

CompoNeRF: Text-guided Multi-object Compositional NeRF with Editable 3D Scene Layout

Haotian Bai*
HKUST(GZ)

Yuanhuiyi Lyu†
HKUST(GZ)

Lutao Jiang‡
HKUST(GZ)

Sijia Li
OPPO

Haonan Lu
OPPO

Xiaodong Lin
OPPO

Lin Wang§
HKUST(GZ)
HKUST

Project homepage: <https://vllis2022.github.io/componerf>

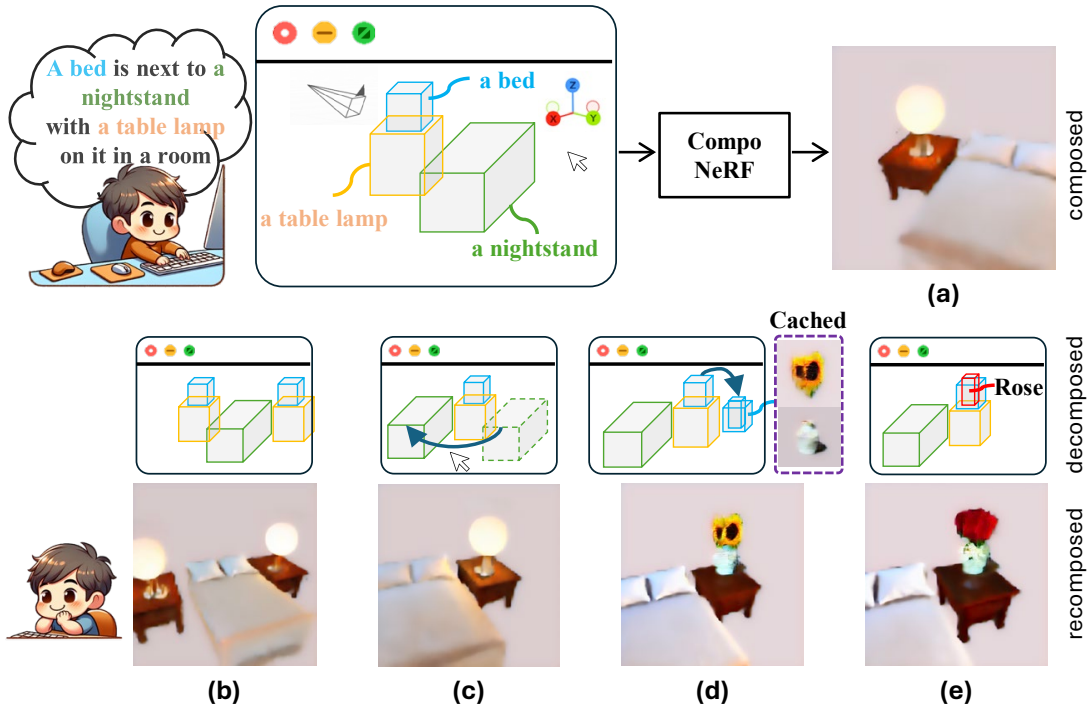


Figure 1: We introduce CompoNeRF, a novel framework that synthesizes coherent multi-object scenes by integrating textual descriptions and layouts. CompoNeRF allows for individual NeRFs, each denoted by a unique prompt color, to be composed, decomposed, and recomposed with ease, streamlining the construction of complex scenes from cached models after decomposition. (a) displays the composed results. (b), (c), (d), (e) are recomposition results after manipulation demos shown above, including duplication, transformation, loading decomposed NeRFs, and semantic editing conducted separately.

ABSTRACT

Text-to-3D form plays a crucial role in creating editable 3D scenes for AR/VR. Recent advances have shown promise in merging neural radiance fields (NeRFs) with pre-trained diffusion models for text-to-3D object generation. However, one enduring challenge is their inadequate capability to accurately parse and regenerate consistent **multi-object** environments. Specifically, these models encounter difficulties in accurately representing quantity and style prompted by multi-object texts, often resulting in a collapse of the rendering fidelity that fails to match the semantic intricacies.

Moreover, amalgamating these elements into a coherent 3D scene is a substantial challenge, stemming from generic distribution inherent in diffusion models. To tackle the issue of ‘guidance collapse’ and further enhance scene consistency, we propose a novel framework, dubbed **CompoNeRF**, by integrating an editable 3D scene layout with object-specific and scene-wide guidance mechanisms. It initiates by interpreting a complex text into the layout populated with multiple NeRFs, each paired with a corresponding subtext prompt for precise object depiction. Next, a tailored composition module seamlessly blends these NeRFs, promoting consistency, while the dual-level text guidance reduces ambiguity and boosts accuracy. Noticeably, our composition design permits decomposition. This enables flexible scene editing and recomposition into new scenes based on the edited layout or text prompts. Utilizing the open-source Stable Diffusion model, CompoNeRF generates multi-object scenes with high fidelity. Remarkably, our framework achieves up to a **54%** improvement by the multi-view CLIP score metric. Our user study indicates that our method has significantly improved semantic accuracy, multi-view consistency, and

*e-mail: haotianwhite@outlook.com

†e-mail: yuanhuiyilv@hkust-gz.edu.cn

‡e-mail: jianglutao98@gmail.com

§Corresponding author, email: linwang@ust.hk

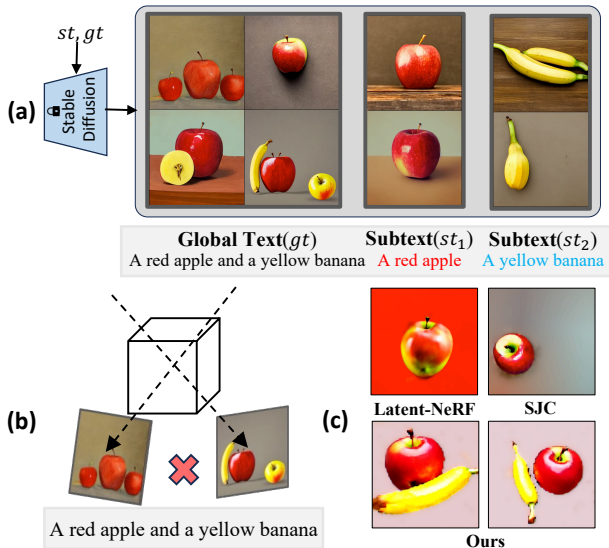


Figure 2: **The guidance collapse issue.** (a) Generation of the multi-object scene involves utilizing the frozen Stable Diffusion. (b) Instances of guidance collapse are observed when using the global text directly. (c) Comparison of rendering results .

individual recognizability for multi-object scene generation.

Index Terms: Multimodal capturing and reconstruction, Text-to-3D, Neural Radiance Field.

1 INTRODUCTION

Consider the last time you wanted to relive an imaginary virtual scene through VR/AR devices. As illustrated in Fig. 1, picture a tranquil bedroom where the day comes to a serene close. Envision a bed adorned with soft pillows, a bedside table, and a lamp casting a soothing light. Our mental imagery often includes a variety of objects and their interplay. How do we then create these scenes from descriptions, such as text prompts, and convert them into cohesive, editable environments within virtual reality?

Recent advances in text-to-image generation have been driven by the integration of vision-language pre-trained models [41, 20] with diffusion processes [12, 36, 44], leading to impressive outcomes. Pioneering text-to-3D approaches [14, 45, 13, 11, 33, 18, 59] have built upon these successes, employing these robust vision-language models to enrich 3D generative models with the structured understanding provided by Neural Radiance Fields (NeRFs) [30, 3, 34]. This synergy [39, 23, 29, 53] facilitates the creation of 3D models which, when rendered from different views, cohere with the learned text-to-image diffusion model distribution, opening new avenues for the direct synthesis of 3D content from textual descriptions.

Textual descriptions can be vague and open to interpretation. Converting these prompts into visual images, particularly for complex scenes with numerous objects, is not a simple task. Specifically, diffusion models like Stable Diffusion [44] have undergone extensive training using large-scale text-image datasets [46]. Despite this, they frequently encounter difficulties when dealing with multiple object texts, particularly when those objects have a limited presence in the training data. As a result, this process often results in images that either leave out certain objects or depict them incorrectly. Fig. 2(a) demonstrates how even in a simple scenario involving just two objects, Stable Diffusion may not consistently preserve the integrity of the scene. In detail, it might overlook an apple or banana or render them with inaccurate colors. Specifically, it may neglect to include an apple or banana, or it might portray them in

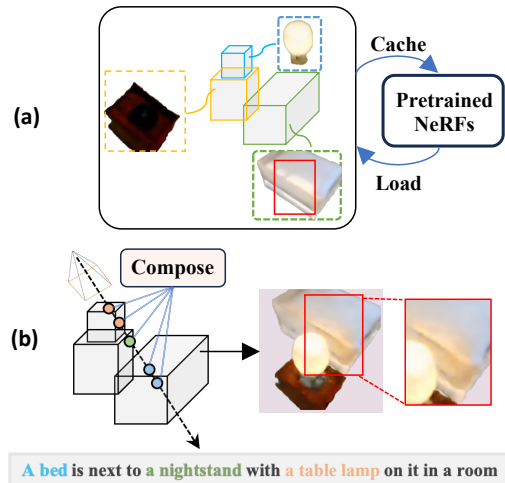


Figure 3: (a) CompoNeRF supports caching and loading to facilitate NeRF composition. (b) The composition module composites multiple NeRFs for coherent scenes. Its enhanced effect is accentuated by the red boxes, showcasing superior scene coherency.

incorrect colors. This problem, referred to as 'guidance collapse' as illustrated in (b), is especially problematic when rendering scenes involving multiple objects based on text prompts. As shown in (c), even advanced models like Latent-NeRF [29] and SJC [53] struggle to accurately generate the configurations described in texts involving multiple objects. This limitation significantly hampers their capability to construct 3D scenes derived from descriptive prompts.

This raises a compelling question: *Can we design a model using diffusion models armed with a generic distribution, not only identify and recreate individual elements in multi-object texts but also amalgamate them into a coherent 3D scene?*

In this paper, we present **CompoNeRF**, a compositional NeRF framework that interprets multi-object text prompts as editable 3D scene layouts with granular text prompts. The procedure, illustrated in Fig. 3(a), begins by identifying individual objects from the textual description and positioning them within customizable 3D bounding boxes. Each box is supported by a distinct NeRF and a subtext label. Thus, as shown at Fig. 1, CompoNeRF is designed to accommodate alterations, allowing for manipulations in the layout—like moving, scaling, or removal, as well as loading decomposed nodes and direct text editing. As depicted in Fig. 3(b), our composition module guarantees that the overall scene is not merely a static collection but an orchestrated assembly. For example, after composition, the lamp illuminates the bed, creating a lively scene with such interaction among objects.

CompoNeRF distinguishes itself with three core capabilities: it *composes* multi-object scenes from textual prompts, *decomposes* by archiving each NeRF for subsequent utilization, and *recomposes* by employing this curated content gallery to rapidly generate elaborate 3D scenes, thereby streamlining the 3D content creation workflow. Next, we employ the averaged CLIP score [52] on rendering views of 3D content against global text prompts, we quantitatively measure the alignment of our generated scenes with their textual prompts. For a more comprehensive evaluation, we carry out a user study to measure **1)** composition correctness regarding both semantic and multi-view consistency; **2)** generation quality in terms of users' overall preferences **3)** recognizability for each component within the scenes. The studies demonstrate our effectiveness in producing detailed and coherent 3D scenes that accurately reflect the given text descriptions.

