

DreamScene360: Unconstrained Text-to-3D Scene Generation with Panoramic Gaussian Splatting

Shijie Zhou^{1*}, Zhiwen Fan^{2*}, Dejia Xu^{2*}, Haoran Chang¹,
Pradyumna Chari¹, Tejas Bharadwaj¹, Suyu You³,
Zhangyang Wang², and Achuta Kadambi¹

¹ University of California, Los Angeles

² University of Texas at Austin

³ DEVCOM Army Research Laboratory

<http://dreamscene360.github.io/>

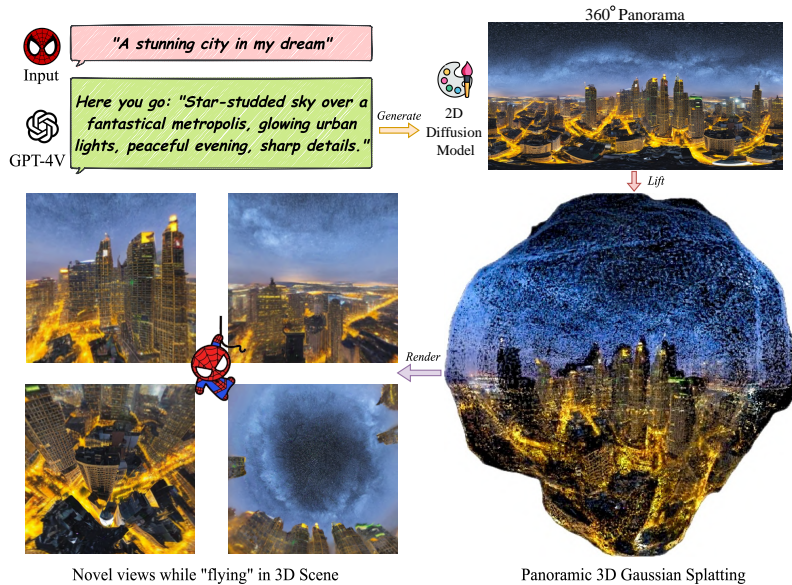


Fig. 1: DreamScene360. We introduce a 3D scene generation pipeline that creates immersive scenes with full 360° coverage from text prompts of any level of specificity.

Abstract. The increasing demand for virtual reality applications has highlighted the significance of crafting immersive 3D assets. We present a text-to-3D 360° scene generation pipeline that facilitates the creation of comprehensive 360° scenes for in-the-wild environments in a matter of minutes. Our approach utilizes the generative power of a 2D diffusion model and prompt self-refinement to create a high-quality and globally coherent panoramic image. This image acts as a preliminary “flat”

* Equal contribution.

(2D) scene representation. Subsequently, it is lifted into 3D Gaussians, employing splatting techniques to enable real-time exploration. To produce consistent 3D geometry, our pipeline constructs a spatially coherent structure by aligning the 2D monocular depth into a globally optimized point cloud. This point cloud serves as the initial state for the centroids of 3D Gaussians. In order to address invisible issues inherent in single-view inputs, we impose semantic and geometric constraints on both synthesized and input camera views as regularizations. These guide the optimization of Gaussians, aiding in the reconstruction of unseen regions. In summary, our method offers a globally consistent 3D scene within a 360° perspective, providing an enhanced immersive experience over existing techniques. Project website at: <http://dreamscene360.github.io/>.

1 Introduction

The vast potential applications of text-to-3D to VR/MR platforms, industrial design, and gaming sectors have significantly propelled research efforts aimed at developing a reliable method for immersive scene content creation at scale. Recent developments in the 2D domain have seen the successful generation or editing of high-quality and adaptable images/videos using large-scale pre-trained diffusion models [48, 51] on large-scale datasets, allowing users to generate customized content on demand.

