

Progressive Text-to-Image Diffusion with Soft Latent Direction

Yuteng Ye, Jiale Cai, Hang Zhou, Guanwen Li,
Youjia Zhang, Zikai Song, Chenxing Gao, Junqing Yu, Wei Yang*

Huazhong University of Science and Technology, Wuhan, China
{yuteng_ye, jaile_cai, henrryzh, juggle_lee, youjiazhang, skyesong, cxg, yjqing, weiyangcs}@hust.edu.cn

Abstract

In spite of the rapidly evolving landscape of text-to-image generation, the synthesis and manipulation of multiple entities while adhering to specific relational constraints pose enduring challenges. This paper introduces an innovative progressive synthesis and editing operation that systematically incorporates entities into the target image, ensuring their adherence to spatial and relational constraints at each sequential step. Our key insight stems from the observation that while a pre-trained text-to-image diffusion model adeptly handles one or two entities, it often falters when dealing with a greater number. To address this limitation, we propose harnessing the capabilities of a Large Language Model (LLM) to decompose intricate and protracted text descriptions into coherent directives adhering to stringent formats. To facilitate the execution of directives involving distinct semantic operations—namely insertion, editing, and erasing—we formulate the Stimulus, Response, and Fusion (SRF) framework. Within this framework, latent regions are gently stimulated in alignment with each operation, followed by the fusion of the responsive latent components to achieve cohesive entity manipulation. Our proposed framework yields notable advancements in object synthesis, particularly when confronted with intricate and lengthy textual inputs. Consequently, it establishes a new benchmark for text-to-image generation tasks, further elevating the field’s performance standards.

Introduction

Text-to-image generation is a vital and rapidly evolving field in computer vision that has attracted unprecedented attention from both researchers and the general public. The remarkable advances in this area are driven by the application of state-of-the-art image-generative models, such as auto-regressive (Ramesh et al. 2021; Wang et al. 2022) and diffusion models (Ramesh et al. 2022; Saharia et al. 2022; Rombach et al. 2022), as well as the availability of large-scale language-image datasets (Sharma et al. 2018; Schuhmann et al. 2022). However, existing methods face challenges in synthesizing or editing multiple subjects with specific relational and attributive constraints from textual prompts (Chefer et al. 2023). The typical defects that oc-

cur in the synthesis results are missing entities, and inaccurate inter-object relations, as shown in ???. Existing work improves the compositional skills of text-to-image synthesis models by incorporating linguistic structures (Feng et al. 2022), and attention controls (Hertz et al. 2022; Chefer et al. 2023) within the diffusion guidance process. Notably, Structured Diffusion (Feng et al. 2022) parse a text to extract numerous noun phrases, Attend-and-Excite (Chefer et al. 2023) strength attention activations associated with the most marginalized subject token. Yet, these remedies still face difficulties when the text description is long and complex, especially when it involves two and more subjects. Furthermore, users may find it necessary to perform subtle modifications to the unsatisfactory regions of the generated image, while preserving the remaining areas.

In this paper, we propose a novel progressive synthesizing/editing operation that successively incorporates entities, that conform to the spatial and relational constraint defined in the text prompt, while preserving the structure and aesthetics in each step. Our intuition is based on the observation that text-to-image models tend to better handle short-sentence prompts with a limited number of entities (1 or 2) than long descriptions with more entities. Therefore, we can parse the long descriptions into short-text prompts and craft the image progressively via a diffusion model to prevent the leakage and missing of semantics.

However, applying such a progressive operation to diffusion models faces two major challenges:

- The absence of a unified method for converting the integrated text-to-image process into a progressive procedure that can handle both synthesis and editing simultaneously. Current strategies can either synthesize (Chefer et al. 2023; Ma et al. 2023) or edit (Kawar et al. 2023; Goel et al. 2023; Xie et al. 2022; Avrahami, Fried, and Lischinski 2022; Yang et al. 2023), leaving a gap in the collective integration of these functions.
- The need for precise positioning and relational entity placement. Existing solutions either rely on user-supplied masks for entity insertion, necessitating manual intervention (Avrahami, Fried, and Lischinski 2022; Nichol et al. 2021), or introduce supplementary phrases to determine the entity editing direction (Hertz et al. 2022; Brooks, Holynski, and Efros 2023), which inadequately addressing spatial and relational dynamics.

*indicates corresponding author.