To encapsulate, our paper makes three key contributions: **(I)** We address the 'guidance collapse' problem in creating multi-object

3D scenes. Our innovative use of editable 3D layouts coupled with multiple localized NeRFs allows for precise direction over individual object representations. Moreover, these localized NeRF models are designed to be storable and reusable, enhancing efficiency in scene composition. **(II)** We introduce a composition module enables fine-tuning of the rendering process and text-based guidance, ensuring both the distinctiveness of individual objects and the holistic integration within the scene. **(III)** We conduct extensive evaluations of CompoNeRF’s performance in multi-object scene generation, employing both qualitative and quantitative assessment for the multi-object text-to-3D task. The rigorous testing confirms that our CompoNeRF outperforms existing models in generating multi-object scenes that closely align with textual prompts.

2 RELATED WORKS

Neural Rendering for 3D Modeling. Recent endeavors have been made to integrate 3D content with NeRF for real-time viewing and interaction in AR/VR. For instance, works such as [43, 19] have provided real-time experience using VR/AR headsets with NeRF rendering, offering immersive virtual experiences. Recent efforts [8] also optimize computational resources by capturing the user’s gaze for enhanced viewing and interaction. The evolution of NeRF has elevated the capabilities of neural rendering. NeRF-based models [31, 24, 34, 25, 3, 51, 2] have redefined volume rendering [16] through the use of coordinate-based MLPs that infer color and density from spatial and directional inputs. Their capacity to produce photo-realistic views has cemented differential volume rendering as a key component in a variety of applications, such as scene relighting [49, 64], dynamic scene reconstruction [10, 40, 58, 50], and scene and avatar editing [26, 61], as well as surface reconstruction [1, 15, 54]. Typically, these approaches rely on a *single* MLP to encode an entire scene, which may introduce ambiguity in differentiating between objects. Our method, in contrast, renders scenes through a composite of multiple NeRFs, each responsible for a distinct part of the scene. This composition takes into consideration the interactions between NeRFs, allowing for the existence of individual objects while maintaining overall coherence.

Text-guided 3D Generative Models. To facilitate 3D asset creation in VR/AR, the integration of vision-language models like CLIP with 3D generative methods has propelled text-guided 3D generation forward. Models that harness these advancements, such as those by Jain et al.[14] and Wang et al.[52], have excelled in aligning 3D renderings with text descriptors but often fall short in detail, limiting realism. Innovative approaches like DreamFusion [39], Magic3D [23], and Latent-NeRF [29] have sought to enhance this through text-to-image diffusion models and score distillation sampling in the latent space, with SJC [53] and Dream-Booth3D [42] further refining the process to address distribution mismatches and enable image-based 3D generation, respectively. Points-to-3D [62] takes a novel route by utilizing 3D point clouds for guidance, whereas Fantasia3D [6] innovatively disentangles geometry and appearance tasks, it employs the Stable Diffusion model for learning geometry and utilizes the Physically-Based Rendering (PBR) material model [28] for appearance learning. Moreover, several studies [22, 56, 63, 27, 21] are dedicated to enhancing the SDS loss to provide more detailed supervision. Departing from these singular approaches, our CompoNeRF introduces a novel approach for the creation of multi-object 3D scenes. It adopts an *object-compositional* strategy, utilizing an editable 3D scene layout that conceptualizes the scene not as a singular but as a constellation of discrete NeRFs. Each NeRF is associated with its spatial 3D bounding box and a corresponding text prompt, allowing for guidance from both the global texts and their subtexts. This dual-text framework ensures that each object is not only individually delineated but also integrated into the composite scene, thereby enhancing the authenticity of the generated 3D scenes.

Object-Compositional Scene Modeling. The creation of new scenes from individual, object-centric components represents a trend in scene generation, as evidenced by existing research [66, 61, 57, 32, 60, 48]. These efforts typically adopt one of two approaches: semantic-based or 3D layout-based. Semantic-based methods enhance object representations by incorporating additional semantic information, such as segmentation labels [66], instance masks [61, 57], or features extracted using pre-trained vision-language models [32]. On the other hand, 3D layout-based approaches, exemplified by NSG [37] and its successors [48, 9], focus on spatial coordinates, using explicit 3D object placement data to guide object and scene composition. Diverging from conventional techniques, our method innovates by utilizing decomposed, object-specific 3D layouts. This approach enables precise control over scene dynamics, encompassing both object-specific text prompt modifications [39, 23] and spatial manipulation. CompoNeRF’s distinctive feature lies in its capability to recompose scenes by interfacing with decomposed NeRFs, thereby accelerating the creation of new scenes. In contrast to the mesh-based method in Fantasia3D, which requires considerable human effort in mesh modification and graphics engine support for editing, CompoNeRF offers a more streamlined process. Our composition module seamlessly integrates components, requiring minimal adjustments in layout or text prompts, followed by fine-tuning existing offline models to align with the global context during training.

3 METHOD

To resolve the issue of guidance collapse, our principal strategy is to *decompose the scene into reusable components and compose/recompose them into a unified and consistent one*. This enables flexible control over the generated content with direct use of prompts and box layouts. As illustrated in Fig. 1, our proposed CompoNeRF confers several key benefits: 1) **Semantic Coherence:** It reliably creates 3D objects with detailed textures and global consistency, exemplified by authentic light interactions, such as reflections on the bed surface. 2) **Modularity and Reusability:** CompoNeRF functions as an ensemble of independently trained NeRF models. These can be efficiently stored and later retrieved from a cached dataset, enabling their reuse in various cases. 3) **Editability:** Our approach allows for flexible scene modification, such as interchanging the lamp for a vase filled with sunflowers or altering its scale, by simply adjusting the box dimensions for later finetuning. This feature enhances flexibility and creative possibilities.

3.1 Preliminaries

Defining individual object bounding boxes as *local frames* and the overall scene coordinate system as the *global frame*, we build the foundation of NeRF and diffusion processes.

3D Representation in Latent Space. Our methodology capitalizes on the state-of-the-art text-to-image generative model—Stable Diffusion as described by Rombach et al[44]. We build upon the Latent-NeRF framework [29], which computes latent colors for individual objects by considering their sample positions within a localized frame. Specifically, it maps a three-dimensional point in local coordinates $x_l = (x_l, y_l, z_l)$ to a volumetric density σ_l and an associated color C_l , expressed as $(C_l, \sigma_l) = f_{\theta_l}(x_l, y_l, z_l)$. Here, f represents a Multi-Layer Perceptron (MLP) characterized by parameters θ_l . This NeRF-generated color is then assessed in the context of the Stable Diffusion model, using text prompts to guide NeRF toward spatially coherent inference with intricate context. **Volume Rendering with Multiple Objects.** We extend the volume rendering process to accommodate multiple objects by assigning each a local frame, denoted as j , with NeRF parameters $\theta_{l,j}$. Drawing from the foundational NeRF approach [31], in each local frame, we integrate the color C_l and density σ_l for points x_l sampled along a ray r_l , emanates from the camera origin o_l in direction

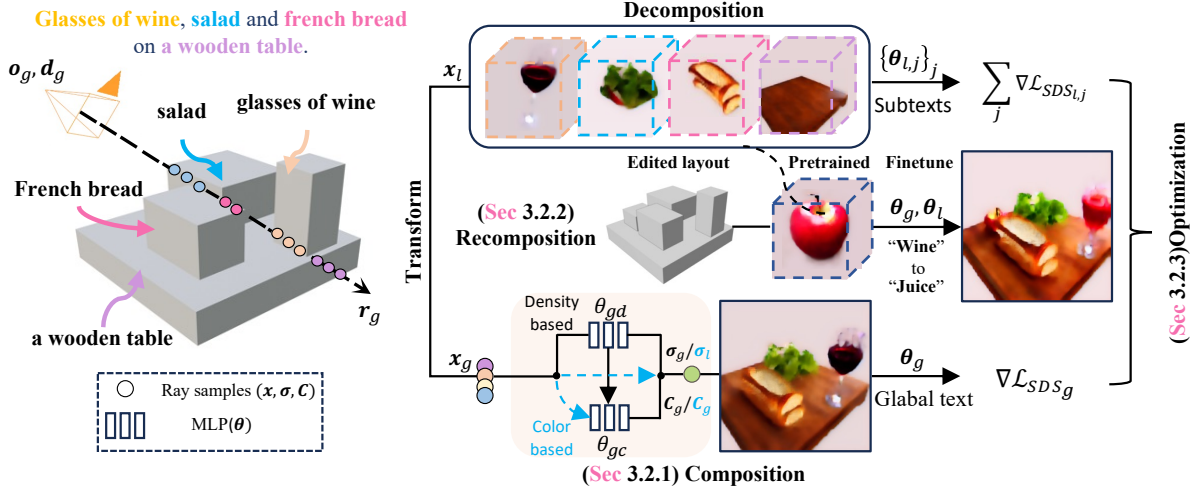


Figure 4: **Framework Overview.** The CompoNeRF model unfolds in three stages: 1) Editing 3D scene, which initiates the process by structuring the scene with 3D boxes and textual prompts; 2) Scene rendering, which encapsulates the composition/recomposition process, facilitating the transformation of NeRFs to a global frame, ensuring cohesive scene construction. Here, we specify design choices between density-based or color-based(without refining density) composition; 3) Joint Optimization, which leverages textual directives to amplify the rendering quality of both global and local views, while also integrating revised text prompts and NeRFs for refined scene depiction.

d_l . This is formalized in the predicted color integration for \hat{C}_l as:

$$\hat{C}_l(r_l) = \sum_{k=1}^N T_{l,k} (1 - \exp(-\sigma_{l,k} \delta_k)) C_{l,k}, \quad (1)$$

where $T_{l,k} = \exp(-\sum_{j=1}^{k-1} \sigma_{l,j} \delta_j)$ represents the transmittance to the k -th of total N sample, calculated exponentially over the cumulative density along r_l , and δ_k is the interval between adjacent samples. To synthesize a coherent scene, we transition from processing individual local frames to a collective global frame. Within this global context, we reconcile object attributes inferred from their individual local NeRFs for refined σ_g, C_g along with $T_{g,k}$. The samples x_g are ordered based on their spatial distances from the origin o_g following the coordinate transformation. We then express the volumetric rendering of a ray r_g integrating m objects within the global frame as follows:

$$\hat{C}_g(r_g) = \sum_{k=1}^{m \times N} T_{g,k} (1 - \exp(-\sigma_{g,k} \delta_k)) C_{g,k}. \quad (2)$$

Score Distillation Sampling. To facilitate the conversion from text descriptions to 3D models, DreamFusion [39] utilizes Score Distillation Sampling (SDS), leveraging the generative capabilities of a diffusion model, denoted as ϕ , to guide the optimization of NeRF parameters, symbolized as θ . Initially, SDS creates a noisy image X_t by infusing a randomly sampled noise ε , which follows a normal distribution $\mathcal{N}(0, I)$, into a NeRF-rendered image X at a given noise level t . The diffusion model ϕ then estimates the noise $\varepsilon_\phi(X_t, t, T)$ from this noisy image, conditioned by the noise level t and an optional text prompt T . The key step in SDS involves calculating the gradient of the loss function, which measures the discrepancy between the estimated noise and the originally added noise:

$$\nabla_{\theta} \mathcal{L}_{\text{SDS}}(X_t, T) = w(t) (\varepsilon_\phi(X_t, t, T) - \varepsilon), \quad (3)$$

where $w(t)$ is a weighting function that adjusts the influence of the gradient based on the noise level. The gradients across all rendered views direct the update of θ , ensuring that the NeRF-generated images align with the text descriptions. Additionally, we incorporate the 'perturb and average' technique from SJC for more robust \mathcal{L}_{SDS} . For a comprehensive understanding of these methods, the

reader is directed to the detailed explanations provided in [39, 53].

