=1}^n$ , text-image pairs sampled from datasets  $\{\tilde{T}_j, \tilde{I}_j\}_{j=1}^n$ , number of diffusion step  $T$ , step range  $[t_1, t_2]$ , diffusion model  $D_\theta$

- 1: **for**  $j = 1, \dots, n$  **do**
- 2:   Compute  $\mathcal{L}(\theta, \tilde{T}_j, \tilde{I}_j)$  based on Eq. 4
- 3:    $\theta \leftarrow \theta + \alpha_1 \frac{\mathcal{L}(\theta)}{\theta}$
- 4:    $t \leftarrow \text{Random}(t_1, t_2)$
- 5:    $z_t \sim \mathcal{N}(0, 1)$
- 6:   **for**  $i = T, \dots, t + 1$  **do**
- 7:      $z_{i-1} \leftarrow D_\theta(z_i)$
- 8:   **end for**
- 9:    $z_{i-1} \leftarrow D_\theta(z_i)$
- 10:    $z_0 \leftarrow z_0(z_{i-1})$
- 11:    $x_0 \leftarrow \text{VAE Decoder}(z_0)$
- 12:   Compute  $\text{ACC}(x_0, I_j)$  based on Eq. 7 and  $\mathcal{L}'(\theta)$  based on Eq. 8
- 13:    $\theta \leftarrow \theta + \alpha_2 \frac{\mathcal{L}'(\theta)}{\theta}$
- 14: **end for**
- 15: **return** network parameters  $\theta$

---

image from the diffusion model, which has been fine-tuned using ReFL. Subsequently, the LVLM is employed to pinpoint instances of misalignment between text and image, such as disparities in object number, attribute binding, spatial relationships and aesthetic quality. Upon identification of misalignments, the LVLM is utilized to guide the process of rectifying the discrepancies by image-editing algorithms until alignment is achieved. Throughout this process, we leverage the editing capabilities of diffusion models. Initially, we employ the Segment Anything Model (SAM) to isolate all objects and backgrounds in the initial image. Following this, a diffusion-based inpainting model is used to modify the relevant objects or the background to achieve congruence with the input text. We delineate the correction of four types of misalignment as follows:

**Object Number.** When the LVLM discerns a discrepancy in the object number in the initial image compared to the input text, it categorizes the variance as either excess or deficit. If the object count exceeds the text description, SAM identifies specific objects, and the inpainting model eliminates the surplus entities. Conversely, if there are fewer objects, SAM targets the background, and the inpainting model introduces additional objects in the background.

**Attribute Binding.** The LVLM identifies instances where an object’s attributes do not align with the descriptions provided in the input text. For instance, if the input text describes “a white dog” while the image depicts a black dog, the LVLM recognizes the color inconsistency. Subsequently, the SAM and LVLM are used to generate a mask for the incorrectly colored dog, and a painting algorithm is employed to replace the black dog with a white one, thus ensuring attribute coherence.

**Spatial Relationship.** When discrepancies in the spatial relationships among multiple objects are identified, the SAM and LVLM are utilized to select the mask of the incorrectly positioned object. Subsequently, we employ the inpainting algorithm to first remove the mislocated object, followed by adding it back in a location that is in alignment with the input text.

**Aesthetic Quality.** If the LVLM identifies that an object within the image is distorted and thus falls short of human performance standards, our initial step is to employ SAM for segmenting the distorted object. Subsequently, an inpainting algorithm is utilized to substitute the distorted object with a normalized version. This approach leverages the capability of the diffusion model, which finds generating singular objects to be a more manageable task.

## 4 EXPERIMENT

In this section, we conduct experiments to validate our proposed method. Initially, we outline the implementation specifics in Section 4.1. Subsequently, we visualize the results from the LVLM-based

Evaluation and LVLM-based Editing in Section 4.2. Additionally, we also conduct comparative assessment between the baseline and the fine-tuned model in Section 4.2.

#### 4.1 IMPLEMENTATION DETAIL

**LVLM-based Evaluation.** In alignment with the TIFA approach, we employ LLama2 to generate question-answer pairs. We then use the Bard, a state-of-the-art LVLM proposed by Google, to process questions alongside the generated image, thereby producing the answers.

**Model Fine-tuning.** Adhering to the ReFL methodology, we source text-image pairs from LAION-5B and extract input text from DiffusionDB for training data. The model is fine-tuned in half-precision with a learning rate set to  $10^{-5}$ , while maintaining a batch size of 64 for each. The sample step range  $[T_1, T_2]$  is defined as  $[1, 10]$ . A key distinction in our approach is that we do not directly optimize the answer accuracy, as it is non-differential. Instead, we focus on optimizing a weighted diffusion loss function, where the answer accuracy serves as the weight.

