# The Fabrication of Reality and Fantasy: Scene Generation with LLM-Assisted Prompt Interpretation

Yi Yao<sup>1</sup>\*<sup>®</sup>, Chan-Feng Hsu<sup>1</sup>\*, Jhe-Hao Lin<sup>1</sup>, Hongxia Xie<sup>2</sup><sup>®</sup>, Terence Lin<sup>1</sup>, Yi-Ning Huang<sup>1</sup>, Hong-Han Shuai<sup>1</sup><sup>®</sup>, and Wen-Huang Cheng<sup>3</sup><sup>®</sup>

<sup>1</sup> National Yang Ming Chiao Tung University, Taiwan, {leo81005.ee10, cfhsu311510211.ee11, hhshuai}@nycu.edu.tw <sup>2</sup> Jilin University, China, <sup>3</sup> National Taiwan University, Taiwan {wenhuang@ntu.edu.tw}

Abstract. In spite of recent advancements in text-to-image generation, limitations persist in handling complex and imaginative prompts due to the restricted diversity and complexity of training data. This work explores how diffusion models can generate images from prompts requiring artistic creativity or specialized knowledge. We introduce the Realistic-Fantasy Benchmark (RFBench), a novel evaluation framework blending realistic and fantastical scenarios. To address these challenges, we propose the Realistic-Fantasy Network (RFNet), a training-free approach integrating diffusion models with LLMs. Extensive human evaluations and GPT-based compositional assessments demonstrate our approach's superiority over state-of-the-art methods. Our code and dataset is available at https://leo81005.github.io/Reality-and-Fantasy/.

**Keywords:** Text-to-image Generation · Realistic-Fantasy Benchmark · Diffusion Model · Large Language Models (LLMs)

## 1 Introduction

Considerable advancements have been made in the field of text-to-image generation, especially with the introduction of diffusion models, *e.g.*, Stable Diffusion [31], GLIDE [23], DALLE2 [30] and Imagen [32]. These models exhibit remarkable proficiency in generating diverse and high-fidelity images based on natural language prompts.

However, despite their impressive capabilities, diffusion models occasionally face challenges in accurately interpreting complex prompts that demand a deep understanding or specialized knowledge [39,41,43]. This limitation becomes particularly apparent with creative and abstract prompts, which require a nuanced

<sup>&</sup>lt;sup>\*</sup> These authors contributed equally to this work



Fig. 1: Text-to-image diffusion models such as Stable Diffusion [31] often struggle to accurately follow prompts that involve scientific and empirical reasoning, metaphorical thinking, role conflicting, or imaginative scenarios. Our method achieves enhanced prompt understanding capabilities and accurately follows these types of prompts.

grasp of context and subtleties. For example, in scenarios where the prompt involves unconventional scenarios like "a rat is hunting a lion", traditional diffusion models might not accurately represent the intended dynamics or relationships between entities (as shown in Fig.1).

A significant obstacle for traditional diffusion models in creating abstract images is the bias present within their training datasets [6, 25]. These datasets often do not include images of scenarios that defy conventional reality, such as a mouse hunting a lion. Traditionally, mitigating these challenges has required costly data collection and complex filtering, as well as model retraining or fine-tuning [8,12,40,42]. Research costs significantly increase due to these long and labor-intensive processes. Moreover, fine-tuning neural networks and model editing can lead to catastrophic forgetting and overall performance degradation [14,33]. Recently, it has been demonstrated that utilizing Large Language Models (LLMs) to aid in the generation process ensures the production of accurate details [5,16,17,28,36,38,40]. Directly integrating these models during the generation phase signifies a more efficient strategy.

In this paper, we want to address the question: how can generative models be improved to better capture imaginative and abstract concepts in images? In response to the existing gap in benchmarks for abstract and creative text-to-image synthesis, our work introduces a novel benchmark, **Realistic-Fantasy Benchmark (RFBench)**. This benchmark is designed to evaluate both *Realistic & Analytical* and *Creativity & Imagination* interpretations in generated images. The *Realistic & Analytical* category includes four sub-categories, focusing on the models' ability to adhere to realism and analytical depth. Images are generated in response to prompts that require not only precision in science but also cultural sensitivity and nuanced expression of symbolic meaning. On the other hand, *Creativity & Imagination*, is segmented into five specific sub-categories based on attribute distinctions, challenges models to navigate the complexities of generating images from prompts that necessitate a high degree of creativity and abstract reasoning.

To empower diffusion models with the capability to generate imaginative and abstract images, we introduce an innovative training-free approach **Realistic-Fantasy Network (RFNet)** that integrates diffusion models with LLMs. Given a prompt describing the desired image, the LLM generates an image layout, which includes bounding boxes for main subjects and background elements, along with textual details to support logic or interpret scientific data. To refine image generation, we further propose the Semantic Alignment Assessment (SAA), ensuring consistency with the scene's objects. This crucial step improves the final image quality. The enhanced details direct the diffusion model, enabling precise object placement through guidance constraints. Our method, leveraging pre-trained models, is compatible with independently trained LLMs and diffusion models, eliminating the need for parameter adjustments.