Moving beyond 2D, the generation of 3D content, particularly 3D scenes, is constrained by the limited availability of annotated 3D image-text data pairs. Consequently, efforts in 3D content creation often rely on leveraging large-scale 2D models. This line of approach facilitates the creation of 3D scenes through a time-consuming distillation process. An example of this is DreamFusion [46], which seeks to distill the object-wise 2D priors from diffusion models into a 3D neural radiance field (NeRF) [34]. However, these approaches often suffer from low rendering quality, primarily due to the multi-view inconsistency of 2D models, and struggle to extend to scene-scale 3D structure with fine details texture creation, particularly for outdoor scenes [46] with outward-facing viewpoints and unbounded scene scale. Another avenue of 3D generation draws insights from explicit representations, such as point clouds and meshes, as demonstrated in LucidDreamer [7] and Text2Room [20]. These methods attempt to bridge the gap between 2D and 3D generation by initializing with an explicit 3D representation, and then progressively expanding the learned 3D representation to encompass a broader field-of-view. However, the progressive optimization frameworks leveraged by these methods struggle to inpaint substantial missing areas, especially when targeting 360° scenes under unconstrained conditions, resulting in notably distorted and disjointed structures. Moreover, the issue of prompt engineering in text-to-image generation [51, 52], becomes more pronounced in text-to-3D generation frameworks [1, 7, 46] that rely on either time-consuming score distillation or complex, multi-step progressive inpainting during the scene generation process, leading to a considerable trial-and-error effort to achieve the desired 3D scene.

To address the above challenges in creating a holistic 360° text-to-3D scene generation pipeline, we introduce **DreamScene360**. Our method initially leverages the generative capabilities of text-to-panorama diffusion models [62] to produce omnidirectional 360° panoramas providing a comprehensive representation of the scene. A self-refining mechanism is used to enhance the image to alleviate prompt engineering, where GPT-4V is integrated to improve the visual quality and the text-image alignment through iterative quality assessment and prompt revision. While the generated panorama images overcome the view consistency issue across different viewpoints, they still lack depth information and any layout priors in unconstrained settings, and contains partial observations due to their single-view nature. To address this, our approach involves initializing scale-consistent scene geometry by employing a pretrained monocular depth estimator alongside an optimizable geometric field, facilitating deformable alignment for each perspective-projected pixel. The gaps, stemming from single-view observations, can be filled by deforming the Gaussians to the unseen regions by creating a set of pseudo-views with a synthesized multi-view effect and the distillation of pseudo geometric and semantic constraints from 2D models (DPT [49] and DINOv2 [43]) to guide the deformation process to alleviate artifacts.

Collectively, our framework, DreamScene360, enables the creation of immersive and realistic 3D environments from a simple user command, offering a novel solution to the pressing demand for high-quality 3D scenes (see the workflow in Fig. 1). Our work also paves the way for more accessible and user-friendly 3D scene generation by reducing the reliance on extensive manual effort.

2 Related Works

2D Assets Generation. The generation of 2D assets allows for incredible creative liberty, and the use of large-scale learning based priors for content generation. Generative Adversarial Networks [15] were the original state of the art for image generation. Variants such as StyleGAN [24] showed the ability for fine-grained control of attributes such as expression through manipulation of the latent representations as well. More recently, denoising diffusion models [9, 18, 56] have been the new state of the art for generative models. Text-guided image diffusion models [53] have shown the capability of generating high-quality images that are faithful to the conditioning text prompts. Subsequent work has made the generation process more efficient by performing denoising in the latent space [51], and by speeding up the denoising process [14, 33, 54]. Techniques such as classifier free guidance [19] have significantly improved faithfulness to text prompts. Additional control of generation has also been shown to be possible, through auxiliary inputs such as layout [74], pose [73] and depth maps [5]. More recently, text-to-image diffusion models have been finetuned to generate structured images, such as panoramas [62]. In this work, we utilize text-to-panorama generation as structured guidance for text to 360° scene generation.

Text-to-3D Scene Generation. In recent times, text-to-3D scene generation has been reinvented with the advent of 2D diffusion models and guidance tech-

niques. Text-to-image diffusion models, along with techniques such as classifier free guidance, have been found to provide strong priors to guide 3D generation methods [46, 57, 63, 64, 68]. These methods are generally more frequently used than direct text-to-3D diffusion models [22, 23]. More recent works [13, 30, 60] successfully generate multi-object compositional 3D scenes. On the other hand, a second class of methods use auxiliary inputs such as layouts [45]. The least constrained text-to-3D methods do not rely on any auxiliary input with the only input being the text prompt describing the 3D scene [7, 11, 20, 32, 44, 70, 72]. Our work requires a text prompt input; however, unlike prior work, we propose using panoramic images as an intermediate input for globally consistent scenes.