3.2 The Proposed CompoNeRF

3.2.1 Composition Module

CompoNeRF is designed to composite multiple NeRFs to reconstruct scenes featuring multiple objects, utilizing guidance from both bounding boxes and textual prompts. Within our framework, depicted in Fig. 4, the Axis-Aligned Bounding Box (AABB) ray intersection test algorithm is applied to ascertain intersections across each box in the global frame. Subsequently, we sample points x_g within the intervals of the ray-box and project them to x_l to deduce the corresponding color C_l and density σ_l within individual NeRF models. These properties are processed through our composition module to infer the global color C_g and density σ_g , crucial for the global rendering. Volume rendering techniques [16] are then employed to procure the rendered views for both local and global frames. We propose dual SDS losses to ensure coherence between the image outputs and their corresponding textual descriptions. Additionally, our approach facilitates recomposition by channeling samples from cached models back into local frames along with the text revision, thereby streamlining the integration.

Global Composition. The independent optimization of each local frame may inadvertently result in a lack of global coherence within the scene. To address this, our scene composition process is designed to integrate these frames, thereby achieving a more consistent result. Before exploring the specifics of the module, it is imperative to discuss two critical design decisions within the composition module, as depicted in Fig. 4. Upon integrating the properties inferred from x_g into the composition module, they are fine-tuned through gradients derived from the global SDS loss. This process leads to a critical consideration: the necessity of refining the global density σ_g . There are two potential methods: **1) Density-based:** The advantage of adjusting σ_g is that it can adjust geometry, thus yielding a scene more congruent with the global text prompt. However, this comes at the cost of potentially compromising the optimal color C_g , as calibrating σ_g introduces more uncertainty for subsequent color refinement as it requires prior density features h . **2) Color-based:** Conversely, directly employing σ_l mitigates this uncertainty but has less geometric control, presenting a balance to

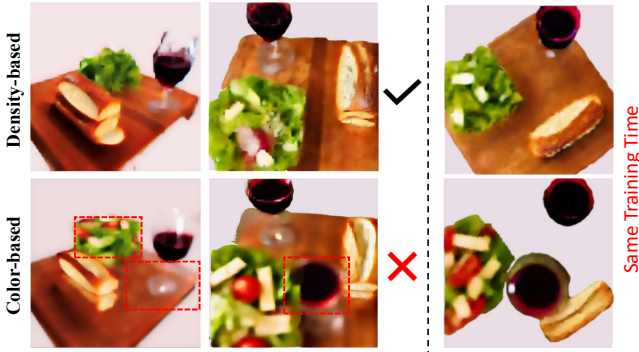


Figure 5: **Design Impact Comparison: Density vs. Color-based Methods.** The top row illustrates the density-based approach’s detailed rendering and quick convergence in the ‘table wine’ scene. The bottom row highlights the color-based method’s enhancements and its drawbacks, such as geometric and shadow inaccuracies, particularly in close-up views and slow convergence.




Figure 6: **Detail of Composition module:** density-based design.

strike in the pursuit of precise scene composition.

After thorough experiments, exemplified in Fig. 5, we have opted for the density-based approach to refine σ_g prioritizing both **accuracy and efficiency**. The test revealed that it excels in rendering intricate details, such as enhanced wood grain textures and more naturally contoured ‘salad’, as accentuated by boxes. This method also demonstrated a swifter convergence rate. Conversely, while the color-based improved reflections and reduced flickering on the ‘wine cup’, it was plagued by issues such as sparse density, which adversely brings holes at the base of the ‘cup’ and the corner of the ‘table’. Furthermore, upon close examination, it becomes evident that shadow artifacts of ‘wine’ on the ‘table’ are pronounced, suggesting that its disadvantages outweigh its advantages.

Network Design. The compositional framework of our network, as delineated in Fig. 6, is predicated on an architecture that employs a suite of MLPs, represented as $\{\theta_l\}_{l=1}^m$, each dedicated to a distinct local frame. To harmonize σ_l and C_l , we incorporate global MLPs, including density calibrator $f_{\theta_{sd}}$ and color calibrator $f_{\theta_{sc}}$. A transformation module complements this system, tasked with maintaining the spatial coherence between the global and local frames. It governs the transformation of sampling points x , ray directions d , and adjacent sampling distances δ . This module also orders the points $\{x_{g,j}\}_j$ by their distance to the global camera origin o_g , ensuring that each local point x_l is accurately matched with its corresponding global point x_g for subsequent volume rendering.

The network design is:

$$\sigma_g = \alpha_d f_{\theta_{sd}}(x_g) + \sigma_l, \quad (4)$$

$$C_g = \alpha_c f_{\theta_{sc}}(h, d_g) + C_l. \quad (5)$$

In contrast to the local frames, the global frame’s color output C_g is inferred based on h and conditional on d_g to enable a view-dependent lighting effect. Residual learning is leveraged here, where σ_l, C_l serve as foundational elements that support the learning of global density σ_g and color C_g . The parameters α_d, α_c are adjustable, allowing fine-tuning of the influence that local components exert on the global outputs. It is imperative to acknowledge that in our color-based method, density calibration is intentionally excluded to concentrate solely on the refinement of color dynamics as shown at Fig. 4. This is achieved by conditioning the process on both spatial and directional global inputs (x_g, d_g) , as demonstrated in the following equations:

$$\sigma_g = \sigma_l, \quad C_g = \alpha_c f_{\theta_{sc}}(x_g, d_g) + C_l. \quad (6)$$

The integration of extra x_g aims to facilitate a fair comparison under same inputs with the density-based. It enhances the visual appeal of effects like the wine cup’s reflection, as demonstrated in Fig. 5. However, this method is not without its compromises. It tends to produce artifacts and is characterized by a slower convergence rate. Additionally, this approach limits the ability to precisely control density, subsequently impacting the intricate geometric details.

3.2.2 Recomposition

Our architecture advances scene reconstruction by providing an intuitive interface for layout manipulation. This capability is crucial for the reconfiguration of scene elements into novel scenes, as depicted in Fig. 4. Here, the input panel allows for adjustments in the attributes of bounding boxes, such as modifying the position and scale of the ‘apple’ bounding box prior to composition. The refinement process further involves sampling ray-box intervals from the global frame, leading to transformed coordinates with the corresponding ray samples that are then incorporated into the pipeline, as demonstrated in Fig. 6. Each bounding box represents a NeRF, providing the flexibility to move, scale, or replace elements as needed. CompoNeRF’s capabilities also extend to textual edits, exemplified by the transformation of ‘wine’ into ‘juice’. Since NeRFs have been well trained, we only finetune θ_g, θ_l to align text prompts to promote consistency of both local and global views. Moreover, the NeRFs once retrained within the edited scene, are also structured to be decomposable and cacheable in future scene compositions.

3.2.3 Optimization

During optimization, our method employs dual text guidance to align rendering results with both global and local textual descriptions. The optimization objective is:

$$\mathcal{L} = \alpha_g \nabla \mathcal{L}_{\text{SDS}}(\hat{X}_g, T) + \alpha_l \sum_{j=1}^m \nabla \mathcal{L}_{\text{SDS}}(\hat{X}_{l,j}, T_{l,j}) + \beta \mathcal{L}_{\text{sparse}},$$

where T signifies the global text prompt, while T_l pertains to a specific object within the global context. The hyperparameters α_g, α_l , and β modulate the respective loss weights. As suggested in [29], we use $\mathcal{L}_{\text{sparse}}$ included to penalize the binary entropy of local NeRFs’ densities, thereby mitigating the issue of extraneous floating radiance. Additionally, incorporating directional cues such as ‘‘front view’’ or ‘‘side view’’ into the input text, as suggested by [39, 29] proves beneficial in specifying camera poses during the training phase, further enhancing the alignment of our generated scenes with the intended perspectives. We refer readers to the pseudo-code in our *suppl.* for our training procedure.