In summary, our key contributions are:

- 1. We have collected a novel **Realistic-Fantasy Benchmark (RFBench)**, which is a meticulously curated benchmark that stands out for its rich diversity of scenarios. It challenges and extends the boundaries of generative model creativity and inference capabilities, establishing a new standard for assessing imaginative data processing.
- 2. To empower diffusion models with the capability to generate imaginative and abstract images, we introduce an innovative training-free approach, **Realistic-Fantasy Network (RFNet)**, that integrates diffusion models with LLMs.
- 3. Through our proposed RFBench, extensive human evaluations coupled with GPT-based compositional assessments have demonstrated our approach's superiority over other state-of-the-art methods.

## 2 Related Work

#### 2.1 Text-guided diffusion models

Diffusion models, utilizing stochastic differential equations, have emerged as effective tools for generating realistic images [1,4,9,22,23,30–32]. DALL-E 2 [30] pioneered the approach of converting textual descriptions into joint image-text embeddings with the aid of CLIP [29]. GLIDE [23] demonstrated that classifier-free guidance [11] is favored by human evaluators over CLIP guidance for generating images based on text descriptions. Imagen [32] follows GLIDE but uses pretrained text encoder instead, further reducing negligible computation burden to the online training of the text-to-image diffusion prior, and can improve sample quality significantly by simply scaling the text encoder.

Although text-to-image capabilities have seen significant development, there has been limited focus on generating images involving *high levels of creativity*,

scientific principles, cultural references, and symbolic meanings. The primary reason is the data bias in the training dataset [21, 24, 34]. Several studies [19, 21, 25] have investigated the impact of data bias on diffusion models, particularly in the context of Text-to-Image generation. Perera *et al.* [25] investigates the bias exhibited by diffusion models across various attributes in face generation. Luccioni *et al.* [19] evaluates bias levels in text-to-image systems regarding gender and ethnicity.

In this work, we introduce a new task: reality and fantasy scene generation. Recognizing the absence of a dedicated evaluation framework for such tasks, we introduce a new benchmark, the Realistic-Fantasy Benchmark (RFBench), which blends scenarios from both realistic and fantastical realms.

#### 2.2 LLMs for image generation

Recently, researchers have explored using LLMs to provide guidance or auxiliary information for text-to-image generation systems [7, 15, 16, 20, 28, 36, 38]. In LMD [16], foreground objects are identified using LLMs, and then images are generated based on the layout determined by the diffusion model. Phung *et al.* [26] proposes attention-refocusing losses to constrain the generated objects on their assigned boxes generated by LLMs. LVD [17] requires LLMs to generate continuous spatial constraints to accomplish video generation. Besides using LLMs to generate spatial layout from user prompts, some studies [28,36,38] investigate integrating LLMs directly into the image generation pipeline. SLD [36] integrates open-vocabulary object detection with LLMs to enhance image editing. RPG [38] integrates LLMs in a closed-loop manner, allowing generated images to continuously improve through LLMs feedback, and uses Chain-of-Thought [35] to further improve generation quality.

As a result of these developments, LLMs can be incorporated into pipelines for the generation of images. In this work, we use LLMs to uncover and elaborate upon the complexities embedded within complex and abstract prompts.

## 3 Our Proposed Realistic-Fantasy Benchmark

In this study, we explore how diffusion models can effectively process and generate imagery from prompts that pose significant challenges due to their reliance on creative thinking or specialized knowledge. Recognizing the absence of a dedicated evaluation framework for such tasks, we introduce a new benchmark, the **Realistic-Fantasy Benchmark (RFBench)**, which blends scenarios from both realistic and fantastical realms.

Benchmark Collection. We focus on two main categories, each with distinct subcategories, Realistic & Analytical and Creativity & Imagination, totaling nine subcategories. Each sub-category is meticulously crafted with around 25 text prompts, leading to an aggregate of 229 unique compositional text prompts designed to test the models against both conventional and unprecedented creative challenges. The collection process, outlined in Fig. 2, employs 
 Table 1: Categories of our proposed Realistic-Fantasy Benchmark (RFBench). The full dataset are show in our supplementary material.