**LVLM-based Editing.** Bard is once again employed as the guiding LVLM for the image editing process. To rectify any misalignment identified by the Bard, we amalgamate the capabilities of SAM and the Blended Diffusion model, utilizing them as the image-editing algorithm.

#### 4.2 EXPERIMENTAL RESULTS

**LVLM-based Evaluation.** We utilize Bard to generate answers by inputting both questions and the corresponding image. Illustrative examples of this process are presented in Figure 3, showcasing instances where images are generated based on the input text “a black dog is standing on a beach.”

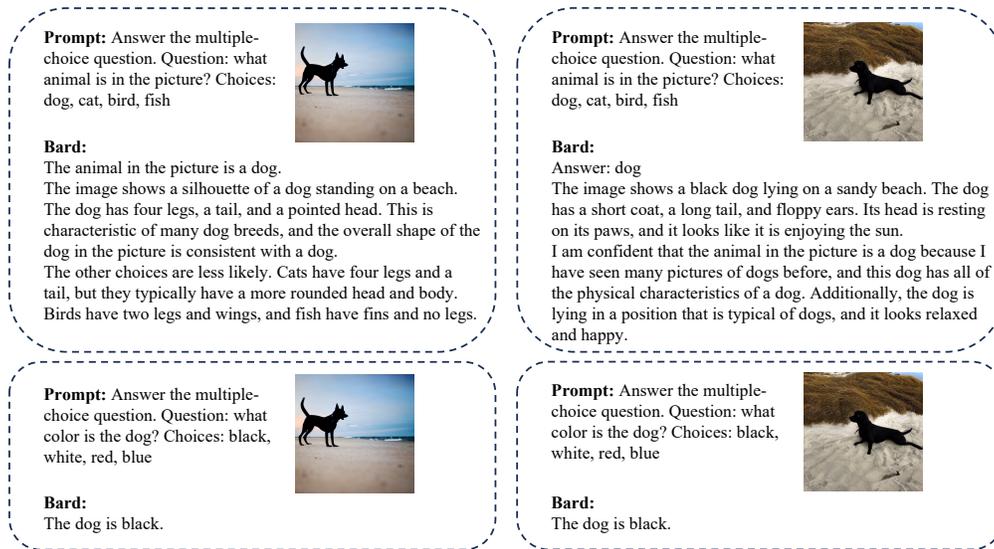


Figure 3: The answers produced by Bard. The images are generated with the text “a black dog is standing on a beach”.

**Model Fine-tuning.** A comparative analysis is conducted between the results derived from Stable Diffusion and those from the fine-tuned model on the dataset from T2ICompBench. The visual representation of these results can be observed in Figure 4. The CLIPScore attributed to the fine-tuned model is 0.3032, which is larger than the 0.3010 score associated with Stable Diffusion.

**LVLM-based Editing.** Employing Bard, we identify misalignments within the generated images and subsequently utilize SAM and Blended Diffusion to rectify these discrepancies. Illustrative examples of this process are presented in Figure 5.

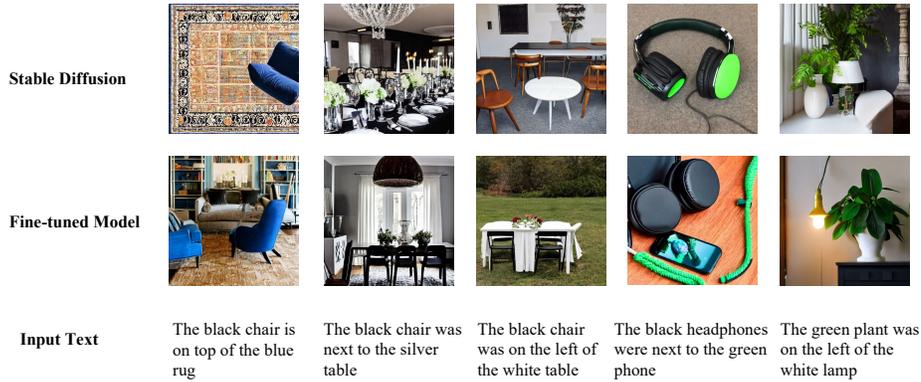


Figure 4: The images generated by Stable Diffusion and the fine-tuned model.