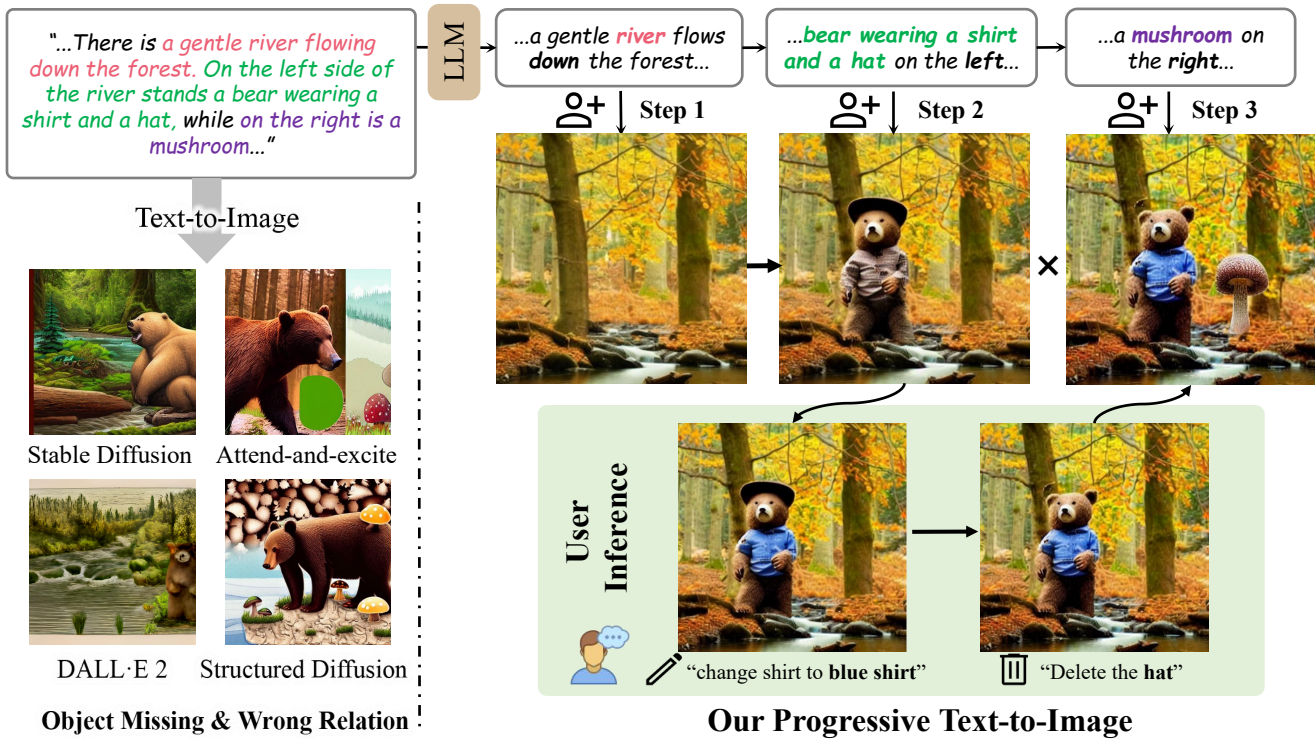


Figure 1: Existing text-to-image synthesis approaches struggle with textual prompts involving multiple entities and specified relational directions. We propose to decompose the protracted prompt into a set of short commands, including synthesis, editing and erasing operations, using a Large Language Model (LLM) and progressively generate the image.

To overcome these hurdles, we present the Stimulus, Response, and Fusion (SRF) framework, assimilating a stimulus-response generation mechanism along with a latent fusion module into the diffusion process. Our methodology involves employing a fine-tuned GPT model to deconstruct complex texts into structured prompts, including synthesis, editing, and erasing operations governed by a unified SRF framework. Our progressive process begins with a real image or synthesized background, accompanied by the text prompt, and applies the SRF method in a step-by-step approach. Unlike previous strategies that aggressively manipulate the cross-attention map (Wu et al. 2023; Ma et al. 2023), our operation guides the attention map via a soft direction, avoiding brusque modifications that may lead to discordant synthesis. Additionally, when addressing relationships like “wearing” and “playing with”, we begin by parsing the relative positions of the objects with the help of GPT, after which we incorporate the relational description and relative positions into the diffusion model to enable object interactions.

In summary, we unveil a novel, progressive text-to-image diffusion framework that leverages the capabilities of a Language Model (LLM) to simplify language description, offering a unified solution for handling synthesis and editing patterns concurrently. This represents an advancement in text-to-image generation and provides a new platform for future research.

Related Work

Our method is closely related to image manipulation and cross-attention control within diffusion models.

Image manipulating refers to the process of digitally manipulating images to modify or enhance their visual appearance. Various techniques can be employed to achieve this end, such as the use of spatial masks or natural language descriptions to guide the editing process towards specific goals. One promising line of inquiry involves the application of generative adversarial networks (GANs) for image domain transfer (Isola et al. 2017; Sangkloy et al. 2017; Zhu et al. 2017; Choi et al. 2018; Wang et al. 2018; Huang et al. 2018; Park et al. 2019; Liu, Breuel, and Kautz 2017; Baek et al. 2021) or the manipulation of latent space (Zhu et al. 2016; Huh et al. 2020; Richardson et al. 2021; Zhu et al. 2020; Wulff and Torralba 2020; Bau et al. 2021). Recently, diffusion models have emerged as the mainstream. GLIDE (Nichol et al. 2021), Blended diffusion (Avrahami, Fried, and Lischinski 2022) and SmartBrush (Xie et al. 2022) replace masked image regions with predefined objects while preserving the inherent image structure. Additionally, techniques such as prompt-to-prompt (Hertz et al. 2022) and instructpix2pix (Brooks, Holynski, and Efros 2023) enable the modification of image-level objects through text alterations. Contrasting previous methods that solely cater to either synthesis or editing, we construct a unified framework that accommodates both.