Categories	Definition	Example		
	Realistic & Analytical			
To evaluate models' comp	prehensive understanding of complex ideas, factual	accuracy, and cultural insights		
Scientific and Empirica	Relates to hypothesis testing, deduction, and sci-	"A drop of water on the Interna-		
Reasoning	entific methodology	tional Space Station."		
Cultural and Tempora	Necessitates understanding and knowledge of par-	"Children in costumes going		
Awareness	ticular cultural or historical events	door-to-door on October 31st."		
Factual or Literal Descrip	-Evaluates the models' capacity to generate images	"A tank that's been sitting on the		
tions	that adhere closely to factual accuracy and real-	beach for 50 years."		
	istic depiction	1		
Conceptual and Metaphori	-Focuses on the models' ability to comprehend and	"A man is as brave as a lion."		
cal Thinking	depict the underlying symbolic messages within	1		
	the prompts	1		
	Creativity & Imagination			
To evaluat	e model's capability in employing creativity and a	bstract thinking		
Common Objects in Un	- Challenges models' capacity to maintain object	"A rubber duck sailing across a		
usual Contexts	integrity while adapting to surreal environments	field of hot lava."		
Imaginative Scenarios	Evaluate the models' abilities to craft scenes in-	"An octopus playing chess with a		
	volving animals or humans in fantastical or un-	seahorse."		
	likely scenarios	1		
Counterfactual Scenarios	Focuses on scenarios that defy conventional ex-	"Fish swimming in the clouds."		
	pectations of reality	1		
Role Reversal or Conflicting	Sistimulates consideration of perspectives and sce-	"A cat is chased by a mouse."		
	narios where typical roles or expectations are in-	I		
	verted	1		
Anthropomorphic Scenarios	Examines the model's ability to imbue inanimate	"A snowman building a friend in		
	objects or phenomena with human-like character-	the blizzard."		
	istics	1		

a hybrid method combining in-context learning and predefined rules, leveraging powerful language models such as ChatGPT and Bard for diverse text prompts creation. By alternating between these models, we achieve a diverse set of responses, capitalizing on the distinct advantages of each LLM. It boosts the variety and complexity of prompts while reducing the reliance on manual labeling.



Fig. 2: The collection pipeline of our proposed RFBench.

**Realistic & Analytical Category.** There are four sub-categories: *Scientific and Empirical Reasoning, Cultural and Temporal Awareness, Factual or Literal Descriptions, and Conceptual and Metaphorical Thinking* (details are shown in the upper part of Tab. 1). These sub-categories are anchored in real-world con-



Fig. 3: Overview of our proposed Realistic-Fantasy Network (RFNet). In stage 1, the user's input prompt is first processed by a LLM to extract the layout and descriptions. The descriptions then go through a text encoder, which is the text-processing component of the CLIP model, and are refined by the SAA to form a better prompt. In stage 2, the refined prompts are fed into the diffusion model for in-depth object generation, which creates each target object with precision. The resulting cross-attention map and mask latent are then utilized for seamless background integration, merging objects into one single image.

texts, emphasizing logical reasoning, accurate data, and an understanding of cultural or historical contexts. They contain scientific exploration, realistic descriptions, and culturally symbolic narratives. This demands that the models not only draw from an extensive knowledge pool but also demonstrate an ability to grasp and articulate underlying concepts.

**Creativity & Imagination Category**. It consists of five sub-categories: Common Objects in Unusual Contexts, Imaginative Scenarios, Counterfactual Scenarios, Role Reversal or Conflicting, and Anthropomorphic Scenarios (details are shown in the lower part of Tab. 1). This evaluation focuses on the model's capacity to innovatively repurpose familiar objects, attribute human-like characteristics to inanimate objects, and generate novel environments for everyday items. This category tests the model's out-of-the-box thinking and imaginative capabilities.

# 4 Our Proposed Realistic-Fantasy Network

In this section, we propose a **Realistic-Fantasy Network (RFNet)** for the benchmark scenario we proposed in the previous section. To thoroughly interpret

the details from the input prompt, we divide our approach into two stages, as shown in Fig. 3. In the first stage, we transform the initial input prompt into a refined version specifically tailored for image generation by LLMs. In the second stage, we utilize a diffusion model through a two-step process to generate outputs with extraordinary details.

## 4.1 LLM-Driven Detail Synthesis

In the first stage of our methodology, we concentrate on utilizing LLMs to uncover and elaborate on the intricacies embedded within the user's input prompt. This process involves specifying task requirements to more accurately define the task and incorporating in-context learning to enhance understanding for LLMs. The enriched response from the LLM encompasses additional information, such as *layout, detailed descriptions, background scenes,* and *negative prompts*<sup>4</sup>. This step is crucial as it aims to mitigate the primary challenge we seek to overcome: the training data bias inherent in current diffusion models. By leveraging the pre-trained LLM for logical reasoning and conjecture, we aim to compensate for the gaps left by these biases, ensuring a more accurate and coherent image generation process.

## 4.2 Semantic Alignment Assessment

As we proceed with generating images using the diffusion model using the details generated by the previous step, there is a critical challenge: the description lists generated by LLMs for one object usually overlook the relationships among them. For example, interpretations of "a lion" could range from being "unaware and asleep" to "frightened and trying to escape." Although both depictions are valid, descriptions such as "unaware" and "trying to escape" can lead to conflicting interpretations, thus complicating the image generation process.

To overcome this challenge, we introduce the **Semantic Alignment Assessment (SAA)** module. This module calculates the relevance between different object vectors, thereby selecting the candidate description that best fits the current scenario. By conducting the cosine similarity among different descriptions, we can navigate the complexities introduced by the LLM's output, selecting the most compatible details for the diffusion model. This step is crucial for maintaining the coherence and accuracy of the generated images, highlighting our novel approach to mitigating the risk of conflicting descriptions. Through this module, we ensure textual precision and compatibility, and provide *clear, consistent instructions* for the subsequent diffusion model to generate visually coherent representations.