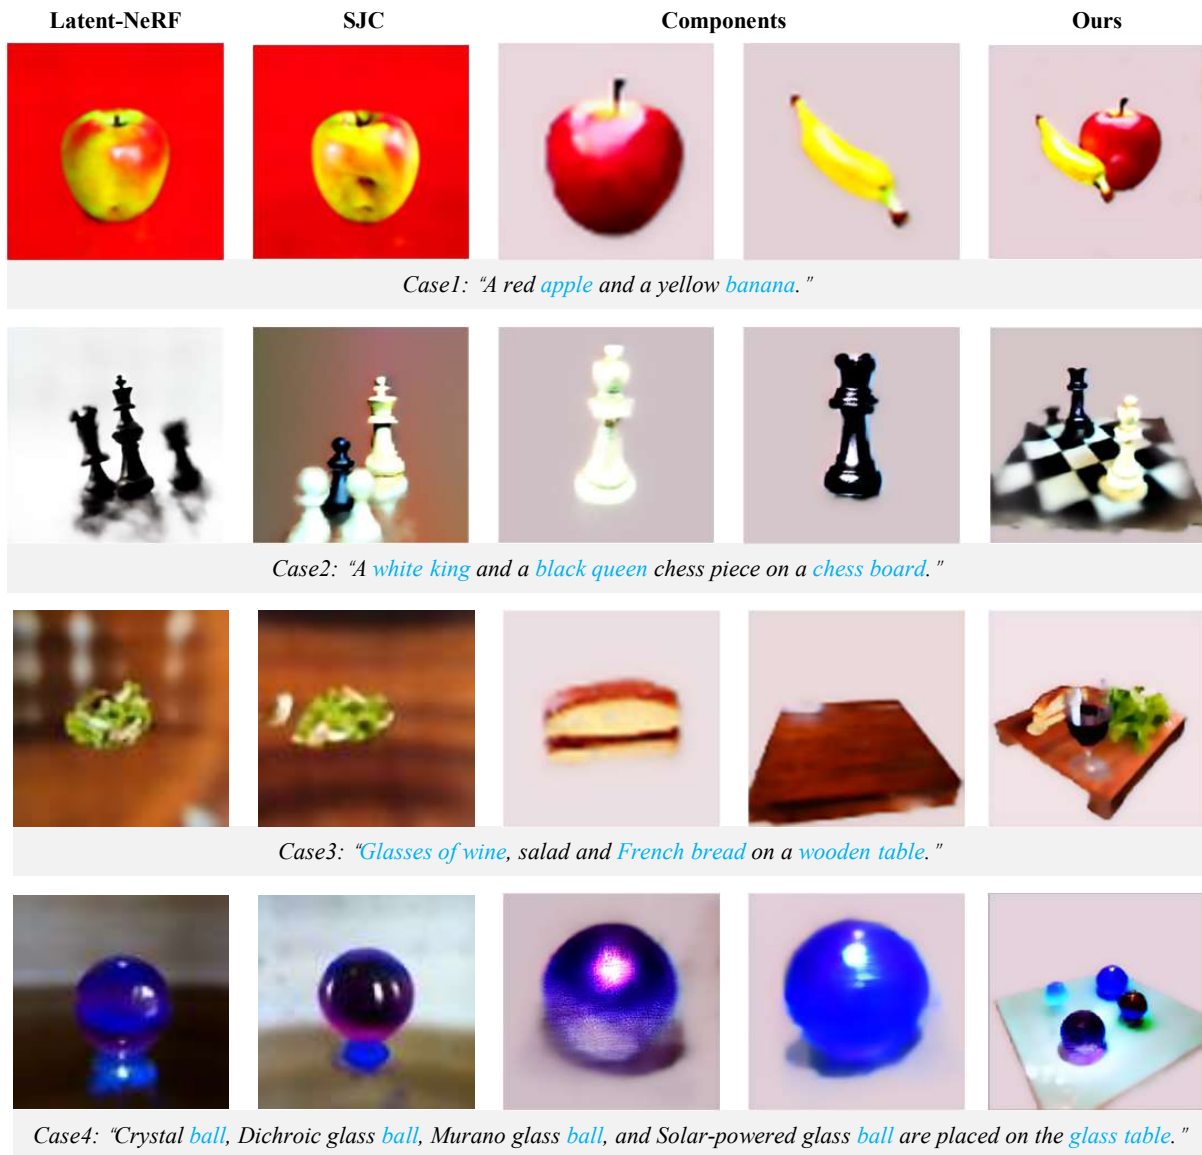


Figure 7: **Qualitative comparison with other text-to-3D methods using multi-object text prompts.** We refer readers to our *suppl.* and video demos for more visual results.

4 EXPERIMENTS

4.1 Qualitative Comparison

In Fig. 7, we conduct qualitative comparisons of 3D assets generated using our method against Latent-NeRF and SJC, all based on the same Stable Diffusion model. Our method exhibits a remarkable ability to generate complex 3D models from a wide array of multi-object text prompts, demonstrating **superior object identity accuracy and enhanced context relevance and richness** compared to its counterparts. In simpler Case 1, our CompoNeRF method adeptly generates two distinct objects: a red *apple* and a yellow *banana*. In stark contrast, competing methods amalgamate the features of both fruits into a singular, indistinct object. Case 2-4, which is more intricate, showcases the capability to render a realistic scene with accurately depicted objects. Conversely, our baselines struggle to produce even recognizable objects.

Method	Case 1	Case 2	Case 3	Case 4
LatentNeRF	27.69	31.19	21.55	29.51
SJC	28.21	30.53	23.33	28.76
CompoNeRF (Ours)	33.37	31.45	36.06	30.98

Table 1: **Quantitative comparison.** For our evaluation metric, we utilize the average of CLIP scores [38, 65, 55] across different views, which serve to assess the similarity between the generated images and the global text prompt.

4.2 Quantitative Comparison

In our study, we employ the CLIP score as the primary evaluation metric to assess the congruence between the generated 3D assets and the associated text prompts. This score, commonly used in text-to-image generation research as noted in studies [38, 65, 55], is derived from the cosine similarity between the embeddings of the text and the image, both encoded by the CLIP model. For 3D assets, we project images from various views and calculate the CLIP score




Figure 8: **The scene editing.** Demonstrated here are the stages of our recombination, utilizing cached source scenes. Each NeRF is individually identified by colorful labels. These decomposed nodes are then positioned in the initial layout and subsequently calibrated to form the final composition. The detailed description of the ambient environment is underscored, enhancing the scene’s realism.

concerning the global text prompt, with the overall CLIP score being the average of these values. As detailed in Tab. 3, our quantitative comparisons demonstrate the superior alignment of our method with the text prompt compared to the Latent-NeRF and SJC approaches in diverse scene configurations. Notably, in the challenging Case 3, our method shows a remarkable 54% improvement. The result of overall enhancement in multi-object scenes underscores the robustness of our global calibration strategy. We refer readers to our *suppl.* for additional results.

4.3 Recomposition for more complex scenes

To further validate CompoNeRF’s performance for more complex multi-object text inputs, we assess its performance using a lengthy sentence describing the color, texture, light, and relationships between scene components. Fig. 8 showcases our refined scene renderings originating from pre-trained source scenes including the ‘lamp’ and the ‘nightstand’ from the ‘bedroom’ scene. Then we add two glass balls ‘Murano’ and ‘Dichroic’ from Case.4 in Fig. 7. We observe that a direct amalgamation of these components can manifest various artifacts at the base of the lamp and incongruous shadows and reflections from the glass ball are notable, detracting from the authenticity of the scene’s ambiance. After composition, reconfigured objects are adeptly integrated, achieving a coherent and consistent global scene. The changes in the materials of the lamp, the nightstand, and their reflective surfaces demonstrate the system’s adaptability to diverse source inputs, which also involves a challenging text prompt containing subtle interplay within a multi-object context. In a nutshell, our method unleashes its potential by composing scenes with reusable NeRF components, which also facilitates the following editing for scenes described by lengthy sentences even with complex relationships among objects.

4.4 Ablation Study

Composition Strategy. Our initial tactic, as shown in Fig. 9(a), attempts to manage this by assigning each object to a distinct NeRF with a shared context, diverging from the traditional single-network scene encodings [39, 23, 29, 53]. This approach aims to improve object recognition, yet it’s prone to ‘guidance collapse’ due to the global text prompts’ inability to offer precise semantic delineation for individual objects. Inspired by more accurate rendering of single objects with targeted subtext prompts as seen in Fig. 2, the subsequent refinement in (b) utilizes specific textual supervision for each object. This tailored guidance confirms that diffusion models

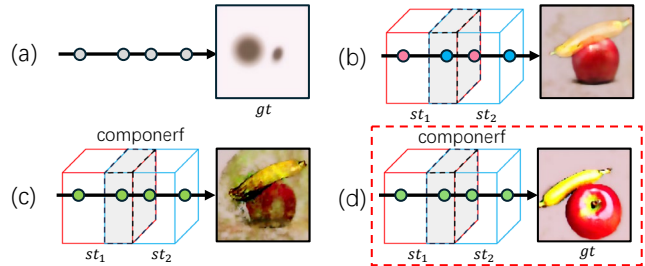


Figure 9: **The composition strategy.** Our proposed strategies for multi-object scene composition align with Eq. (2). The areas of NeRF overlap are indicated in gray. The green nodes represent composited samples. Our design is highlighted by the dashed box.





Figure 10: **Effect of adjusting camera distances.** Scaling camera distance can benefit rendering quality.

more effectively render individual objects when provided with explicit textual for each, thereby overcoming the ‘guidance collapse’ issue in complex multi-object scenarios. Nevertheless, the direct rendering of NeRFs often results in incorrect occlusions within their overlapping regions, indicating a deficiency in global refinement. To improve it, as detailed in (d), we refine sampling points that were originally guided by texts corresponding to individual objects, now with the incorporation of global textual guidance. This supplementary global oversight guarantees a harmonious rendering that upholds the unique identities of the objects while fostering overall compositional unity. Its vital role is underscored by its absence in (c), where its omission brings collapsed results. We refer readers to our *suppl.* for more ablation studies on this composition strategy.



“The game of *black queen* and *white king* chess was played on a *table*, with a *glass of wine*, and a *table lamp* providing the necessary illumination.”

Figure 11: **Recomposition process.** Individual components are labeled in **blue**, while **red** boxes emphasize areas of contrast.

Optimizing Diffusion Model Guidance for Scene Resolution. A critical aspect is the need to strike a balance between the overall scene’s resolution and the rendering details. For example, when a single object is placed within a vast scene, its rendering results may not be clear as it occupying only a few pixels. This small pixel footprint can limit the amount of gradient information received during backpropagation. Fig. 10 illustrates this scenario using the same text prompt but with varying scales of the global frames, ranging from 0.3 to 0.7. The results underscore a key insight: the more pixel rays an object interacts with, the better the quality of its rendering. This finding is particularly relevant for large-scale scene rendering, where multiple local frames coexist within the same space. A small object in such a setup may receive minimal ray-box interactions, potentially leading to training inefficiencies or collapse. It’s also important to consider that our scene models are optimized in the latent space of the Stable Diffusion model, which has a feature resolution of 64×64 . However, we decode these latent color features into RGB images with a resolution of 128×128 . This discrepancy affects the density of rays throughout the space and, by extension, the number of objects that can be effectively rendered.

Influence of Global MLP Size on Composition Capabilities. The complexity of a scene directly influences the required configuration of the parameters θ_g within the composition module. In Figure 12, we experiment with varying the number of layers in the MLPs responsible for both densities $f_{\theta_{s_d}}$ and color calibration $f_{\theta_{s_c}}$. The results indicate that an insufficiently representative MLP can fail to preserve the distinct identities of individual NeRFs. For instance, the ‘white king chess’ piece struggles to manifest its characteristic whiteness, leading the global calibrators to compensate inadequately by projecting a flattened representation onto the chessboard surface. By increasing the number of MLP layers, we observe a notable improvement in the accurate portrayal of each object’s identity and overall scene quality. For example, the contextual details of the chessboard, like its grid pattern, are rendered more clearly and naturally. Based on these findings, we recommend fine-tuning this hyperparameter to align the composition module’s capabilities with the specific complexity of the scene.

Global Context Assimilation by Local NeRFs and Composition Module. Despite the primary embedding of the context within the composition module, local NeRFs exhibit an ability to partially learn these global attributes. For instance, Fig. 11 shows that initially, the local NeRF optimization does not occur in isolation; the table from Case 3 in Fig. 7, for example, bears a residual shadow from the ‘French bread’ in the original configuration as depicted in the upper left corner. As the training progresses, these anomalies are resolved, and the local NeRFs gradually assimilate aspects of the global texture, albeit to a limited extent. For instance, while the black and white pattern of the chessboard is predominantly captured by the global composition module, the local representation of the

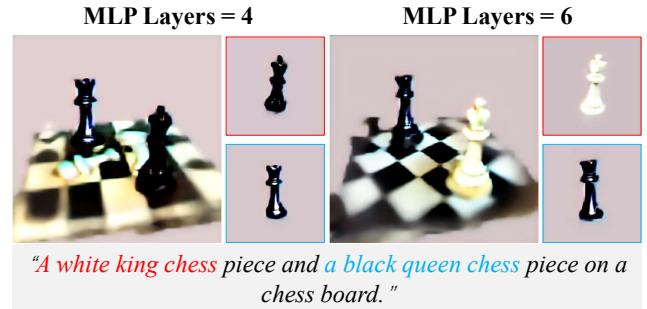


Figure 12: **Effect of MLP sizes.** Comparison of different parameter sizes for composition.


table, highlighted in red box in their upper left corners, remains unchanged. However, as iterations advance, despite there’s no NeRF that is responsible for ‘chess board’, the global frame begins to discern it as the necessary environment and replicate the underlying ‘chessboard’ pattern. This reveals that NeRFs can initially embed global environmental context, while the composition module can possibly merge some necessary local patterns for consistency. On the other hand, the early stop may benefit the potential degradation in specific elements, such as the ‘wine’, which worsens as training progresses, as evidenced by the local frames’ comparison in the upper left corner. Concurrently, the global rendering depicts the ‘wine’ as nearly imperceptible. This deterioration hints at the possibility that continued optimization may inadvertently diminish the representation of certain objects.