Objects and positional relationships are manifested within

Text Decomposition

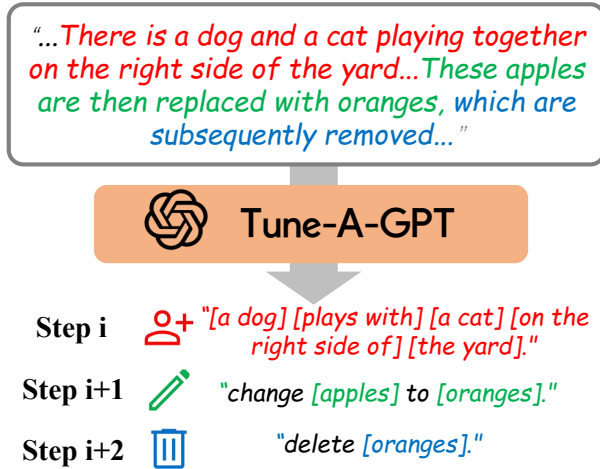


Figure 2: We employ a fine-tuned GPT model to deconstruct a comprehensive text into structured prompts, each classified under synthesis, editing, and erasing operations.

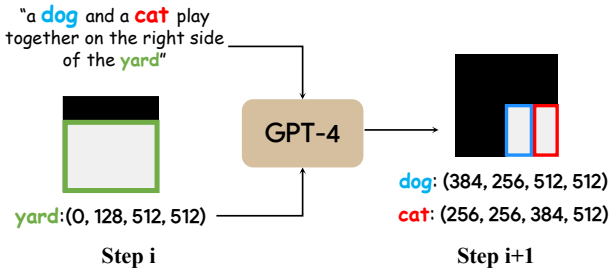


Figure 3: For the synthesis operation, we generate the layout indicated in the prompt from a frozen GPT-4 model, which subsequently yields the new bounding box coordinates for object insertion.

the cross attention map of the diffusion model. Inspired by this observation (Feng et al. 2022), techniques have been devised to manipulate the cross attention map for image synthesis or editing. Prompt-to-Prompt approach (Hertz et al. 2022) aims at regulating spatial arrangement and geometry through the manipulation of attention maps derived from textual prompts. Structured Diffusion (Feng et al. 2022) utilizes a text parsing mechanism to isolate numerous noun phrases, enhancing the corresponding attention space channels. The Attend-and-Excite approach (Chefer et al. 2023) amplifies attention activations linked to the most marginalized subject tokens. Directed Diffusion (Ma et al. 2023) proposes an attention refinement strategy through the utilization of a weak and strong activation approach. The main difference between our layout generation and the layout prediction approaches is that our method enables precise increment generation and intermediate modifications, i.e., we gradually change the layout instead of generating one layout at once. As for background fusion, we use a soft mask to ensure the object’s integrity.

Method

Problem Formulation

we elaborate upon our innovative progressive text-to-image framework. Given a multifaceted text description \mathcal{P} and a real or generated background \mathcal{I} , our primary goal is to synthesize an image that meticulously adheres to the modifications delineated by \mathcal{P} in alignment with \mathcal{I} . The principal challenge emerges from the necessity to decode the intricacy of \mathcal{P} , manifesting across three complex dimensions:

- The presence of multiple entities and attributes escalates the complexity of the scene, imposing stringent demands on the model to generate representations that are not only accurate but also internally coherent and contextually aligned.
- The integration of diverse positional and relational descriptions calls for the model to exhibit an advanced level of understanding and to employ sophisticated techniques to ascertain precise spatial configuration, reflecting both explicit commands and implied semantic relations.
- The concurrent introduction of synthesis, editing, and erasing operations introduces additional layers of complexity to the task. Managing these intricate operations within a unified model presents a formidable challenge, requiring a robust and carefully designed approach to ensure seamless integration and execution.

We address these challenges through a unified progressive text-to-image framework that: (1) employs a fine-tuned GPT model to distill complex texts into short prompts, categorizing each as synthesis, editing, or erasing mode, and accordingly generating the object mask; (2) sequentially processes these prompts within the same framework, utilizing attention-guided generation to capture position-aware features with soft latent direction, and subsequently integrates them with the previous stage’s outcomes in a subtle manner. This approach synthesizes the intricacies of text-to-image transformation into a coherent, positionally aware procedure.

Text Decomposition

\mathcal{P} may involve multiple objects and relations, we decompose \mathcal{P} into a set of short prompts, which produces an image accurately representing \mathcal{P} when executed sequentially. As illustrated in fig. 2, we fine-tune a GPT with OpenAI API (OpenAI 2023) to decompose \mathcal{P} into multiple structured prompts, denoted as $\{\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_n\}$. Each \mathcal{P}_i falls into one of the three distinct modes: **Synthesis mode**: “[object 1] [relation] [object 2] [position] [object 3]”, **Editing mode**: “change [object 1] to [object 2]”, and **Erasing mode**: “delete [object]”. In pursuit of this aim, we start by collecting full texts using ChatGPT (Brown et al. 2020) and then manually deconstruct them into atomic prompts. Each prompt has a minimal number of relations and is labeled with synthesis/editing/erasing mode. Using these prompts and their corresponding modes for model supervision, we fine-tune the GPT model to enhance its decomposition and generalization ability.