### 4.3 Comprehensive Image Synthesis

In the second stage of our proposed RFNet, following LMD [16], we propose a two-step generation process for imaginative and abstract concepts. As shown

<sup>&</sup>lt;sup>4</sup> One detailed sample can be found in our supplementary material.



Fig. 4: Comprehensive Image Synthesis. In step 1, utilizing the prompt refined by the SAA module, the frozen stable diffusion model generates each foreground object independently. During the denoising phase, the cross-attention map is extracted and saved for Guidance Constraint in the next step. In addition, a Suppression Constraint is also added in step 2 to minimize influence between different objects.

in Fig. 4, in the first step, we focus on generating each *foreground* object with comprehensive details. In the second step, we integrate the objects generated in the first step into corresponding *background* derived from the initial prompt. This structured approach ensures a cohesive integration of detailed foreground objects into a contextually relevant background, enhancing the overall effectiveness of our framework.

## Step 1: In-Depth Object Generation.

We re-organize the SAA description lists to consider both the layout of specific objects and their descriptions. By concatenating the background prompt with the target object and its relevant descriptions, we set up the input prompt as "[background prompt] with [target object], [descriptions]" (e.g., "A grassland scene with a rat, roaring, with big mouth and sharp teeth, leap out at..."). Following LMD [16], the initial latent representation for each target object is fixed to facilitate the fusion of various objects into a cohesive background scene.

During generation, the diffusion model uses cross-attention layers to manage the influence of textual information on the visual output, allowing precise control over image details. The cross-attention map's constraint function integrates objects within the bounding box by enhancing cross-attention inside the box for accurate object representation while minimizing it outside the box. This function guides the update of the noised latent vector during denoising to ensure spatial conditions match predefined specifications. The constraint function is defined as:

$$\mathcal{L}_{obj}(\mathbf{A}, i, v) = [1 - \operatorname{Topk}_{u}(\mathbf{A}_{uv} \cdot \mathbf{m}_{i})] + [\operatorname{Topk}_{u}(\mathbf{A}_{uv} \cdot (1 - \mathbf{m}_{i}))], \quad (1)$$

where  $\mathbf{m}_i$  denotes the binary mask of the bounding box associated with object i, performing element-wise multiplication over the cross-attention map  $\mathbf{A}$ . The cross-attention map  $\mathbf{A}$  is aggregated by summing the contributions across all layers. For object i, the operation Topk<sub>u</sub> computes the mean of the top-k values within the spatial dimension u. Prior to each denoising step, the latent is refined by minimizing the constraint function:

$$z'_t \leftarrow z_t - \alpha \cdot \nabla_{z_t} \sum_{v \in V} \mathcal{L}(\mathbf{A}, i, v),$$
 (2)

$$z_{t-1} \leftarrow DiffusionStep(z'_t, \mathcal{P}^{(i)}), \tag{3}$$

where  $\alpha$  denotes the hyperparameter that controls the magnitude of the gradient update, and V contains the set of token indices for the target object in the prompt. As in the diffusion step, the updated  $z'_t$  along with the modified prompt  $\mathcal{P}^{(i)}$  of object *i*, are served as the inputs to the diffusion model.

After the generation, the cross-attention map derived from each target object is then converted into a saliency mask. This mask is applied to the latent representation of the target object through element-wise multiplication at each step of the denoising process. Both the cross-attention map and the masked latent representation of the target object between each denoising step are transmitted to the next step for background integration.

Step 2: Seamless Background Integration. This step involves fusing the generated results with the background while preserving the high-quality generation achieved in the first step. To accomplish this, we first replace the generated latent  $z'_t$  with the masking latent  $z'_t^{(masked,i)}$  for each object *i*:

$$z'_t \leftarrow Replacement(z'_t, z^{(masked, i)}_t), \ \forall i.$$
 (4)

Following the approach established in Step 1, the initial latent representation  $z'_T$  is also fixed with the initial latent of each target object in the first step. This alignment ensures a seamless integration of the object into the background. The replacement approach is performed within timestep rT, where  $r \in [0, 1]$ , reserving principles from LMD [16]. According to LMD, the diffusion model determines the position of the objects in the early denoising steps, while adjusting details in the later steps. This helps us to preserve exceptional control over the layout.