Efficient 3D Scene Representation. 3D scene representations include a wide variety of techniques, including point clouds [4], volumetric representations [31, 41], and meshes [71]. More recently, however, learning-based scene representations have gained prominence. Implicit representations such as neural radiance fields [35] have shown high quality rendering and novel view synthesis capability. Additional learning-based representations, that span both explicit [12] and mixed [39] have been proposed to speed up learning. Most recently, 3D Gaussian splatting [25] has enabled fast learning and rendering of high quality radiance fields using a Gaussian kernel-based explicit 3D representation. Since then, several works have emerged enabling sparse view [66, 76] and compressed [10, 27, 38, 40, 42] 3D Gaussian representations, as well as representation of multidimensional feature fields [75]. A variation involves learning a Gaussian-based radiance field representation using panoramic images as input [2]. In this work, we propose a method for text to 360° 3D scene generation, by using panorama images as an intermediate representation.

3 Methods

In this section, we detail the proposed DreamScene360 architecture (Fig. 2). DreamScene360 initially generates a 360° panorama utilizing a self-refinement process, which ensures robust generation and aligns the image with the text semantically (Sec. 3.1). The transformation from a flat 2D image to a 3D model begins with the initialization from a panoramic geometric field (Sec. 3.2), followed by employing semantic alignment and geometric correspondences as regularizations for the deformation of the 3D Gaussians (Sec. 3.3).

3.1 Text to 360° Panoramas with Self-Refinement

Panoramic images provide an overview of the entire scene in one image, which is essential for generating 360° 3D scenes with global consistency.

360° Panoramic Image Generation. A crucial requirement for the generated panorama is ensuring continuity between the leftmost and rightmost sides of the image. We utilize the diffusion process from MultiDiffuser [3] to generate a

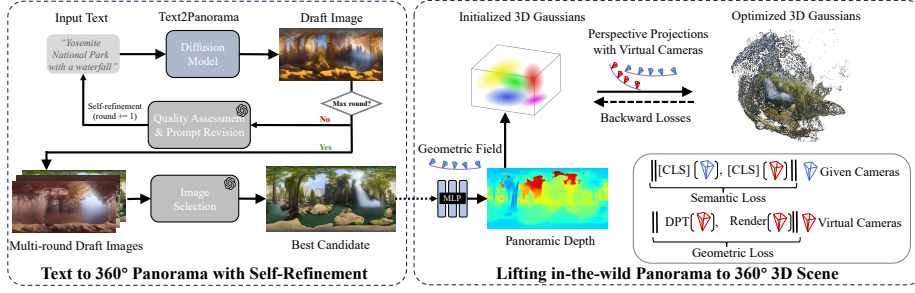


Fig. 2: Overall Architecture. Beginning with a concise text prompt, we employ a diffusion model to generate a 360° panoramic image. A self-refinement process is employed to produce the optimal 2D candidate panorama. Subsequently, a 3D geometric field is utilized to initialize the Panoramic 3D Gaussians. Throughout this process, both semantic and geometric correspondences are employed as guiding principles for the optimization of the Gaussians, aiming to address and fill the gaps resulting from the single-view input.

panoramic image I_0 of resolution $H \times 2H$ based on a pre-trained diffusion model Φ . Starting from a noisy image I_T , we iteratively denoise the image by solving an optimization problem for each of several patches of the image, selected via a sliding window mechanism. For each patch $P_i(I_t)$, we ensure the distance against their denoised version $\Phi(P_i(I_t))$ is minimized. Though each patch may pull the denoising process in a different direction, the fused result is the weighted average of the update of each sample, whose closed form is written as follows,

$$\Phi(I_{t-1}) = \sum_{i=1}^n \frac{P_i^{-1}(W_i)}{\sum_{j=1}^n P_j^{-1}(W_j)} \otimes P_i^{-1}(\Phi(P_i(I_t))), \quad (1)$$

where W_i refers to per pixel weight, set to 1 in our experiments.

We use StitchDiffusion [62] as the pretrained 2D diffusion model, where a stitch method is employed in the generation process for synthesizing seamless 360° panoramic images. Trained on a curated paired image-text dataset containing 360° panoramas, a customized LoRA [21] module is incorporated into the MultiDiffusion [3] process. At each denoising timestamp, we not only diffuse at the original resolution $H \times (2H + 2W)$ via MultiDiffuser Φ , but also stitch the leftmost and rightmost image regions ($H \times W$ each) and diffuse the concatenated patch as well to ensure consistency at the border regions. Finally, we consider the center cropped region of $H \times 2H$ as the final 360° panoramic image. It is worth noting that Dreamscene360 is versatile in practice and can also adapt to other text-to-panorama diffusion models.