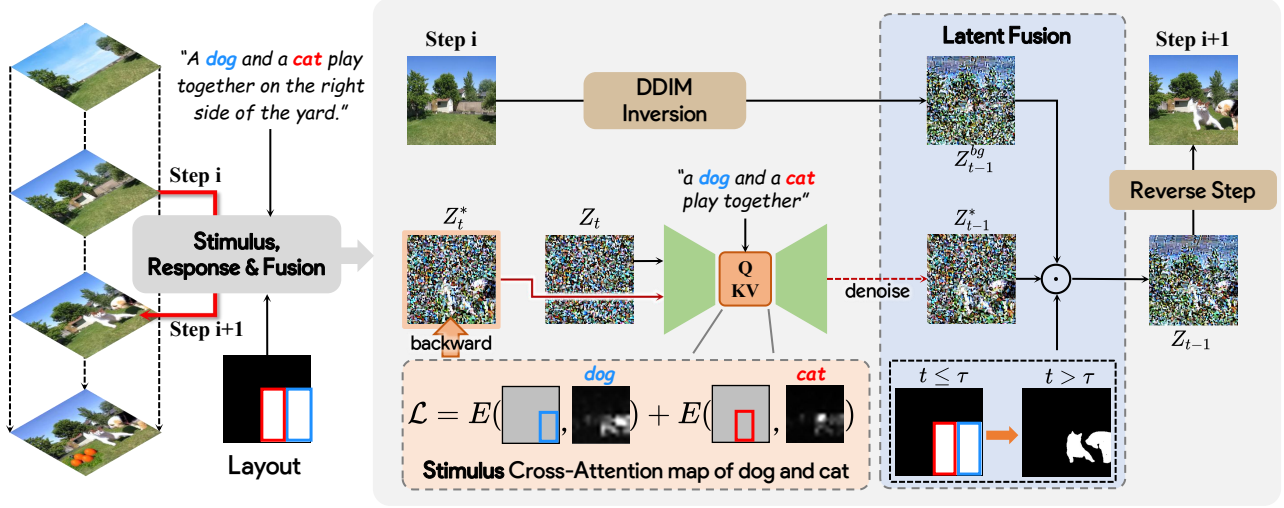


Figure 4: Overview of our unified framework emphasizes progressive synthesis, editing, and erasing. In each progressive step, A random latent z_t is directed through the cross-attention map in inverse diffusion. Specifically, we design a soft stimulus loss that evaluates the positional difference between entity attention and the target mask region, leading to a gradient for updating the latent z_{t-1}^* as a latent response. Subsequently, another forward diffusion pass is applied to denoise z_t^* , yielding deriving z_{t-1}^* . In the latent fusion phase, we transform the previous i -th image into a latent code z_{t-1}^{bg} using DDIM inversion. The blending of z_{t-1}^* with z_{t-1}^{bg} incorporates a dynamic evolving mask, which starts with a layout box and gradually shifts to cross-attention. Finally, z_{t-1}^* undergoes multiple diffusion reverse steps and results in the $(i+1)$ -th image.

Operational Layouts. For the synthesis operation, as shown in fig. 3, we feed both the prompt and a reference bounding box into a frozen GPT-4 API. This procedure produces bounding boxes for the target entity that will be used in the subsequent phase. We exploit GPT-4’s ability to extract information from positional and relational text descriptors. For example, the phrase “cat and dog play together” indicates a close spatial relationship between the “cat” and “dog”. Meanwhile, “on the right side” suggests that both animals are positioned to the right of the “yard”. For the editing and erasing operations, we employ Diffusion Inversion (Mokady et al. 2023) to obtain the cross-attention map of the target object, which serves as the layout mask. For example, when changing “apples” to “oranges”, we draw upon the attention corresponding to “apples”. On the other hand, to “delete the oranges”, we focus on the attention related to “oranges”. Notably, this approach avoids the need to retrain the diffusion model and is proficient in managing open vocabularies. we denote generated layout mask as \mathcal{M} for all operations in following sections for convention.

In the following section, we provide a complete introduction to the synthesis operation. At last, we exhibit that the editing and erasing operations only differ from the synthesis operation in parameter settings.

Stimulus & Response

With the synthesis prompt \mathcal{P}_i to be executed and its mask configuration \mathcal{M}_i . The goal of Latent Stimulus & Response is to enhance the positional feature representation on \mathcal{M} . As illustrated in fig. 4, this is achieved by guided cross-attention generation. Differing from the approaches (Ma et al. 2023;

Wu et al. 2023), which manipulate attention through numerical replacement, we modulate the attention within mask regions associated with the entity in \mathcal{P}_i via a soft manner. Rather than directly altering the attention, we introduce a stimulus to ensure that the object attention converges to the desired scores. Specifically, we formulate a stimulus loss function between the object mask \mathcal{M} and the corresponding attention A as:

$$\mathcal{L}_s = \sum_{i=1}^n (\text{softmax}(A_t^i) - \delta \cdot \mathcal{M}^i) \quad (1)$$

where A_t^i signifies the cross-attention map of the i -th object at the t -th timestep. \mathcal{M}^i denotes the mask of the i -th object. δ represents the stimulus weights. The intent of stimulus attention leans towards a spatial-wise generation process. This is achieved by backpropagating the gradient of the stimulus loss function, as defined in Eq. 1, to update the latent code. This process serves as a latent response to the stimulated attention, which can be formally expressed as:

$$z_t^* \leftarrow z_t - \alpha_t \cdot \nabla_{z_t} \mathcal{L}_s \quad (2)$$

In the above equation, z_t^* represents the updated latent code and α_t denotes the learning rate. Finally, we execute another forward pass of the stable diffusion model using the updated latent code z_t^* to compute z_{t-1}^* for the subsequent denoising step. Based on eq. (1) and eq. (2), we observe consistent spatial behavior in both the cross-attention and latent spaces. For a more detailed analysis, we refer to fig. 5 and find this property contributes to producing faithful and position-aware image representations.