Furthermore, we incorporate a specialized constraint function designed to enhance the integration of generated objects with their background, distinguished by two key components: **guidance constraint** and **suppression constraint**. As shown in Fig. 4, the guidance constraint is engineered to reduce the cross-attention within each bounding box relative to the original object's attention. With the purpose of seamlessly integrating with the detailed object generated in the first step. Conversely, the *suppression constraint* works to minimize cross-attention outside the bounding box, thereby mitigating interference among multiple objects when processed together, as illustrated in Eq. (5). These constraint functions mark a departure from conventional methods that predominantly use loss to fix the layout.

$$\mathcal{L}_{bg}(\mathbf{A}', \mathbf{A}^{(i)}, i, v) = \underbrace{\beta \cdot \sum_{u} \left| (\mathbf{A}_{uv}' - \mathbf{A}_{uv}^{(i)}) \cdot \mathbf{m}_{i} \right|}_{\text{guidance constraint}} + \underbrace{\gamma \cdot \operatorname{Topk}_{u}(\mathbf{A}_{uv}' \cdot (1 - \mathbf{m}_{i}))}_{\text{suppression constraint}}, \forall i,$$
(5)

where  $\mathbf{A}'$  represents the cross-attention map post the substitution procedure and  $\mathbf{A}^{(i)}$  is the cross-attention map of object *i* extracted from the diffusion step in Eq. (3). The hyperparameters  $\beta$  and  $\gamma$  indicate the intensity of guidance constraint and suppression constraint, respectively. It is noteworthy that adjusting  $\beta$  amplifies the importance of the guidance loss within the overall loss function. This adjustment ensures the latent representation is precisely aligned with the object generated in the initial phase, thereby guaranteeing a high level of accuracy and consistency in the integration process. Upon completion of all denoising steps, the latent  $z_0$  is fed into the decoder to produce the final image.

Our strategy focuses on maintaining the integrity and coherence of the foreground objects generated, emphasizing the preservation of their quality and interaction with the background. By doing so, the generated visual elements are fidelity-aware and contextually appropriate.

## 5 Experiments

#### 5.1 Implementation Details

**Experimental Setup.** In this work, we choose versions 1.4 and 2.1 of Stable Diffusion [31] as the text-to-image baseline model. The number of denoising steps is set as 50 with a fixed guidance scale of 7.5, and the synthetic images are in a resolution of  $512 \times 512$ . All experiments are conducted on the NVIDIA RTX 3090 GPU with 24 GB memory.

**Evaluation Metrics.** We generate 32 images for each text prompt in **RF-Bench** for automatic evaluation. We selected the following two metrics: (1) GPT4-CLIP <sup>5</sup>. By utilizing GPT4 for captioning and calculating CLIP text-text cosine similarity, GPT4-CLIP ensures a more precise reflection of the intended meanings between images and prompts. (2) GPT4Score. Inspired by [13,18], we adopt GPT4Score to evaluate image alignment with text prompts, where GPT4 rates images on a 0-100 scale based on their fidelity to the prompts, enabling precise assessment of model-generated visuals against specified criteria<sup>6</sup>.

**Comparison with Existing Methods.** We benchmark our proposed RFNet against various open-source scene generation methods, including Stable Diffusion [31], Attend and Excite [3], LMD [16], BoxDiff [37], MultiDiffusion [2], and SDXL [27]. Notably, all methods, including ours, utilize Stable Diffusion 2.1 as the foundational model, ensuring a fair comparison.

<sup>&</sup>lt;sup>5</sup> We adopt GPT4-CLIP due to BLIP-CLIP's [3] limitations in accurately capturing image meanings through generated captions.

<sup>&</sup>lt;sup>6</sup> The widely recognized metric, CLIPScore [10, 29] exhibits limitations in evaluating our task. For detailed examples, please see the supplementary materials.

**Table 2:** Benchmarking on GPT4-CLIP and GPT4Score. R & A, C & I, Avg represent Realistic & Analytical category, Creativity & Imagination category, and average of both categories, respectively. The red text indicates the improvement ratio of our method compared to Stable Diffusion.

	GPT4-CLIP			GPT4Score		
Model	R & A	C & I	Avg	R & A	С&І	Avg
Stable Diffusion [31]	0.573	0.552	0.561	0.667	0.440	0.541
MultiDiffusion [2]	0.510	0.510	0.510	0.517	0.493	0.504
Attend and Excite [3]	0.523	0.560	0.546	0.633	0.520	0.570
LLM-groundedDiffusion [16]	0.457	0.536	0.501	0.550	0.600	0.578
BoxDiff [37]	0.532	0.553	0.543	0.583	0.520	0.548
SDXL [27]	0.536	0.619	0.582	0.567	0.587	0.578
RFNet(ours)	0.587 <mark>(2%</mark> ↑)	0.623 (13%)	0.607 (8%)	0.833 (25%)	$0.627(43\%\uparrow)$	0.719 (33%†)

#### 5.2 Quantitative Evaluation

**Evaluation on RFBench.** As evidenced in Tab. 2, our approach significantly outperforms other methods for both Realistic & Analytical and Creativity & Imagination tasks. For Realistic & Analytical task, our method seamlessly integrates LLM-based insights, achieving a remarkable accuracy improvement. Unlike Attend-and-excite, which focuses on semantic guidance, our method ensures precise adherence to detailed and complex prompt requirements. For the Creativity & Imagination, which demands high degrees of creativity and abstract conceptualization, our method outperforms others by not only adhering to the imaginative aspects of prompts but also maintaining coherent structure and contexts. For instance, SDXL, while adept at high-resolution image synthesis, occasionally lacks in capturing the nuanced creativity intended in prompts; our method fills this gap effectively. Similarly, LMD, though enhancing prompt understanding through LLMs, sometimes struggles with the scientific reasoning required for Realistic& Analytical tasks.