5 USER STUDY

In order to comprehensively evaluate the quality of the quality of 3D assets generated using our method, we conducted a user study involving 15 participants. The study consisted of three tasks: a composition correctness task, a generative quality evaluation task, and an object identification task. The specific design of each task will be elaborated in subsequent sections. Data collection involved both quantitative analysis, through subjective ratings, and qualitative research, carried out via semi-structured interviews.

5.1 Participants

The study included 11 participants, with 54.5% falling within the 18-24 age group and 45.5% in the 25-34 age range. The gender distribution was 63.6% male and 36.4% female, reflecting a balanced representation. In terms of prior experience, 18.2% of participants had some familiarity with 3D generative models, while having experience with Blender within the last six months. As depicted in Figure 13, the experimental procedure involved participants viewing videos showcasing the 3D assets generated by the model. This



Q: On a scale of 1 (low) to 7 (high), how would you rate the quality of the generated images from the inputs?

Figure 13: User study with our questions demo.

approach allowed participants to evaluate the outputs in a controlled and standardized manner, ensuring consistent exposure to the same visual stimuli across all participants.

5.2 Task design

- Composition Correctness Evaluation.** In this task, we assess the consistency of the generated 3D assets across two views, focusing on both semantic consistency and multi-view consistency. We collected four groups of samples, each comprising 3D assets generated by Latent-NeRF, SJC, and our method, using the same text prompt. Participants were asked to evaluate the generative consistency from both semantic and multi-view perspectives on a scale from 1 (low) to 7 (high).
- Generative Quality Evaluation.** We provided four groups of generated 3D assets (refer to the supplementary material) to each participant. For each group, the 3D assets were created using Latent-NeRF, SJC, and our method, all based on the same text prompt. Participants were then asked to assess the quality of these generated assets. Finally, we evaluated our method’s capability for multi-object generation and combination by calculating the match rate between the objects identified by the participants and those described in the prompt.
- Object Identification.** For this task, we selected four samples of 3D assets generated using our method, comprising a total of seven objects. Participants were then asked to identify the objects depicted in these assets. To evaluate the multi-object generation and combination capabilities of our approach, we calculated the accuracy of the participants’ object recognition.

5.3 Measurements

In the first task, we utilized a 7-point Likert scale to measure participants’ perceptions across two dimensions, including semantic consistency and multiview consistency. For the second task, we also used the 7-point Likert scale to evaluate the generative quality and made a comparison with existing works, including Latent-NeRF and SJC. Lastly, we evaluate the effectiveness of our method for multi-object combinations by counting the accuracy of subjects in finding the objects provided in the prompt from the generated 3d assets. We also encourage participants to provide feedback on their emotional responses and immersive experiences during the overall experience. For details of questions, please refer to *suppl.*

5.4 Results

As shown in Fig. 14(a), our method demonstrates significantly improved semantic consistency across all four cases, markedly outperforming previous approaches. As shown in Fig. 14(b), benefiting from the reduction of semantic chaos, the multi-view consistency is also better than previous methods. As shown in Fig. 14(c), in terms of generative quality evaluation, including the rendering quality of videos and object details, the performance of the three methods is comparable due to their utilization of the same Stable Diffusion version and identical training resolutions. As shown in Fig. 14(d), the experiments demonstrate that the majority of our generated 3D

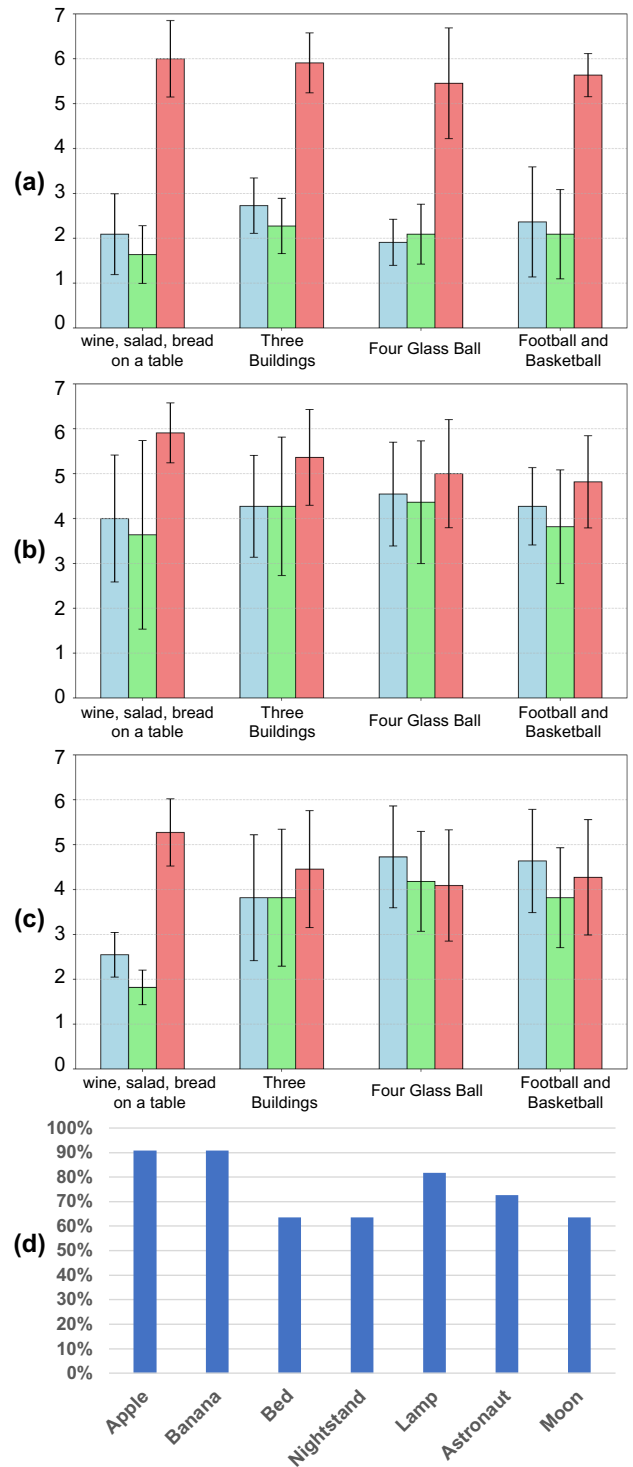


Figure 14: (a) Generative Consistency Evaluation (Semantic); (b) Generative Consistency Evaluation (Multi-view); (c) Generative Quality Evaluation; (d) Object Identification Task.

scenes are correctly recognized, affirming the semantic consistency between the generated content and the input prompt.

Discussion: Our observations indicate that users prioritize the level of detail in individual objects when assessing the quality of generated scenes, even if the objects’ semantics do not match the given

text prompts. For example, as shown at our *suppl.*, when presented with the prompt 'a football and a basketball,' users favored our baseline results, which depicted a single 'football' or 'basketball' with more intricate details, thus appearing more realistic. In contrast, our generated scenes, which included more objects, were perceived as less clear due to the increased complexity in rendering.

6 CONCLUSION AND FUTURE WORK

In this work, we have proposed a novel framework for multi-object text-guided compositional 3D scene generation with an editable 3D scene layout. Our framework interpreted a multi-object text prompt as a collection of localized NeRFs, each associated with a spatial box and an object-specific text prompt, which were then composited to render the entire scene view. We have further enhanced the framework with a specialized composition module for global consistency, effectively mitigating the issue of guidance collapse in the multi-object generation. Utilizing Stable Diffusion model, we have demonstrated that our method, the first to apply a compositional NeRF design to the text-to-3D task, can produce high quality 3D models that feature multiple objects and perform well compared with contemporaneous methods. Looking ahead, we have explored a promising application of CompoNeRF in the realm of scene editing, which allows for the reuse of trained models in scene recomposition. This capability opens up new possibilities and identifies a rich vein of future work to be pursued in this domain.

REFERENCES

- [1] D. Azinović, R. Martin-Brualla, D. B. Goldman, M. Nießner, and J. Thies. Neural rgb-d surface reconstruction. In *CVPR*, pp. 6290–6301, 2022. 3
- [2] H. Bai, Y. Lin, Y. Chen, and L. Wang. Dynamic plenotree for adaptive sampling refinement in explicit nerf. *ArXiv*, abs/2307.15333, 2023. 3
- [3] J. T. Barron, B. Mildenhall, M. Tancik, P. Hedman, R. Martin-Brualla, and P. P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *ICCV*, pp. 5835–5844. IEEE, 2021. 2, 3
- [4] J. T. Barron, B. Mildenhall, M. Tancik, P. Hedman, R. Martin-Brualla, and P. P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5835–5844, 2021. 14
- [5] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5460–5469, 2021. 14
- [6] R. Chen, Y. Chen, N. Jiao, and K. Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023. 3
- [7] D. Cohen-Bar, E. Richardson, G. Metzer, R. Giryes, and D. Cohen-Or. Set-the-scene: Global-local training for generating controllable nerf scenes. Mar 2023. 14
- [8] N. Deng, Z. He, J. Ye, B. Duinkharjav, P. Chakravarthula, X. Yang, and Q. Sun. Fov-nerf: Foveated neural radiance fields for virtual reality, 2022. 3
- [9] D. Epstein, B. Poole, B. Mildenhall, A. A. Efros, and A. Holynski. Disentangled 3d scene generation with layout learning. *arXiv preprint arXiv:2402.16936*, 2024. 3
- [10] C. Gao, A. Saraf, J. Kopf, and J.-B. Huang. Dynamic view synthesis from dynamic monocular video. In *ICCV*, pp. 5712–5721, 2021. 3
- [11] J. Gao, T. Shen, Z. Wang, W. Chen, K. Yin, D. Li, O. Litany, Z. Gocic, and S. Fidler. Get3d: A generative model of high quality 3d textured shapes learned from images. 2022. 2
- [12] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2
- [13] F. Hong, M. Zhang, L. Pan, Z. Cai, L. Yang, and Z. Liu. Avatarclip: Zero-shot text-driven generation and animation of 3d avatars. *arXiv preprint arXiv:2205.08535*, 2022. 2
- [14] A. Jain, B. Mildenhall, J. T. Barron, P. Abbeel, and B. Poole. Zero-shot text-guided object generation with dream fields. In *CVPR*, pp. 867–876, 2022. 2, 3
- [15] L. Jiang, R. Ji, and L. Zhang. Sdf-3dgan: A 3d object generative method based on implicit signed distance function. *arXiv preprint arXiv:2303.06821*, 2023. 3
- [16] J. T. Kajiya and B. P. Von Herzen. Ray tracing volume densities. *ACM SIGGRAPH*, 18(3):165–174, 1984. 3, 4
- [17] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 12
- [18] H.-H. Lee and A. X. Chang. Understanding pure clip guidance for voxel grid nerf models. *arXiv preprint arXiv:2209.15172*, 2022. 2
- [19] C. Li, S. Li, Y. Zhao, W. Zhu, and Y. Lin. Rt-nerf: Real-time on-device neural radiance fields towards immersive ar/vr rendering, 2022. 3
- [20] J. Li, D. Li, C. Xiong, and S. Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, pp. 12888–12900. PMLR, 2022. 2
- [21] Z. Li, M. Hu, Q. Zheng, and X. Jiang. Connecting consistency distillation to score distillation for text-to-3d generation. *arXiv preprint arXiv:2407.13584*, 2024. 3
- [22] Y. Liang, X. Yang, J. Lin, H. Li, X. Xu, and Y. Chen. Luciddreamer: Towards high-fidelity text-to-3d generation via interval score matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6517–6526, 2024. 3
- [23] C.-H. Lin, J. Gao, L. Tang, T. Takikawa, X. Zeng, X. Huang, K. Kreis, S. Fidler, M.-Y. Liu, and T.-Y. Lin. Magic3d: High-resolution text-to-3d content creation. *arXiv preprint arXiv:2211.10440*, 2022. 2, 3, 7
- [24] D. B. Lindell, J. N. Martel, and G. Wetzstein. Autoint: Automatic integration for fast neural volume rendering. In *CVPR*, pp. 14556–14565, 2021. 3
- [25] L. Liu, J. Gu, K. Z. Lin, T.-S. Chua, and C. Theobalt. Neural sparse voxel fields. *ArXiv*, abs/2007.11571, 2020. 3
- [26] L. Liu, M. Habermann, V. Rudnev, K. Sarkar, J. Gu, and C. Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Transactions on Graphics (TOG)*, 40(6):1–16, 2021. 3
- [27] Z. Ma, Y. Wei, Y. Zhang, X. Zhu, Z. Lei, and L. Zhang. Scaledreamer: Scalable text-to-3d synthesis with asynchronous score distillation. *arXiv preprint arXiv:2407.02040*, 2024. 3
- [28] S. McAuley, S. Hill, N. Hoffman, Y. Gotanda, B. Smits, B. Burley, and A. Martinez. Practical physically-based shading in film and game production. In *ACM SIGGRAPH 2012 Courses*, Aug 2012. doi: 10.1145/2343483.2343493 3
- [29] G. Metzer, E. Richardson, O. Patashnik, R. Giryes, and D. Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. *arXiv preprint arXiv:2211.07600*, 2022. 2, 3, 5, 7, 12, 15
- [30] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2
- [31] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, vol. 12346, pp. 405–421. Springer, 2020. 3
- [32] A. Mirzaei, Y. Kant, J. Kelly, and I. Gilitschenski. Laterf: Label and text driven object radiance fields. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III*, pp. 20–36. Springer, 2022. 3
- [33] N. Mohammad Khalid, T. Xie, E. Belilovsky, and T. Popa. Clip-mesh: Generating textured meshes from text using pretrained image-text models. In *SIGGRAPH Asia 2022 Conference Papers*, pp. 1–8, 2022. 2
- [34] T. Müller, A. Evans, C. Schied, and A. Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022. 2, 3, 12
- [35] T. Müller, A. Evans, C. Schied, and A. Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41:1–15, 2022. 14
- [36] A. Q. Nichol and P. Dhariwal. Improved denoising diffusion proba-