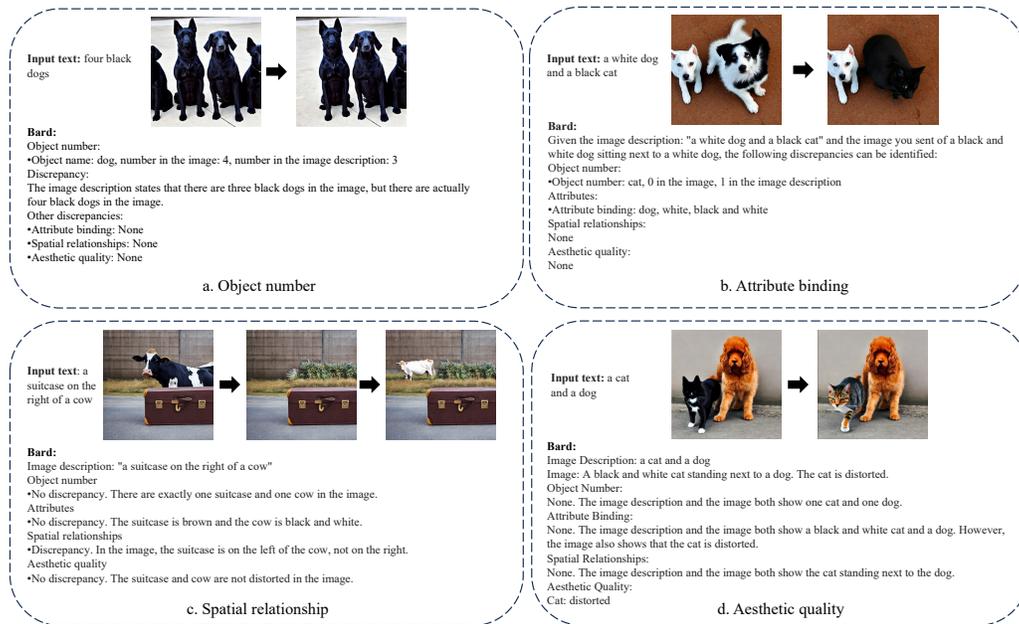


Figure 5: Visualization of LVLm-guided Editing.

## 5 CONCLUSION

**Limitations.** The efficacy of our method is constrained by the performance of the LVLms. The current version of Bard is still not very accurate. Nevertheless, given the rapid advancements in the field of LVLms, we will refine our approach in alignment with the development of more sophisticated LVLms.

**Conclusion.** In this paper, our primary objective is to enhance the quality of composable image generation using LVLms. Our methodology consists of three key components. Initially, we leverage LVLms to assess the alignment between the generated image and the input text. Following this, we fine-tune the diffusion models utilizing LVLm-based evaluations. In the subsequent inference phase, we deploy LVLms to detect any discrepancies between the text and the image, and an image-editing algorithm is engaged to amend these misalignments. Our empirical investigations substantiate the effectiveness of our approach in improving compositional image generation.

## REFERENCES

- Aishwarya Agarwal, Srikrishna Karanam, KJ Joseph, Apoorv Saxena, Koustava Goswami, and Balaji Vasan Srinivasan. A-star: Test-time attention segregation and retention for text-to-image synthesis. *arXiv preprint arXiv:2306.14544*, 2023.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022.
- Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Text-to-image generation via masked generative transformers. *arXiv preprint arXiv:2301.00704*, 2023.
- Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–10, 2023.
- Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023.
- Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis, 2021.
- 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. *Advances in Neural Information Processing Systems*, 34:19822–19835, 2021.
- Ming Ding, Wendi Zheng, Wenyi Hong, and Jie Tang. Cogview2: Faster and better text-to-image generation via hierarchical transformers. *arXiv preprint arXiv:2204.14217*, 2022.
- Hanze Dong, Wei Xiong, Deepanshu Goyal, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. Raft: Reward ranked finetuning for generative foundation model alignment. *arXiv preprint arXiv:2304.06767*, 2023a.
- Hao Dong, Simiao Yu, Chao Wu, and Yike Guo. Semantic image synthesis via adversarial learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5706–5714, 2017.
- Xiaoyi Dong, Jianmin Bao, Yinglin Zheng, Ting Zhang, Dongdong Chen, Hao Yang, Ming Zeng, Weiming Zhang, Lu Yuan, Dong Chen, et al. Maskclip: Masked self-distillation advances contrastive language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10995–11005, 2023b.
- Patrick Esser, Robin Rombach, Andreas Blattmann, and Bjorn Ommer. Imagebart: Bidirectional context with multinomial diffusion for autoregressive image synthesis. *Advances in Neural Information Processing Systems*, 34:3518–3532, 2021a.
- Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12873–12883, 2021b.
- Guian Fang, Zutao Jiang, Jianhua Han, Guansong Lu, Hang Xu, and Xiaodan Liang. Boosting text-to-image diffusion models with fine-grained semantic rewards, 2023.
- Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xiangyu Yue, et al. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*, 2023.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Google. Bard. <https://bard.google.com>, 2023.