Notably, for Realistic & Analytical category, our approach shows a 61% performance increase over MultiDiffusion on GPT4Score. Meanwhile, in Creativity & Imagination task, we observe a substantial enhancement, outperforming Stable Diffusion by over 43%. In light of the above, our method is unique in its ability to bridge the gap between realistic reasoning and imagination, creating a new benchmark for text-to-image generation.

**Evaluation on DrawBench**. We also evaluate our method on Draw-Bench [32], a comprehensive and challenging benchmark for text-to-image models. Similar to us, DrawBench also includes some Creativity & Imagination prompts, and we evaluate our method with Imagen [32] on these prompts. As shown in Tab. 3, our approach significantly outperforms Imagen on most prompt settings, demon-

**Table 3:** GPT4Score comparison with Imagen on DrawBench subset.

Prompt	Imagen	Ours
A bird scaring a scarecrow	0.069	0.275
A blue coloured pizza	0.425	0.125
A fish eating a pelican	0.000	0.000
A horse riding an astronaut	0.000	0.000
A panda making latte art	0.050	0.250
A pizza cooking an oven	0.700	0.831
A shark in the desert	0.194	0.713
An elephant under the sea	0.300	0.900
Hovering cow abducting aliens	0.025	0.144
Rainbow coloured penguin	0.394	0.519
1 0		

strating the generalization ability of our model.



Fig. 5: Qualitative comparison on RFBench. The compared models include (a) Stable Diffusion, (b) MultiDiffusion, (c) Attend and Excite, (d) LMD, (e) BoxDiff, (f) SDXL, (g) Ours (Best viewed in color and zoom in. More samples can be found in our supplementary material.)

#### 5.3 Qualitative Evaluation

In the qualitative comparison of text-guided image generation, we select some advanced baseline methods, including Attend-and-excite [3], BoxDiff [37], LMD [16], MultiDiffusion [2], and SDXL [27]. Attend-and-excite focuses on enhancing the semantic understanding of prompts through attention mechanisms, while BoxDiff introduces a novel approach to text-to-image synthesis with box-constrained diffusion without the need for explicit training. MultiDiffusion proposes a method for fusing multiple diffusion paths to achieve greater control over the image generation process, and SDXL aims at improving the capabilities of latent diffusion models for synthesizing high-resolution images. As shown in Fig. 5, our method, produces more precise editing results than the aforementioned methods. This is attributed to our **In-Depth Object Generation** and **Seamless Background Integration** strategy. It ensures outstanding fidelity in outcomes and flawlessly retains the semantic structure of the source image, highlighting our approach's superior capability in complex editing tasks.

#### 5.4 User study

Through an extensive user study, we benchmarked our model against other methods to assess real human preferences for the generated images. Utilizing our newly proposed benchmark, the RFBench, we selected a diverse set of 27 prompts and generated six images per prompt to ensure a broad representation of the model's capabilities. Detailed feedbacks were collected from 120 participants, evaluating each image for visual quality and text prompt fidelity <sup>7</sup>. These criteria are critical, which measure the image's quality and correctness of semantics in the synthesized image. Participants rated images on a scale from  $\{1, 2, 3, 4, 5\}$ , with scores normalized by dividing by 5. We calculated the average score across all images and participants.

As illustrated in Fig. 6, participants uniformly favored our model's output, recognizing it as superior in both quality and alignment with the textual descriptions.



Fig. 6: Comparison between our method and other advanced methods on RFBench. The image-text *Alignment* and *Fidelity* of our method are highly preferred by users.

### 5.5 Ablation study

Impact of Various Constraints. To validate the impact of guidance constraint and suppression constraint, we perform ablation studies on different combinations of constraints, and the results are listed in Tab. 4. As shown, the baseline model (Stable Dif-

Table 4:	Ablation	studies	on	various	com-
ponents o	n RFBen	ch.			

SAA	guidance	suppression	GPT4Score
			0.295
		$\checkmark$	0.407
	$\checkmark$		0.554
	$\checkmark$	$\checkmark$	0.572
$\checkmark$	$\checkmark$	$\checkmark$	0.719

fusion) achieves a 0.295 in terms of GPT4Score without any constraints. As guidance constraint and suppression constraint work complementary to restrict the cross-attention of objects inside the conditional boxes, a higher GPT4Score

<sup>&</sup>lt;sup>7</sup> Details of survey samples can be found in our supplementary material.

of 0.572 is achieved on the generated images. Both proposed constraints are effective in controlling image quality and layout of synthesized foreground objects. **Impact of Semantic Alignment Assessment (SAA) Module.** As aforementioned, using conflict descriptions in the denoise step may potentially affect image synthesis. The quantitative evaluation is presented in Tab. 4. In the absence of SAA, the model attains a GPT4Score of 0.572. Similarly, with SAA, the model reaches a GPT4Score of 0.719. This indicates a lack of consistency between the semantics generated in the images and the provided text prompts, leading to a reduction in image quality. It is important to note that the inclusion of SAA significantly enhances the clarity of the images obtained. One visual illustration can be found in Fig. 7.