- bilistic models. In *ICML*, pp. 8162–8171. PMLR, 2021. 2
- [37] J. Ost, F. Mannan, N. Thuerey, J. Knodt, and F. Heide. Neural scene graphs for dynamic scenes. In *CVPR*, pp. 2856–2865, 2021. 3
- [38] G. Parmar, K. Kumar Singh, R. Zhang, Y. Li, J. Lu, and J.-Y. Zhu. Zero-shot image-to-image translation. In *ACM SIGGRAPH 2023 Conference Proceedings*, pp. 1–11, 2023. 6, 15
- [39] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 2, 3, 4, 5, 7
- [40] A. Pumarola, E. Corona, G. Pons-Moll, and F. Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *CVPR*, pp. 10318–10327, 2021. 3
- [41] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pp. 8748–8763. PMLR, 2021. 2
- [42] A. Raj, S. Kaza, B. Poole, M. Niemeyer, N. Ruiz, B. Mildenhall, S. Zada, K. Aberman, M. Rubinstein, J. Barron, et al. Dreambooth3d: Subject-driven text-to-3d generation. *arXiv preprint arXiv:2303.13508*, 2023. 3
- [43] S. Rojas, J. Zarzar, J. C. Perez, A. Sanakoyeu, A. Thabet, A. Pumarola, and B. Ghanem. Re-nerf: Real-time rendering of nerfs across devices, 2023. 3
- [44] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10684–10695, 2022. 2, 3, 12
- [45] A. Sanghi, H. Chu, J. G. Lambourne, Y. Wang, C.-Y. Cheng, M. Fumero, and K. R. Malekshan. Clip-forge: Towards zero-shot text-to-shape generation. In *CVPR*, pp. 18603–18613, 2022. 2
- [46] C. Schuhmann, R. Beaumont, R. Vencu, C. Gordon, R. Wightman, M. Cherti, T. Coombes, A. Katta, C. Mullis, M. Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402*, 2022. 2
- [47] T. Shen, J. Gao, K. Yin, M.-Y. Liu, and S. Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. In *Neural Information Processing Systems*, 2021. 14
- [48] Y. Song, C. Kong, S. Lee, N. Kwak, and J. Lee. Towards efficient neural scene graphs by learning consistency fields. *arXiv preprint arXiv:2210.04127*, 2022. 3
- [49] P. P. Srinivasan, B. Deng, X. Zhang, M. Tancik, B. Mildenhall, and J. T. Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *CVPR*, pp. 7495–7504, 2021. 3
- [50] E. Tretschk, A. Tewari, V. Golyanik, M. Zollhöfer, C. Lassner, and C. Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *ICCV*, pp. 12959–12970, 2021. 3
- [51] D. Verbin, P. Hedman, B. Mildenhall, T. E. Zickler, J. T. Barron, and P. P. Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *CVPR*, pp. 5481–5490. IEEE, 2022. 3
- [52] C. Wang, M. Chai, M. He, D. Chen, and J. Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *CVPR*, pp. 3835–3844, 2022. 2, 3
- [53] H. Wang, X. Du, J. Li, R. A. Yeh, and G. Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. *arXiv preprint arXiv:2212.00774*, 2022. 2, 3, 4, 7, 12, 14, 15
- [54] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 34:27171–27183, 2021. 3
- [55] S. Wang, C. Saharia, C. Montgomery, J. Pont-Tuset, S. Noy, S. Pellegrini, Y. Onoe, S. Laszlo, D. J. Fleet, R. Soricut, et al. Imagen editor and editbench: Advancing and evaluating text-guided image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18359–18369, 2023. 6, 15
- [56] Z. Wang, C. Lu, Y. Wang, F. Bao, C. Li, H. Su, and J. Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [57] Q. Wu, X. Liu, Y. Chen, K. Li, C. Zheng, J. Cai, and J. Zheng. Object-compositional neural implicit surfaces. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVII*, pp. 197–213. Springer, 2022. 3
- [58] W. Xian, J.-B. Huang, J. Kopf, and C. Kim. Space-time neural irradiance fields for free-viewpoint video. In *CVPR*, pp. 9421–9431, 2021. 3
- [59] J. Xu, X. Wang, W. Cheng, Y.-P. Cao, Y. Shan, X. Qie, and S. Gao. Dream3d: Zero-shot text-to-3d synthesis using 3d shape prior and text-to-image diffusion models. *arXiv preprint arXiv:2212.14704*, 2022. 2
- [60] Y. Xu, M. Chai, Z. Shi, S. Peng, I. Skorokhodov, A. Siarohin, C. Yang, Y. Shen, H.-Y. Lee, B. Zhou, et al. Discoscene: Spatially disentangled generative radiance fields for controllable 3d-aware scene synthesis. *arXiv preprint arXiv:2212.11984*, 2022. 3
- [61] B. Yang, Y. Zhang, Y. Xu, Y. Li, H. Zhou, H. Bao, G. Zhang, and Z. Cui. Learning object-compositional neural radiance field for editable scene rendering. In *ICCV*, pp. 13779–13788, 2021. 3
- [62] C. Yu, Q. Zhou, J. Li, Z. Zhang, Z. Wang, and F. Wang. Points-to-3d: Bridging the gap between sparse points and shape-controllable text-to-3d generation. *arXiv preprint arXiv:2307.13908*, 2023. 3
- [63] X. Yu, Y.-C. Guo, Y. Li, D. Liang, S.-H. Zhang, and X. Qi. Text-to-3d with classifier score distillation. *arXiv preprint arXiv:2310.19415*, 2023. 3
- [64] X. Zhang, P. P. Srinivasan, B. Deng, P. Debevec, W. T. Freeman, and J. T. Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (TOG)*, 40(6):1–18, 2021. 3
- [65] Z. Zhang, L. Han, A. Ghosh, D. N. Metaxas, and J. Ren. Sine: Single image editing with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6027–6037, 2023. 6, 15
- [66] S. Zhi, T. Laidlow, S. Leutenegger, and A. J. Davison. In-place scene labelling and understanding with implicit scene representation. In *ICCV*, pp. 15838–15847, 2021. 3

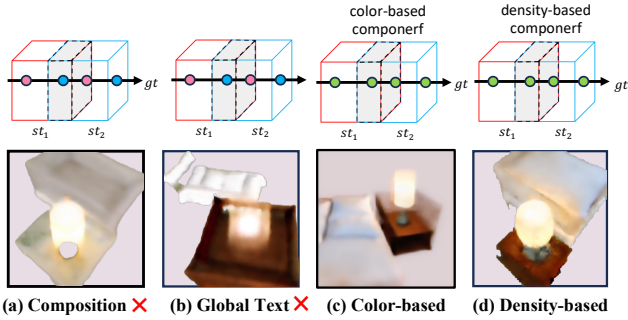


Figure 15: Ablation study on module designs with the scene bedroom. (a) without global calibration. (b) without global text loss. (c) color-based design. (d) our density-based design.

We provide more details of the proposed method and experimental results in the supplementary material. Sec.1 and Sec.2 provide the algorithm and more implementation details. Sec.3 provides more insights into our CompoNeRF model. Sec.4 adds more details of visualization results. Sec.5 lists our attached material details for both scene reconstruction and editing.

7 ALGORITHM

The detailed algorithm of training our proposed CompoNeRF is shown in Algorithm 1.

8 IMPLEMENTATION DETAILS

For score distillation sampling, we use the v1-4 checkpoint of Stable Diffusion based on the latent diffusion model [44]. We utilize the code-base [29] for 3D representation and grid encoder from Instant-NGP [34] as our NeRF model. The global MLP consists of 4 or 6 Linear layers with 64 hidden channels. In the training loss, we set $\alpha_g = 100$, $\alpha_l = 100$, and $\beta = 5e^{-4}$ if without specification. Our 3D scenes are optimized with a batch size of 1 using the Adam [17] optimizer on a single RTX3090. Our global frame is centered at the world origin and has a normalized side length of $[-1,1]$. To generate camera positions, we uniformly sample points on a hemisphere that covers the global frame with a random radius between 1.0 and 1.5. The camera distance can also be scaled in the way discussed in the main paper. Plus, cameras are oriented to look toward the objects. During optimization, the camera field of view is randomly sampled between 40 and 70 degrees. At test time, the field of view is fixed at 60 degrees. We use the Adam optimizer and perform gradient descent at a learning rate of 0.001 for 5,000 steps in simple prompts, such as ‘‘apple and banana’’, and 8,000 steps in more complex prompts for better quality. We follow the implementation of SJC [53] to perform the averaging implicitly, relying on the optimizer’s momentum state when applying the perturb-and-average scoring strategy during training.