Figure 5: Visual results generated by Stable Diffusion and Stimulus & Response. Stable Diffusion shows noticeable problems in positional generation (top), semantic and attribute coupling (middle), and object omission (bottom), while ours delivers precise outcomes.

Latent Fusion

Recalling that z_{t-1}^* denotes the latent feature of the target object, our next task is to integrate them seamlessly with the image from the preceding stage. For this purpose, we first convert the previous image into latent code by DDIM inversion, denoted as z_t^{bg} . Then for timestep t , we take a latent fusion strategy (Avrahami, Lischinski, and Fried 2022) between z_t^{bg} and z_t^* , which is formulated as:

$$z_{t-1} = \widehat{\mathcal{M}} \cdot z_{t-1}^* + (1 - \widehat{\mathcal{M}}) \cdot z_{t-1}^{bg} \quad (3)$$

where $\widehat{\mathcal{M}}$ acts as a latent mask to blend the features of target objects with the background. In the synthesis operation, employing a uniform mask across all steps can be too restrictive, potentially destroying the object’s semantic continuity. To mitigate this, we introduce a more soft mask, ensuring both object integrity and spatial consistency. Specifically, during the initial steps of diffusion denoising, we use layout mask \mathcal{M} to provide spatial guidance. Later, we shift to an attention mask \mathcal{M}_{attn} , generated by averaging and setting a threshold on the cross-attention map, to maintain object cohesion. This process is denoted as:

$$\widehat{\mathcal{M}}(\mathcal{M}_{attn}, \mathcal{M}, t) = \begin{cases} \mathcal{M} & \text{if } t \leq \tau \\ \mathcal{M}_{attn} & \text{if } t > \tau \end{cases} \quad (4)$$

Here, τ serves as a tuning parameter balancing object integrity with spatial coherence. The above response and fusion process is repeated for a subset of the diffusion timesteps, and the final output serves as the image for the next round generation.

Editing and Erasing Specifications. Our editing and erasing operation differs in parameter setting: we set $\widehat{\mathcal{M}}$ in eq. (1) as editing/erasing reference attention. we set $\widehat{\mathcal{M}}$ in eq. (3) as the editing/erasing mask in all diffusion steps for detailed, shape-specific modifications.

Experiment

Baselines and Evaluation. Our experimental comparison primarily concentrates on Single-Stage Generation and Progressive Generation baselines. (1) We refer to **Single-Stage Generation** methods as those that directly generate images from input text in a single step. Current methods include Stable Diffusion (Rombach et al. 2022), Attend-and-excite (Chefer et al. 2023), and Structured Diffusion (Feng et al. 2022). We compare these methods to analyze the efficacy of our progressive synthesis operation. We employ GPT to construct 500 text prompts that contain diverse objects and relationship types. For evaluation, we follow (Wu et al. 2023) to compute **Object Recall**, which quantifies the percentage of objects successfully synthesized. Moreover, we measure **Relation Accuracy** as the percentage of spatial or relational text descriptions that are correctly identified, based on 8 human evaluations. (2) We define **Progressive Generation** as a multi-turn synthesis and editing process that builds on images from preceding rounds. Our comparison encompasses our comprehensive progressive framework against other progressive methods, which includes Instruct-based Diffusion models (Brooks, Holynski, and Efros 2023) and mask-based diffusion models (Rombach et al. 2022; Avrahami, Fried, and Lischinski 2022). To maintain a balanced comparison, we source the same input images from SUN (Xiao et al. 2016) and text descriptions via the GPT API (OpenAI 2023). Specifically, we collate five scenarios totaling 25 images from SUN, a dataset that showcases real-world landscapes. Each image is paired with the text description, which ensures: 1. Integration of synthesis, editing, and easing paradigms; 2. Incorporation of a diverse assortment of synthesized objects; 3. Representation of spatial relations (e.g., top, bottom, left, right) and interactional relations (e.g., “playing with”, “wearing”). For evaluation, we utilize Amazon Mechanical Turk (AMT) to assess image fidelity. Each image is evaluated based on the fidelity of the generated objects, their relationships, the execution of editing instructions, and the alignment of erasures with the text descriptions. Images are rated on a fidelity scale from 0 to 2, where 0 represents the lowest quality and 2 signifies the highest. With two evaluators assessing each generated image, the cumulative score for each aspect can reach a maximum of 100.

Implementation Details. Our framework builds upon Stable Diffusion (SD) V-1.4. During the Stimulus & Response stage, we assign a weight of δ equals 0.8 in eq. (1), and set t equals 25 and α_t equals 40 in eq. (2). We implement the stimulus procedure over the 16×16 attention units and integrate the Iterative Latent Refinement design (Chefer et al. 2023). In the latent fusion stage, the parameter τ is set to a value of 40.