Multi-Round Self-Refinement. Unlike previous works that generate 3D scenes through a time-consuming score distillation process [64] or progressive inpainting [7], our work uses panorama to achieve user-friendly “one-click” 3D scene generation. We integrate GPT-4V to facilitate iterative self-refinement during

the generation process, a feature that was challenging to implement in previous baselines due to the absence of global 2D image representations. Here, we draw inspiration from perspective image generation, Idea2Img [69], and implement a self-improvement framework aiming for better text-image aligned panoramic image generation. Starting from a user-provided rough prompt, we leverage GPT-4V to provide feedback and prompt revision suggestions according to the results generated by the previous step. During each round, GPT-4V judges the generated image quality in terms of object counts, attributes, entities, relationships, sizes, appearance, and overall similarity with the original user-specified prompt. A score from 0-10 is assigned to each draft image and the one image with highest score is provided to GPT-4V for additional improvement. GPT-4V will then produce an improved text prompt based on the issues observed in the current generation results, and the new prompt will be used for another round of panorama generation. After a number of rounds, we collect the image with the highest visual quality score judged by GPT-4V in the whole generation process as our final panorama image.

Our 3D scene generation framework, therefore, enjoys a user-friendly self-improvement process, without the need for troublesome prompt engineering for users as in previous methods [7, 46, 64], but is able to obtain a high-quality, visually pleasing, text-aligned, and 360° consistent panoramic images that can be later converted into an immersive 3D scene via Panoramic Gaussian Splatting, which is detailed in the next sections.

3.2 Lifting in-the-wild Panorama to 360 Scene

Transforming a single image, specifically an in-the-wild 360° panoramic image, into a 3D model poses significant challenges due to inadequate observational data to regularize the optimization process, such as those required in 3D Gaussian Splatting (see Fig. 7). Rather than beginning with a sparse point cloud (3DGS), we initialize with a dense point cloud utilizing pixel-wise depth information from the panoramic image of resolution $H \times W$, which are then refined towards more precise spatial configurations, ensuring the creation of globally consistent 3D representations, robust to viewpoint changes.

Monocular Geometric Initialization. Given a single panoramic image \mathbf{P} , we project it onto N perspective tangent images with overlaps $\{(\mathbf{I}_i \in \mathbb{R}^{H \times W \times 3}, \mathbf{P}_i \in \mathbb{R}^{3 \times 4})\}_{i=1}^N$, following the literature’s suggestion that 20 tangent images adequately cover the sphere’s surface as projected by an icosahedron [50]. Unlike indoor panoramas, which benefit from structural layout priors in optimization, we employ a monocular depth estimator, DPT [49], to generate a monocular depth map $\mathbf{D}_i^{\text{Mono}}$ providing a robust geometric relationship. Nevertheless, these estimators still possess affine ambiguity, lacking a known scale and shift relative to metric depth. Addressing this efficiently is crucial for precise geometric initialization.

Global Structure Alignment. Previous studies in deformable depth alignment [16, 50, 61] and pose-free novel view synthesis [6] have explored aligning scales across multiple view depth maps derived from monocular depth estimations. In this context, we utilize a learnable global geometric field (MLPs), inspired by [50, 61], supplemented by per-view scale and per-pixel shift parameters: $\{(\alpha_i \in \mathbb{R}, \beta_i \in \mathbb{R}^{H \times W})\}_{i=1}^N$. We define the view direction (\mathbf{v}) for all pixels on the perspective images deterministically. The parameters of MLPs Θ are initialized with an input dimension of three and an output dimension of one. The parameters $\{(\alpha_i \in \mathbb{R})\}_{i=1}^N$ are initialized to ones, and $\{(\beta_i \in \mathbb{R}^{H \times W})\}_{i=1}^N$ to zeros. With these optimizable parameters, we define our optimization goal as follows:

$$\min_{\alpha, \beta, \Theta} \left\{ \|\alpha \cdot \mathbf{D}^{\text{Mono}} + \beta - \text{MLPs}(\mathbf{v}; \Theta)\|_2^2 + \lambda_{\text{TV}} \mathcal{L}_{\text{TV}}(\beta) + \lambda_{\alpha} \|\gamma(\alpha) - 1\|^2 \right\} \quad (2)$$

$$\text{where } \mathcal{L}_{\text{TV}}(\beta) = \sum_{i,j} ((\beta_{i,j+1} - \beta_{i,j})^2 + (\beta_{i+1,j} - \beta_{i,j})^2)$$

Here, λ_{TV} and λ_{α} are regularization coefficients to balance the loss weight, $\gamma(\cdot)$ is the softplus function. We set FoV as 80° during the optimization and use the predicted depth from the MLPs for the subsequent Gaussian optimization, as it can provide the depth of any view direction.