9 DISCUSSION

More ablations on the design of composition module. In Fig. 15, we present further results from the ablation study of our composition module. As outlined in our main manuscript, our preference for a density-based approach is due to its effective and precise calibration of global density. For example, the ‘bedroom’ scene builds upon the discussion from Fig.2(b.2) in the main paper. The complementary study in Fig. 15(a) demonstrates that direct global text supervision without compositional integration leads to a loss of material context, washing out the ‘bed’ and ‘nightstand’ in white. Conversely, Fig. 15(b) illustrates that omitting global text and relying solely on subtext supervision retains the familiar context of a ‘white sheet’ bed and a polished tan ‘nightstand’. However, this approach

Text: *A bed is next to a nightstands where a table lamp on it in a room.*




Figure 16: Multi view results with the text prompts on Fig. 15. Sub-texts for individual NeRFs are highlighted in bold.

introduces geometric inconsistencies, such as an overly tall nightstand and a lamp lacking a base, along with an absence of light reflection in the surrounding space. The application of our composition module, depicted in Fig. 15(c) and 8(d), reveals that the density-based design affords enhanced control over density and, consequently, finer geometry. Instead of an empty space above the nightstand, the design aims to adjust the nightstand’s height to achieve scene harmony, although it still exhibits limitations in controlling density, leading to subdued floating radiance. Fig. 16 provides a comprehensive visual comparison within the ‘bedroom’ context. When global calibration is absent, as seen in Fig. 16(a), the scene is plagued by sparse holes and a loss of color and texture detail. Neglecting the global branch entirely, as shown in Fig. 16(b), results in a lack of global consistency, evident in the disproportionate size of the ‘nightstand’ relative to the scene. Finally, the color-based solution in Fig. 16(c) fails to effectively correct the geometry, introducing additional artifacts. In contrast, the full model in Fig. 16(d) exhibits a marked improvement in these aspects.

Addressing the multi-face issue with enhanced prompts. Much like Latent-NeRF and SJC, our CompoNeRF framework encounters the multi-face challenge, where guidance from the Stable Diffusion model may result in conflicting facial features for certain objects, as illustrated in Figure 17. The reason lies in the fact that diffusion model does not always provide reliable guidance that aligns with the desired orientation corresponding to the camera’s viewpoint during sampling. To mitigate the multi-face problem, stronger constraints can be introduced to promote geometric consistency within the 3D representation. CompoNeRF incorporates mesh constraints, akin to those utilized in Latent-NeRF, offering a more detailed 3D layout compared to traditional bounding boxes. As demonstrated in Figure 17, the implementation of exact mesh constraints markedly mitigates the multi-face issue, though it may come at the expense of detail and adaptability. Nevertheless, the requirement for accurate mesh input necessitates considerable manual editing, which may reduce the method’s range of applications. Despite this, our approach illustrates that the 3D scene layout can be readily adapted

Algorithm 1 Training for CompoNeRF

Input: a pre-trained text-to-image diffusion model ϕ , multi-object text prompt T and a set of boxes for 3D scene layout. **Output:** learned parameters of local NeRFs $\{\theta_{l,i}\}_{i=1}^m$ and Global MLP θ_g .

```
1: for Iter = 0 < MaxIter do
2:   Sample  $H \times W$  rays from the random camera position and add the directional prompt into  $T$ .
3:   for  $i = 0$  to  $H \times W$  do
4:     Calculate the ray-box intersection for ray  $r_i$  to get  $m_i$  hits.
5:     for  $j = 0$  to  $m_i$  do
6:       Sample  $N$  points with normalized location in the  $j$  hit local frame.
7:       Calculate color  $C_l$  and density  $\sigma_l$  for each point from  $\theta_{l,j}$ .
8:       Calculate the volumetric rendering color  $\hat{C}_{l,i,j}$  of the local Frame for ray  $r_i$ .
9:     end for
10:  end for
11:  Map all points into global locations and sort them according to the depth.
12:  Calculate the calibrated color  $C_{g,i}$ ,  $\sigma_{g,i}$  via Eq. 4 and Eq. 5 for each point.
13:  Calculate the global volumetric rendering color  $\hat{C}_{g,i}$  for each ray  $r_i$ .
14: end for
15: Generate the local view from  $\{\hat{C}_{l,i,j}\}_{i=1}^{H \times W}$  and the global view from  $\{\hat{C}_{g,i}\}_{i=1}^{H \times W}$ 
16: Perform score distillation sampling on the local render view and the global render view.
17: Update network parameters via an Adam optimizer.
```

Eng: Decompose local NeRFs and cache them into offline dataset.

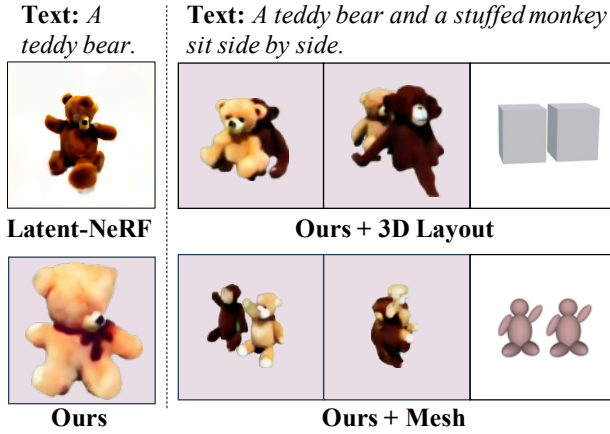


Figure 17: **(Left)** we observe the multi-face problem, e.g., duplicated face views with geometry collapse in all methods, even in single-object cases. **(Right)** We provide mesh as guidance instead of box layouts to solve this problem, which further proves our method’s versatility and effectiveness.

to accommodate a broader range of input prompts. Further study is needed to solve the persistent multi-face issue in the text-to-3D tasks.

Adjusting Layout to address Floating Artifacts. The process of scene composition begins with strategically positioning NeRFs within the predefined layout. An overlap of object bounding boxes is critical, as highlighted in Fig.2 of the main document, to facilitate the generation of convincing scenes. In our investigations, demonstrated in Fig. 18, we identified a ‘floating’ issue when bounding box overlaps are absent. This issue may stem from the regularization behavior within NeRFs, where the radiance fields—specifically, the regions responsive to gradient interactions, symbolized by ellipses centered on the boxes—fail to intersect. Such non-interaction can pose challenges, as it does not provide the necessary contiguous context for the global semantics to incorporate these objects seamlessly. To rectify this, one straightforward approach we recommend is the judicious repositioning of bound-




Figure 18: An ablation study examines layout editing and the floating issue. The upper row shows the layout editing and the lower row indicates rendering views. **(Left)** Renderings exhibit floating objects due to a suboptimal layout. **(Right)** Improved outcomes following layout refinements.


ing boxes to introduce overlaps. For example, a slight downward adjustment of a box can instigate detectable overlaps during training, facilitating better integration. This insight opens up a potential avenue of research into the interplay between layout configurations and NeRFs, offering the possibility of more nuanced control over scene dynamics without the need for explicit layout modifications.

Analysis of Failure Cases in Scene Composition and Editing. Our composition module may sometimes fail to produce coherent scenes, often due to limited text description distributions within the training data of diffusion models as illustrated in Fig. 19. This can be mitigated by adjusting the loss weights governing the global and local guidance, such as Fig. 20. In scene composition, the ‘computer station’ lacks accessories like cables and wires, the ‘headphones’ are misshapen, and the ‘computer screen’ lacks a base.

Methods	Diffusion Model	3D Representation	Scene Rendering	Input Prompt	Scene Editing	recomposition
DreamFusion	Imagen	Mip-NeRF 360 [5]	Object-centric	Text	T	✗
Magic3D	eDiff-I + SD	Instant-NGP [35]	Object-centric	Text	T	✗
DreamBooth3D	DreamBooth+DreamFusion	Mip-NeRF [4]	Object-centric	Text+Images	T	✗
Points-to-3D	ControlNet+Point-E	Instant-NGP	Object-centric	Text+Image	T	✗
Fantasia3D	SD+PBR	DMTET [47]	Object-centric	Text/Fine Shape	T/M/S/R	✓
Latent-NeRF	SD V1.4	Instant-NGP	Object-centric	Text+Fine Shape	T	✗
SJC [53]	SD V1.5	voxel radiance field	Object-centric	Text	T	✗
Set-the-Scene [7]	SD V2.0	-	Object-compositional	Text+3D Layout	T/M/S/R	✓
Ours	SD V1.4	Instant-NGP	Object-compositional	Text+3D Layout	T/M/S/R	✓


Table 2: **Comparison of our method with the related works for text-to-image generation.** SD denotes Stable Diffusion. For scene editing, we use T(editing object with text), M(moving object), S(scaling object), and R(removing object) for short.

Text: The computer station with a large monitor, a mechanical keyboard and gaming mouse, and noise-cancelling headphones.



Composition Decomposition Stable Diffusion


Text: After a long day in space, the tired astronaut is standing by the bed.



Init Recomposed Stable Diffusion

Figure 19: The failure cases are categorized as follows: **(Left)** issues encountered during scene reconstruction, and **(Right)** challenges arising in scene editing. The outputs from Stable Diffusion selected for illustration represent the most frequently occurring types generated by the model.


Text: The *boy* was gleefully riding his *bike* down the *hill*, feeling the wind rush past his face.



$$\alpha_l = 2\alpha_g$$

$$\alpha_l = \alpha_g$$

Text: The *coffee* sat next to a *plate* of warm *croissants*, creating a cozy breakfast scene on the rustic *table*.



$$\alpha_l = 2\alpha_g$$

$$\alpha_l = \alpha_g$$

Figure 20: The failure cases with their color-labeled NeRFs components are shown beside. The textual guidance manipulation on global/local weights is shown below.

Scene recomposition similarly shows the 'astronaut' and 'bed' placed together without sensible global calibration. Moreover, the Stable Diffusion model's depiction of human figures often suffers from geometric distortions, potentially due to the multi-face problem, as shown in Fig. 20. These failures are mostly due to uncommon layouts or the rarity of certain objects in text-to-image datasets. Repeated global text prompts on Stable Diffusion and examination of numerous samples have failed to yield images that align with our objectives. This challenge extends beyond guidance collapse, reflecting the scarcity of certain objects in the model's outputs. CompNeRF's effectiveness is inherently linked to the performance of large-scale text-to-image models, restricting its capabilities to generating primarily conventional scenes with well-defined global features. To address these limitations, we can strategically adjust the weights of global textual guidance $\alpha_g \nabla \mathcal{L}_{SDS_g}$ and local textual guidance $\alpha_l \nabla \mathcal{L}_{SDS_l}$. This adjustment aims to find an equilibrium between the consistency of the overall scene and the accuracy of in-

dividual components. For example, increasing α_g enhances global consistency, as evidenced in Figure 20. However, this can inadvertently lead to objects assimilating extraneous global context. In the 'boy riding bike' scenario, a heightened α_g may result in the 'bike' being erroneously represented as both a human figure and a bike. Similarly, in the 'breakfast' scene, amplifying global context might result in a more proportionate table, yet it complicates the distinction between the 'croissant' and its individual NeRF representation.