Qualitative and Quantitative Results

Qualitative and Quantitative Comparisons with Single-Generation Baselines. fig. 6 reveals that traditional baseline methods often struggle with object omissions and maintaining spatial and interactional relations. In contrast, our progressive generation process offers enhanced image fidelity and controllability. Additionally, we maintain finer details in the generated images, such as the shadows of the

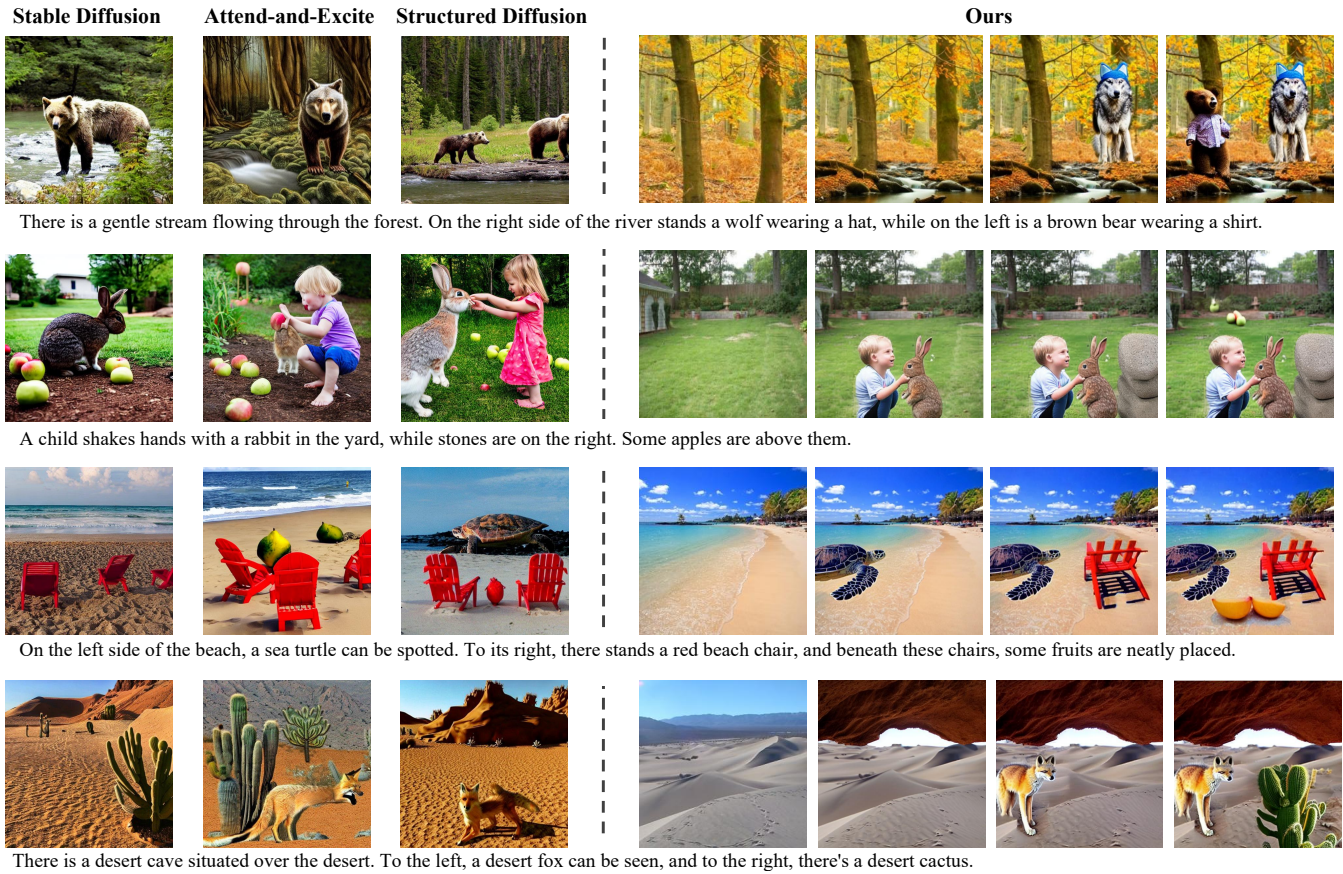


Figure 6: Qualitative comparison with Single-Stage baselines. Common errors in the baselines include missing objects and mismatched relations. Our method demonstrates the progressive generation process.

“beach chair”. Result in table 1 indicates that our method outperforms the baselines in both object recall and relation accuracy.

Qualitative and Quantitative Comparisons with Progressive Generation Baselines. In fig. 8, baseline methods often fail to synthesize full objects and may not represent relationships as described in the provided text. Moreover, during editing and erasing operations, these methods tend to produce outputs with compromised quality, showcasing unnatural characteristics. It’s worth noting that any missteps or inaccuracies in the initial stages, such as those seen in InstructPix2Pix, can cascade into subsequent stages, exacerbating the degradation of results. In contrast, our proposed method consistently yields superior results through every phase. The results in table 2 further cement our method’s dominant performance in synthesis, editing, and erasing operations, as underscored by the impressive rating scores.

Ablation Study

Ablation study of method components is shown in table 3. Without latent fusion, we lose continuity from prior generation stages, leading to inconsistencies in object synthesis and placement. On the other hand, omitting the Stimulus & Response process results in a lack of positional

Method	Object Recall ↑	Relation Accuracy ↑
Stable Diffusion	40.7	19.8
Structured Diffusion	43.5	21.6
Attend-and-excite	50.3	23.4
Ours	64.4	50.8

Table 1: Quantitative comparison with Single-Stage Generation baselines.

awareness, making the synthesis less precise. Both omissions manifest as significant drops in relation and entity accuracies, emphasizing the synergistic importance of these components in our approach.

The analysis of Stimulus & Response in the editing operation is highlighted in fig. 7. Compared to Stable Diffusion, Stimulus & Response not only enhances object completeness and fidelity but also demonstrates a broader diversity in editing capabilities. The loss curve indicates that Stimulus & Response aligns more closely with the reference cross-attention, emphasizing its adeptness in preserving the original structure.