3.3 Optimizing Monocular Panoramic 3D Gaussians

While 3D Gaussians initialized with geometric priors from monocular depth maps provide a foundational structure, they are inherently limited by the lack of parallax inherent to single-view panoramas. This absence of parallax — critical for depth perception through binocular disparity — along with the lack of multiple observational cues typically provided by a baseline, poses substantial challenges in accurately determining spatial relationships and depth consistency. Thus, it is imperative to employ more sophisticated and efficient generative priors that function well in unconstrained settings and enable the scaling up and extend to any text prompts.

3D Gaussian Splatting. 3D Gaussian Splatting (3DGS) utilizes the multi-view calibrated images using Structure-from-Motion [55], and optimizes a set of Gaussians with center $\mathbf{x} \in \mathbb{R}^3$, an opacity value $\alpha \in \mathbb{R}$, spherical harmonics (SH) coefficients $\mathbf{c} \in \mathbb{R}^C$, a scaling vector $\mathbf{s} \in \mathbb{R}^3$ and a rotation vector $\mathbf{q} \in \mathbb{R}^4$ represented by a quaternion. Upon projecting the 3D Gaussians into a 2D space, the color C of a pixel is computed by volumetric rendering, which is performed using front-to-back depth order [26]:

$$C = \sum_{i \in \mathcal{N}} c_i \alpha_i T_i, \quad (3)$$

where $T_i = \prod_{j=1}^{i-1} (1 - \alpha_j)$, \mathcal{N} is the set of sorted Gaussians overlapping with the given pixel, T_i is the transmittance, defined as the product of opacity values of previous Gaussians overlapping the same pixel.