Ultimately, fine-tuning loss weight parameters is a delicate process that can mitigate identity issues, yet it demands careful calibration to maintain a harmony between scene integrity and the authenticity of each component. The limited representation of certain objects in pre-trained models remains a substantial obstacle, underscoring the need for further investigation into the issue of inadequate guidance in complex scene generation.

Influence of SD version. As depicted in Fig. 21, there is a significant gap between the results using SD V1.5 and SD V1.4, sug-

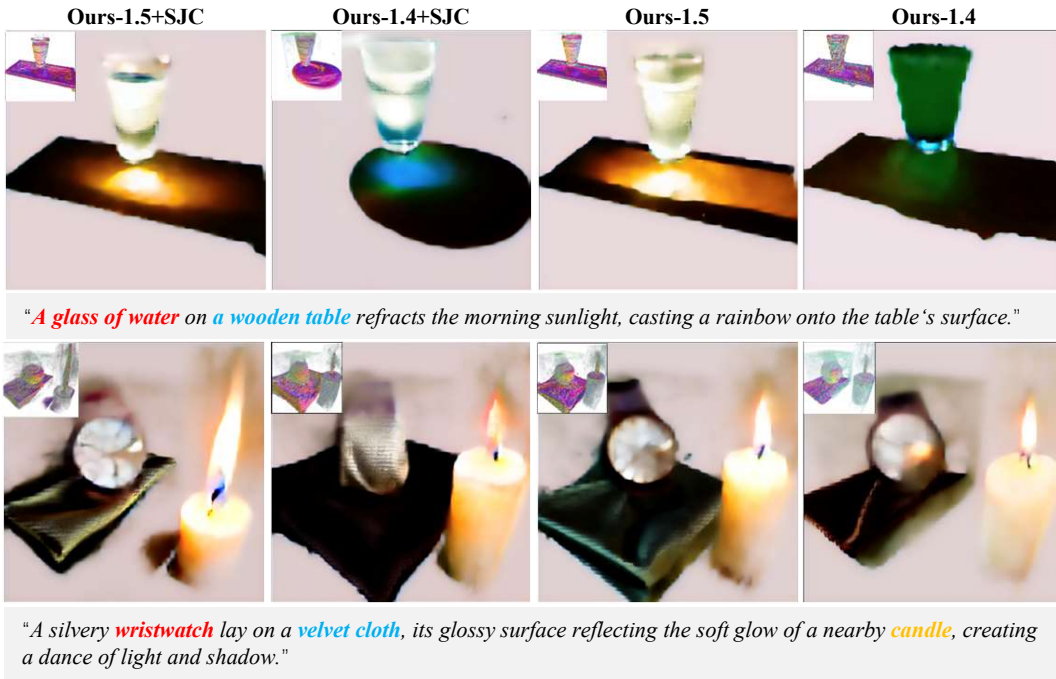


Figure 21: The influence of different weight versions of Stable Diffusion.

Method	Case 5	Case 6	Case 7	Case 8
LatentNeRF	25.16	27.07	26.32	27.43
SJC	23.55	27.84	27.41	25.62
CompoNeRF (Ours)	26.13	32.71	28.44	28.96

Table 3: **More quantitative comparisons.** For our evaluation metric, we utilize the average of CLIP scores [38, 65, 55] across different views, which serve to assess the similarity between the generated images and the global text prompt.

gesting that upgrading to a more advanced diffusion model could further enhance the quality of the generated content.

10 MORE VISUALIZATION RESULTS

We provide the multi-view qualitative results from CompoNeRF in Fig. 22. Note that we increase the resolution of the image latent features from 64×64 to 128×128 during inference for better results. We also have attached video results in the supplemental materials for each case and the baseline Latent-NeRF [29] and SJC [53]. Please see the attached video for rotating frame results.

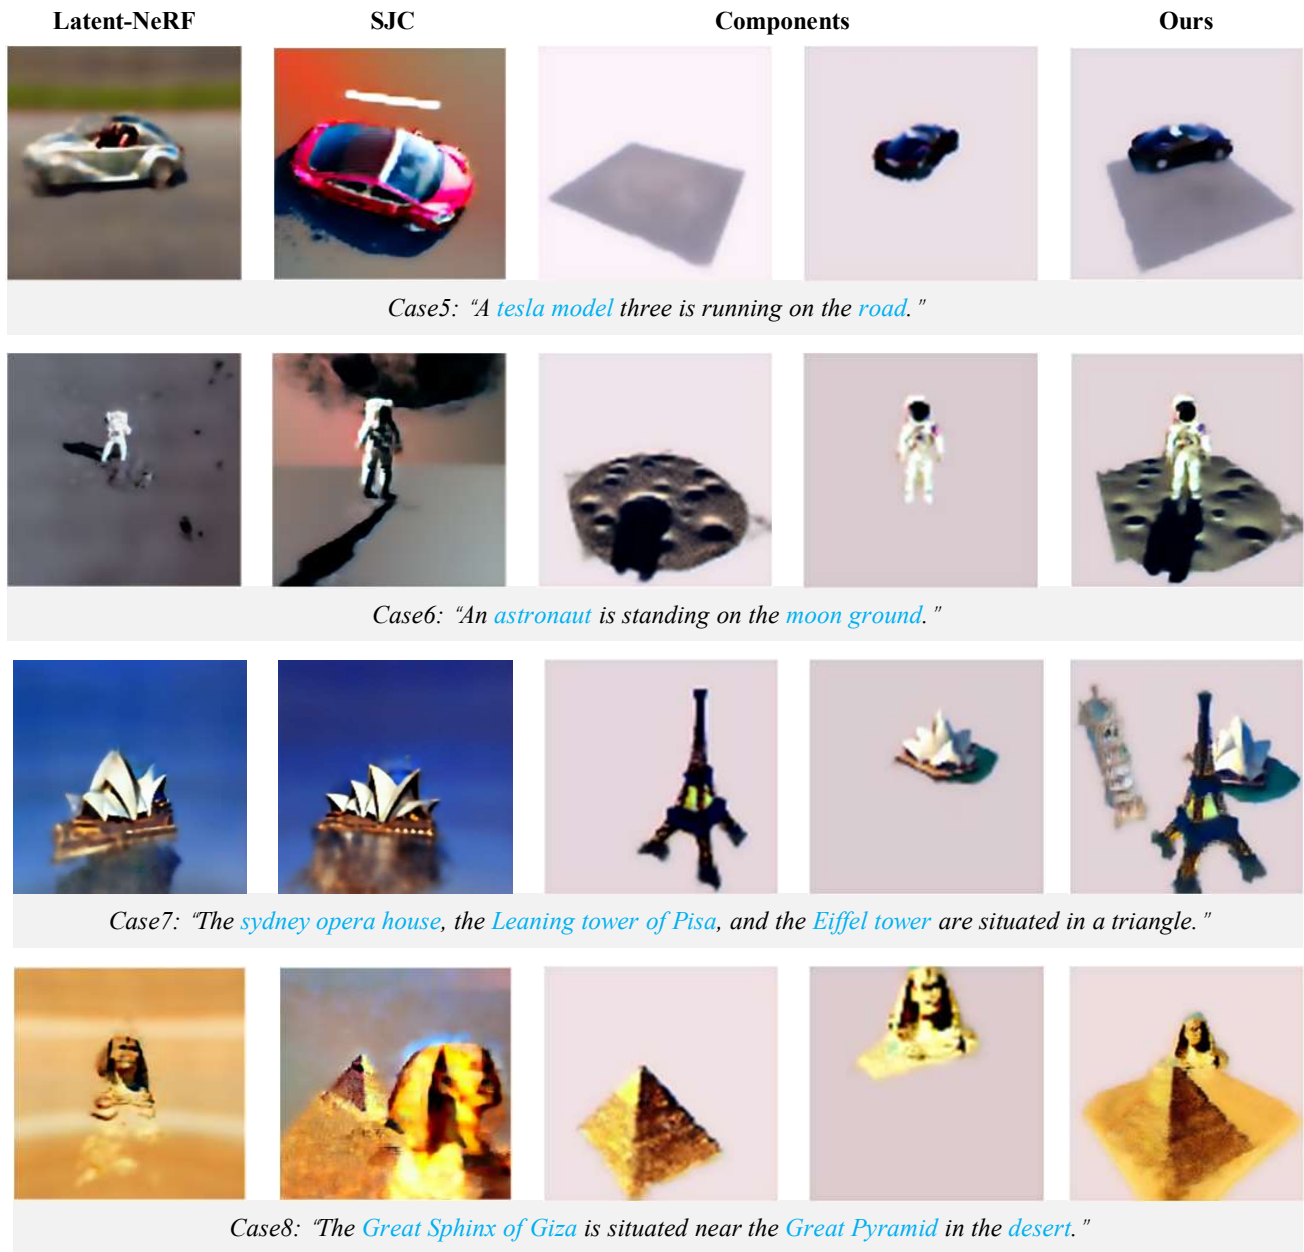


Figure 22: More qualitative results using multi-object text prompts.



Figure 23: More qualitative results using multi-object text prompts.

11 CASE DETAILS

1. Text prompts :

- (a) A tesla model three is running on the road.
- (b) An astronaut is standing on the moon ground.
- (c) A red apple and a yellow banana.
- (d) A white king and a black queen chess piece on a chess board.
- (e) Glasses of wine, salad and french bread on a wooden table.
- (f) The sydney opera house, the Leaning tower of Pisa, and the Eiffel tower are situated in a triangle.
- (g) The Great Sphinx of Giza is situated near the Great Pyramid in the desert.
- (h) Crystal ball, Dichroic glass ball, Murano glass ball, and Solar-powered glass ball are placed on the glass table.
- (i) A bed is next to a nightstand with a table lamp on it in a room.
- (j) A bunch of sunflowers in a barnacle encrusted clay vase.
- (k) A silvery wristwatch lay on a velvet cloth, its glossy surface reflecting the soft glow of a nearby candle, creating a dance of light and shadow.
- (l) A glass of water on a wooden table refracts the morning sunlight, casting a rainbow onto the table's surface.

2. Scene Editing:

- (a) A bed is next to a nightstand with a vase of sunflowers on it in a room.
- (b) A bed is next to nightstands with table lamps on them in a room.
- (c) On the polished surface of the mahogany nightstand, a lamp with a shade of woven silk cast a warm glow over a Murano glass ball and its companion, a shimmering Dichroic glass ball, creating a dance of colors against the darkening twilight.
- (d) Glasses of juice, salad, apple, and French bread on a wooden table.
- (e) The game of black queen and white king chess was played on a table, with a table lamp providing the necessary illumination.

3. Questions for User Study:

- (a) On a scale of 1 (low) to 7 (high), how would you rate the semantic consistency of the generated 3D assets? (Composition Correctness Evaluation task).
- (b) On a scale of 1 (low) to 7 (high), how would you rate the multi-view consistency of the generated 3D assets? (Composition Correctness Evaluation task).
- (c) On a scale of 1 (low) to 7 (high), how would you rate the quality of the generated images from the inputs? (Generative Quality Evaluation task).
- (d) What objects do you find in the 3D assets? (Object Identification task).