- Ziyu Guo, Renrui Zhang, Xiangyang Zhu, Yiwen Tang, Xianzheng Ma, Jiaming Han, Kexin Chen, Peng Gao, Xianzhi Li, Hongsheng Li, et al. Point-bind & point-llm: Aligning point cloud with multi-modality for 3d understanding, generation, and instruction following. *arXiv preprint arXiv:2309.00615*, 2023.
- Jiaming Han, Renrui Zhang, Wenqi Shao, Peng Gao, Peng Xu, Han Xiao, Kaipeng Zhang, Chris Liu, Song Wen, Ziyu Guo, et al. Imagebind-llm: Multi-modality instruction tuning. *arXiv preprint arXiv:2309.03905*, 2023.
- Yaru Hao, Zewen Chi, Li Dong, and Furu Wei. Optimizing prompts for text-to-image generation. *arXiv preprint arXiv:2212.09611*, 2022.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 7514–7528, 2021.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance, 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.*, 23: 47–1, 2022.
- Yushi Hu, Benlin Liu, Jungo Kasai, Yizhong Wang, Mari Ostendorf, Ranjay Krishna, and Noah A Smith. Tifa: Accurate and interpretable text-to-image faithfulness evaluation with question answering. *arXiv preprint arXiv:2303.11897*, 2023.
- Kaiyi Huang, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation. *arXiv preprint arXiv:2307.06350*, 2023.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pp. 4904–4916. PMLR, 2021.
- Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. Scaling up gans for text-to-image synthesis. *arXiv preprint arXiv:2303.05511*, 2023.
- Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11523–11532, 2022.
- Kimin Lee, Hao Liu, Moonkyung Ryu, Olivia Watkins, Yuqing Du, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, and Shixiang Shane Gu. Aligning text-to-image models using human feedback. *arXiv preprint arXiv:2302.12192*, 2023.
- Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip HS Torr. Controllable text-to-image generation. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp. 2065–2075, 2019.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pp. 12888–12900. PMLR, 2022.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
- Yanguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv preprint arXiv:2110.05208*, 2021.

- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023.
- Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.
- OpenAI. Chatgpt. <https://chat.openai.com>, 2023a.
- OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023b.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744, 2022.
- 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. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pp. 8821–8831. PMLR, 2021.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in neural information processing systems*, pp. 14866–14876, 2019.
- Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *International Conference on Machine Learning*, pp. 1060–1069. PMLR, 2016.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, June 2022.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022.
- Axel Sauer, Tero Karras, Samuli Laine, Andreas Geiger, and Timo Aila. Stylegan-t: Unlocking the power of gans for fast large-scale text-to-image synthesis. *arXiv preprint arXiv:2301.09515*, 2023.
- Ming Tao, Hao Tang, Songsong Wu, Nicu Sebe, Xiao-Yuan Jing, Fei Wu, and Bingkun Bao. Df-gan: Deep fusion generative adversarial networks for text-to-image synthesis. *arXiv preprint arXiv:2008.05865*, 2020.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, Ł. Kaiser, and I. Polosukhin. Attention Is All You Need. In *Neural Information Processing Systems*, 2017.
- Xiaoshi Wu, Keqiang Sun, Feng Zhu, Rui Zhao, and Hongsheng Li. Better aligning text-to-image models with human preference. *arXiv preprint arXiv:2303.14420*, 2023.

- Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation, 2023.
- Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1316–1324, 2018.
- Xingqian Xu, Zhangyang Wang, Eric Zhang, Kai Wang, and Humphrey Shi. Versatile diffusion: Text, images and variations all in one diffusion model. *arXiv preprint arXiv:2211.08332*, 2022.
- Hui Ye, Xiulong Yang, Martin Takac, Rajshekhar Sunderraman, and Shihao Ji. Improving text-to-image synthesis using contrastive learning. *arXiv preprint arXiv:2107.02423*, 2021.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023.
- Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 5907–5915, 2017.
- Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1947–1962, 2018.
- Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free adaption of clip for few-shot classification. In *European Conference on Computer Vision*, pp. 493–510. Springer, 2022.
- Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*, 2023a.
- Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Hao Dong, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*, 2023b.
- Zhu Zhang, Jianxin Ma, Chang Zhou, Rui Men, Zhikang Li, Ming Ding, Jie Tang, Jingren Zhou, and Hongxia Yang. M6-ufc: Unifying multi-modal controls for conditional image synthesis. *arXiv preprint arXiv:2105.14211*, 2021.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.
- Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5802–5810, 2019.