**Fig. 7:** The illustration of final generated results of minimum (min) and maximum (max) similarity descriptions during the SAA. It can be observed that prompts with higher similarity yield images of higher quality, which is significantly better than the one with the lowest similarity.

# 6 Conclusion and Future Work

In this research, we present a novel challenge: generating scenes that blend reality and fantasy. We investigate the capacity of diffusion models to create visuals from prompts that demand high levels of creativity or specific knowledge. Noting the lack of a specific evaluation mechanism for such tasks, we establish the Realistic-Fantasy Benchmark (RFBench), combining elements of both realistic and imaginary scenarios. To address the task of generating realistic and fantastical scenes, we introduce a unique, training-free, two-tiered method, Realistic-Fantasy Network (RFNet), that combines diffusion models with large language models (LLMs). Our approach, evaluated through the RFBench using thorough human assessments and GPT-based compositional evaluations, has proven to be superior to existing cutting-edge techniques. Given the novelty of our task, future research could develop additional evaluation metrics beyond those used in this study, enhancing the assessment of generated scenes.

## Acknowledgment

This work is partially supported by the National Science and Technology Council, Taiwan under Grants NSTC-112-2221-E-A49-059-MY3 and NSTC-112-2221-E-A49-094-MY3 as well as the partial support from Ministry of Science and Technology of Taiwan under the grant numbers: MOST-109-2221-E-009-114-MY3, MOST-110-2221-E-A49-164, MOST-109-2223-E-009-002-MY3, MOST-110-2218-E-A49-018 and MOST-111-2634-F-007-002.

## References

- Anciukevičius, T., Xu, Z., Fisher, M., Henderson, P., Bilen, H., Mitra, N.J., Guerrero, P.: Renderdiffusion: Image diffusion for 3d reconstruction, inpainting and generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12608–12618 (2023)
- Bar-Tal, O., Yariv, L., Lipman, Y., Dekel, T.: Multidiffusion: Fusing diffusion paths for controlled image generation. In: International Conference on Machine Learning (2023)
- Chefer, H., Alaluf, Y., Vinker, Y., Wolf, L., Cohen-Or, D.: Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. ACM Transactions on Graphics (TOG) 42(4), 1–10 (2023)
- Feng, W., He, X., Fu, T.J., Jampani, V., Akula, A.R., Narayana, P., Basu, S., Wang, X.E., Wang, W.Y.: Training-free structured diffusion guidance for compositional text-to-image synthesis. In: International Conference on Learning Representations (2023), https://openreview.net/forum?id=PUIqjT4rzq7
- Feng, W., Zhu, W., Fu, T.j., Jampani, V., Akula, A., He, X., Basu, S., Wang, X.E., Wang, W.Y.: Layoutgpt: Compositional visual planning and generation with large language models. In: Advances in Neural Information Processing Systems. vol. 36 (2024)
- Friedrich, F., Brack, M., Struppek, L., Hintersdorf, D., Schramowski, P., Luccioni, S., Kersting, K.: Fair diffusion: Instructing text-to-image generation models on fairness. arXiv preprint arXiv:2302.10893 (2023)
- Gani, H., Bhat, S.F., Naseer, M., Khan, S., Wonka, P.: Llm blueprint: Enabling text-to-image generation with complex and detailed prompts. In: International Conference on Learning Representations (2024)
- Golnari, P.A.: Lora-enhanced distillation on guided diffusion models. arXiv preprint arXiv:2312.06899 (2023)
- Gong, J., Foo, L.G., Fan, Z., Ke, Q., Rahmani, H., Liu, J.: Diffpose: Toward more reliable 3d pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13041–13051 (2023)
- Hessel, J., Holtzman, A., Forbes, M., Bras, R.L., Choi, Y.: Clipscore: A referencefree evaluation metric for image captioning. arXiv preprint arXiv:2104.08718 (2021)
- Ho, J., Salimans, T.: Classifier-free diffusion guidance. In: Advances in Neural Information Processing Systems Workshop (2021), https://openreview.net/forum? id=qw8AKxfYbI
- Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. In: International Conference on Learning Representations (2022)