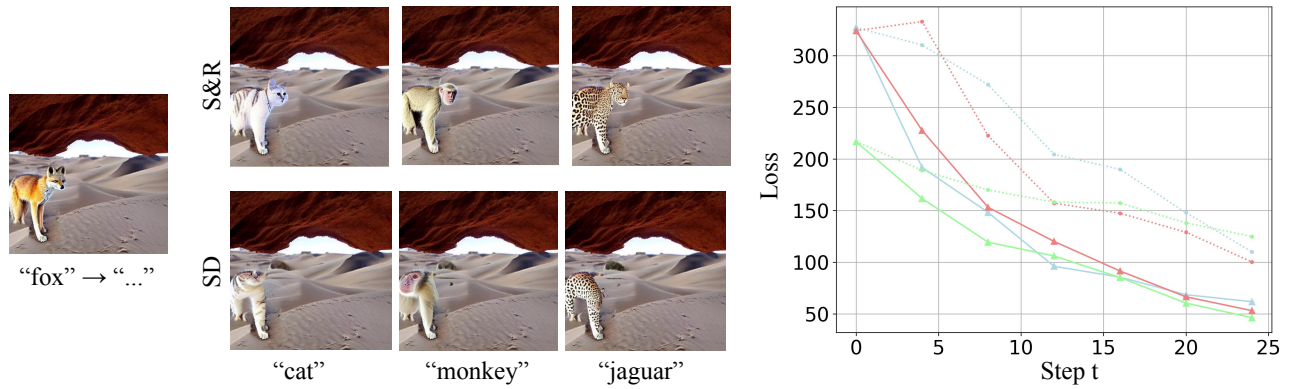


Figure 7: The analysis of Stimulus & Response in the editing operation. The left side shows a visual comparison between SD (Stable Diffusion) and S&R (Stimulus & Response). The right side presents the convergence curve of cross-attention loss during diffusion sampling steps. The loss is computed as the difference between reference attention and model-generated attention. In the right figure, red, blue, and green colors represent the objects “jaguar”, “cat”, and “monkey” respectively. Solid lines indicate SD loss, while dashed lines represent S&D loss.

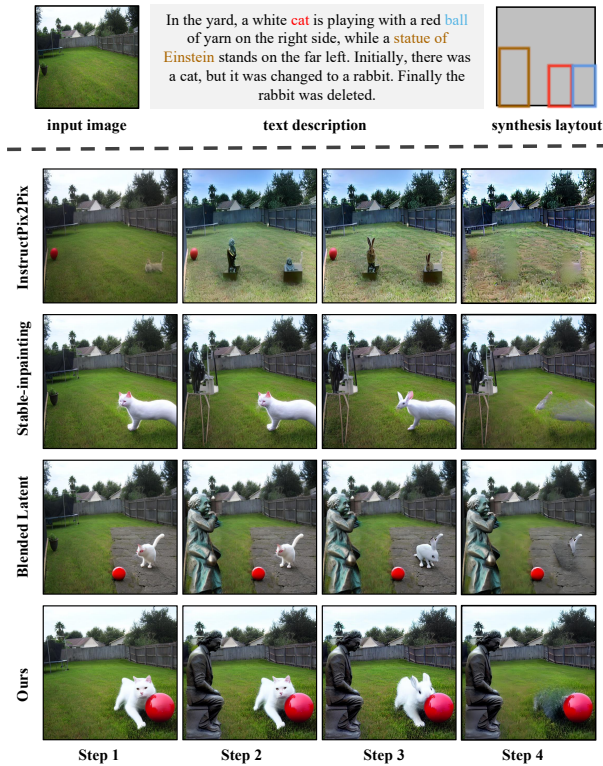


Figure 8: Qualitative comparison with Progressive Generation baselines. The first two phases illustrate object synthesis operation, where target objects are color-coded in both the text and layout. Subsequent phases depict object editing and erasing processes, wherein a cat is first transformed into a rabbit and then the rabbit is removed.

Method	Synthesis		Editing	Erasing
	Object	Relation		
InstructPix2Pix	19	24	32	29
Stable-inpainting	64	54	65	45
Blended Latent	67	52	67	46
Ours	74	60	72	50

Table 2: Quantitative comparison of our method against Progressive Generation baselines, using rating scores.

Method Variant	Object Recall ↑	Relation Accuracy ↑
w/o LF	38.8	21.8
w/o S&R	58.3	45.2
Ours	64.4	50.8

Table 3: Ablation study. LF and S&R represent Latent Fusion and Stimulus & Response respectively.

Conclusion

In this study, we addressed the prevailing challenges in the rapidly advancing field of text-to-image generation, particularly the synthesis and manipulation of multiple entities under specific constraints. Our innovative progressive synthesis and editing methodology ensures precise spatial and relational representations. Recognizing the limitations of existing diffusion models with increasing entities, we integrated the capabilities of a Large Language Model (LLM) to dissect complex text into structured directives. Our Stimulus, Response, and Fusion (SRF) framework, which enables seamless entity manipulation, represents a major stride in object synthesis from intricate text inputs.

One major limitation of our approach is that not all text can be decomposed into a sequence of short prompts. For instance, our approach finds it challenging to sequentially parse text such as “a horse under a car and between a cat and a dog.” We plan to gather more training data and labels of this nature to improve the parsing capabilities of GPT.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (NSFC No. 62272184). The computation is completed in the HPC Platform of Huazhong University of Science and Technology.