## A ADDITIONAL RELATED WORK

**Large Vision-Language Models (LVLMs).** Driven by the increasing diversity of large-scale data, developing powerful LVLMs has gained significant attention and progress in recent years. Early efforts, such as CLIP (Radford et al., 2021), ALIGN (Jia et al., 2021), and follow-up works (Li et al., 2021; Dong et al., 2023b; Zhang et al., 2022), adopt vision-language contrastive pre-training on extensive web-scale data, emerging superior generalization performance for zero-shot evaluation. With the popularity of large language models (LLMs) (OpenAI, 2023a;b), recent LVLMs tend to incorporate pre-trained LLMs with visual understanding capabilities. With advanced training strategies, BLIP series (Li et al., 2022; 2023) learn a Q-Former network to bridge between frozen image encoders and LLMs, which exhibit robust visual reasoning power. Trained by image-text interleaved data, Flamingo (Alayrac et al., 2022) obtains impressive few-shot learning capacity and enriches the display form of vision-language reasoning.

In contrast to the powerful but close-source GPT-4 (OpenAI, 2023b) and Bard (Google, 2023), a new branch of LVLMs is based on the open-source LLaMA (Touvron et al., 2023), and endows it with image understanding ability by visual instruction tuning. Therein, LLaMA-Adapter series (Zhang et al., 2023a; Gao et al., 2023; Han et al., 2023) introduce zero-initialized attention mechanisms, and conduct multi-modal parameter-efficient fine-tuning. LLaVA (Liu et al., 2023) introduces a high-quality visual instruction dataset to fully fine-tune the entire LLaMA, while MiniGPT-4 (Zhu et al., 2023) only adopts a projection layer for vision-language alignment. There are also many inspiring LVLM works for exploring different tuning strategies (Ye et al., 2023), collecting more diverse datasets (Chen et al., 2023), and incorporating multi-modality (Guo et al., 2023).

In this paper, as the first work, we leverage the robust vision-language reasoning of LVLMs to enhance compositional text-to-image generation. We select Bard developed by Google to first provide answers based on visual questioning, and then point out the misalignment between text prompts and generated images. Experiments have shown the effectiveness of our approach for unleashing the potential of LVLMs for improving compositional text-to-image generation.

## B ADDITIONAL VISULIZATION

A more detailed comparison of the results between Stable Diffusion and the fine-tuned model is illustrated in Figure 6, where the images generated from the fine-tuned model exhibit a higher degree of alignment with the input text.

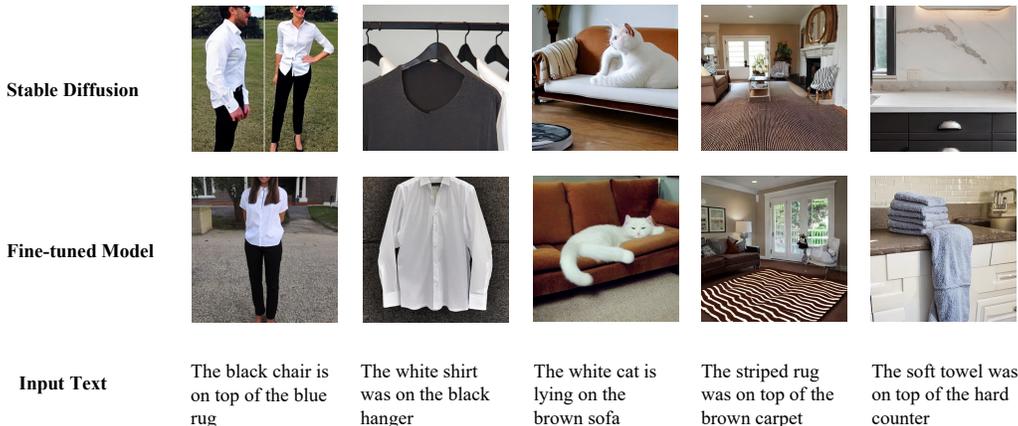


Figure 6: The images generated by Stable Diffusion and the fine-tuned model.