- 16 Y. Yao et al.
- Huang, K., Sun, K., Xie, E., Li, Z., Liu, X.: T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation. In: Advances in Neural Information Processing Systems. vol. 36 (2024)
- Kemker, R., McClure, M., Abitino, A., Hayes, T., Kanan, C.: Measuring catastrophic forgetting in neural networks. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018)
- 15. Leigang Qu, Shengqiong Wu, H.F.L.N.T.S.C.: Layoutllm-t2i: Eliciting layout guidance from llm for text-to-image generation. In: Proceedings of the ACM International Conference on Multimedia (2023)
- Lian, L., Li, B., Yala, A., Darrell, T.: Llm-grounded diffusion: Enhancing prompt understanding of text-to-image diffusion models with large language models. arXiv preprint arXiv:2305.13655 (2023)
- 17. Lian, L., Shi, B., Yala, A., Darrell, T., Li, B.: Llm-grounded video diffusion models. arXiv preprint arXiv:2309.17444 (2023)
- Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. In: Advances in Neural Information Processing Systems (2023)
- Luccioni, S., Akiki, C., Mitchell, M., Jernite, Y.: Stable bias: Evaluating societal representations in diffusion models. In: Advances in Neural Information Processing Systems. vol. 36 (2024)
- Mantri, K.S.I., Sasikumar, N.: Interactive fashion content generation using llms and latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshop (2023)
- Naik, R., Nushi, B.: Social biases through the text-to-image generation lens. In: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society. p. 786–808. AIES '23, Association for Computing Machinery, New York, NY, USA (2023). https://doi.org/10.1145/3600211.3604711
- Nair, N.G., Cherian, A., Lohit, S., Wang, Y., Koike-Akino, T., Patel, V.M., Marks, T.K.: Steered diffusion: A generalized framework for plug-and-play conditional image synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 20850–20860 (2023)
- Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In: Proceedings of Machine Learning Research. pp. 16784–16804 (2022)
- Orgad, H., Kawar, B., Belinkov, Y.: Editing implicit assumptions in text-to-image diffusion models. In: 2023 IEEE/CVF International Conference on Computer Vision. pp. 7030–7038. IEEE Computer Society, Los Alamitos, CA, USA (oct 2023). https://doi.org/10.1109/ICCV51070.2023.00649
- Perera, M.V., Patel, V.M.: Analyzing bias in diffusion-based face generation models. arXiv preprint arXiv:2305.06402 (2023)
- 26. Phung, Q., Ge, S., Huang, J.B.: Grounded text-to-image synthesis with attention refocusing. arXiv preprint arXiv:2306.05427 (2023)
- Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. In: International Conference on Learning Representations (2024), https: //openreview.net/forum?id=di52zR8xgf
- Qin, J., Wu, J., Chen, W., Ren, Y., Li, H., Wu, H., Xiao, X., Wang, R., Wen, S.: Diffusiongpt: Llm-driven text-to-image generation system. arXiv preprint arXiv:2401.10061 (2024)

- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical textconditional image generation with clip latents. arXiv preprint arXiv:2204.06125 1(2), 3 (2022)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10684–10695 (2022)
- 32. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-toimage diffusion models with deep language understanding. In: Advances in Neural Information Processing Systems. vol. 35, pp. 36479–36494 (2022)
- Smith, J.S., Hsu, Y.C., Zhang, L., Hua, T., Kira, Z., Shen, Y., Jin, H.: Continual diffusion: Continual customization of text-to-image diffusion with c-lora. arXiv preprint arXiv:2304.06027 (2023)
- Su, X., Ren, Y., Qiang, W., Song, Z., Gao, H., Wu, F., Zheng, C.: Unbiased image synthesis via manifold-driven sampling in diffusion models. arXiv preprint arXiv:2307.08199 (2023)
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. In: Advances in Neural Information Processing Systems. vol. 35, pp. 24824–24837 (2022)
- Wu, T.H., Lian, L., Gonzalez, J.E., Li, B., Darrell, T.: Self-correcting llm-controlled diffusion models. arXiv preprint arXiv:2311.16090 (2023)
- Xie, J., Li, Y., Huang, Y., Liu, H., Zhang, W., Zheng, Y., Shou, M.Z.: Boxdiff: Textto-image synthesis with training-free box-constrained diffusion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7452–7461 (2023)
- Yang, L., Yu, Z., Meng, C., Xu, M., Ermon, S., Cui, B.: Mastering text-to-image diffusion: Recaptioning, planning, and generating with multimodal llms. arXiv preprint arXiv:2401.11708 (2024)
- Yang, L., Zhang, Z., Song, Y., Hong, S., Xu, R., Zhao, Y., Zhang, W., Cui, B., Yang, M.H.: Diffusion models: A comprehensive survey of methods and applications. ACM Comput. Surv. 56(4) (2023). https://doi.org/10.1145/3626235
- Yang, Z., Wang, J., Gan, Z., Li, L., Lin, K., Wu, C., Duan, N., Liu, Z., Liu, C., Zeng, M., Wang, L.: Reco: Region-controlled text-to-image generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition. pp. 14246–14255 (2023)
- Zhang, C., Zhang, C., Zhang, M., Kweon, I.S.: Text-to-image diffusion model in generative ai: A survey. arXiv preprint arXiv:2303.07909 (2023)
- Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023)
- Zhang, T., Wang, Z., Huang, J., Tasnim, M.M., Shi, W.: A survey of diffusion based image generation models: Issues and their solutions. arXiv preprint arXiv:2308.13142 (2023)