References

- Avrahami, O.; Fried, O.; and Lischinski, D. 2022. Blended latent diffusion. *arXiv preprint arXiv:2206.02779*.
- Avrahami, O.; Lischinski, D.; and Fried, O. 2022. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18208–18218.
- Baek, K.; Choi, Y.; Uh, Y.; Yoo, J.; and Shim, H. 2021. Rethinking the truly unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14154–14163.
- Bau, D.; Andonian, A.; Cui, A.; Park, Y.; Jahanian, A.; Oliva, A.; and Torralba, A. 2021. Paint by word. *arXiv preprint arXiv:2103.10951*.
- Brooks, T.; Holynski, A.; and Efros, A. A. 2023. Instruct-pix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18392–18402.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chefer, H.; Alaluf, Y.; Vinker, Y.; Wolf, L.; and Cohen-Or, D. 2023. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *arXiv preprint arXiv:2301.13826*.
- Choi, Y.; Choi, M.; Kim, M.; Ha, J.-W.; Kim, S.; and Choo, J. 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8789–8797.
- Feng, W.; He, X.; Fu, T.-J.; Jampani, V.; Akula, A.; Narayana, P.; Basu, S.; Wang, X. E.; and Wang, W. Y. 2022. Training-Free Structured Diffusion Guidance for Compositional Text-to-Image Synthesis. *arXiv preprint arXiv:2212.05032*.
- Goel, V.; Peruzzo, E.; Jiang, Y.; Xu, D.; Sebe, N.; Darrell, T.; Wang, Z.; and Shi, H. 2023. Pair-diffusion: Object-level image editing with structure-and-appearance paired diffusion models. *arXiv preprint arXiv:2303.17546*.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.
- Huang, X.; Liu, M.-Y.; Belongie, S.; and Kautz, J. 2018. Multimodal unsupervised image-to-image translation. In *Proceedings of the European conference on computer vision (ECCV)*, 172–189.
- Huh, M.; Zhang, R.; Zhu, J.-Y.; Paris, S.; and Hertzmann, A. 2020. Transforming and projecting images into class-conditional generative networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, 17–34. Springer.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125–1134.
- Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2023. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6007–6017.
- Liu, M.-Y.; Breuel, T.; and Kautz, J. 2017. Unsupervised image-to-image translation networks. *Advances in neural information processing systems*, 30.
- Ma, W.-D. K.; Lewis, J.; Kleijn, W. B.; and Leung, T. 2023. Directed Diffusion: Direct Control of Object Placement through Attention Guidance. *arXiv preprint arXiv:2302.13153*.
- Mokady, R.; Hertz, A.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2023. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6038–6047.
- Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.
- OpenAI. 2023. GPT-4 Technical Report. *arXiv:2303.08774*.
- Park, T.; Liu, M.-Y.; Wang, T.-C.; and Zhu, J.-Y. 2019. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2337–2346.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, 8821–8831. PMLR.
- Richardson, E.; Alaluf, Y.; Patashnik, O.; Nitzan, Y.; Azar, Y.; Shapiro, S.; and Cohen-Or, D. 2021. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2287–2296.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E.; Ghasemipour, S. K. S.; Ayan, B. K.; Mahdavi, S. S.; Lopes, R. G.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*.

- Sangkloy, P.; Lu, J.; Fang, C.; Yu, F.; and Hays, J. 2017. Scribbler: Controlling deep image synthesis with sketch and color. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5400–5409.
- Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402*.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2556–2565.
- Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8798–8807.
- Wang, Z.; Liu, W.; He, Q.; Wu, X.; and Yi, Z. 2022. Clipgen: Language-free training of a text-to-image generator with clip. *arXiv preprint arXiv:2203.00386*.
- Wu, Q.; Liu, Y.; Zhao, H.; Bui, T.; Lin, Z.; Zhang, Y.; and Chang, S. 2023. Harnessing the spatial-temporal attention of diffusion models for high-fidelity text-to-image synthesis. *arXiv preprint arXiv:2304.03869*.
- Wulff, J.; and Torralba, A. 2020. Improving inversion and generation diversity in stylegan using a gaussianized latent space. *arXiv preprint arXiv:2009.06529*.
- Xiao, J.; Ehinger, K. A.; Hays, J.; Torralba, A.; and Oliva, A. 2016. Sun database: Exploring a large collection of scene categories. *International Journal of Computer Vision*, 119: 3–22.
- Xie, S.; Zhang, Z.; Lin, Z.; Hinz, T.; and Zhang, K. 2022. SmartBrush: Text and Shape Guided Object Inpainting with Diffusion Model. *arXiv preprint arXiv:2212.05034*.
- Yang, B.; Gu, S.; Zhang, B.; Zhang, T.; Chen, X.; Sun, X.; Chen, D.; and Wen, F. 2023. Paint by example: Exemplar-based image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18381–18391.
- Zhu, J.; Shen, Y.; Zhao, D.; and Zhou, B. 2020. In-domain gan inversion for real image editing. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16*, 592–608. Springer.
- Zhu, J.-Y.; Krähenbühl, P.; Shechtman, E.; and Efros, A. A. 2016. Generative visual manipulation on the natural image manifold. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, 597–613. Springer.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 2223–2232.