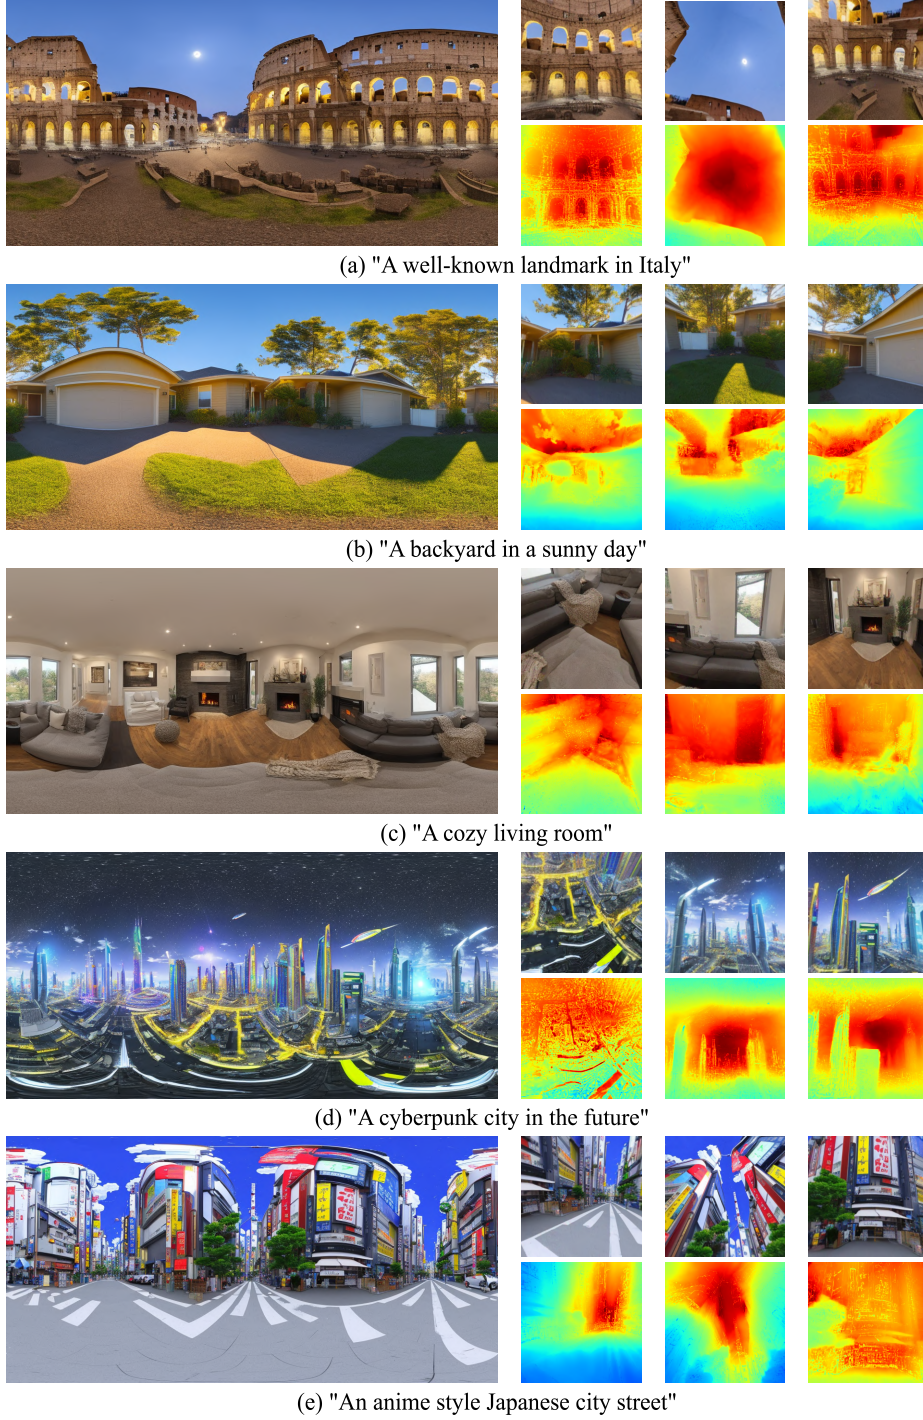


Fig. 3: Diverse Generation. We demonstrate that our generated 3D scenes are diverse in style, consistent in geometry, and highly matched with the simple text inputs.

Synthesize Parallax with Virtual Cameras. We emulate parallax by synthesizing virtual cameras that are unseen in training but are close to the input panoramic viewpoint. We methodically create these cameras to simulate larger movements. This procedure is quantitatively described by incrementally introducing perturbation to the panoramic viewpoint coordinates (x, y, z) formalized as:

$$(x', y', z') = (x, y, z) + \delta(d_x, d_y, d_z) \quad (4)$$

where (x', y', z') denotes the new virtual camera positions, and $\delta(d_x, d_y, d_z)$ denotes the progressively increasing perturbations in each coordinate direction, under uniform distribution over $[-0.05, +0.05] \times \gamma$, where $\gamma \in \{1, 2, 4\}$ stands for 3-stage progressive perturbations, emulating the camera’s movement from the original point.

Distilling Semantic Similarities. Previous research has underscored the significance of capturing the appearance of objects/scenes to establish robust feature correspondences across different views [17, 29], which is crucial for tasks like co-segmentation and 2D pre-training. Inspired by these studies, we aim to establish a connection in visual feature correspondence between training views and synthetically generated virtual views. This is achieved by enforcing feature-level similarity, guiding the 3D Gaussians to fill the geometric gaps in the invisible regions effectively. To elaborate, we generate a perspective image I_i from the panoramic 3D Gaussians at a specific training camera viewpoint P_i and create another synthesized image I'_i using Eq 4 that is proximate to P_i . We employ the [CLS] token from the pre-trained DINOv2 [43] model to encapsulate compact semantic features through the equation:

$$\mathcal{L}_{sem} = 1 - \text{Cos}([\text{CLS}](I_i), [\text{CLS}](I'_i)), \quad (5)$$

where Cos denotes the cosine similarity. This approach is inspired by [59, 67], assuming that the [CLS] token from a self-supervised, pre-trained Vision Transformer (ViT) can capture the high-level semantic attributes of an image. Such a mechanism is instrumental in identifying and leveraging similarities between two adjacent rendered images, facilitating our end-to-end 3DGS optimization process.

Regularizing Geometric Correspondences. While leveraging appearance cues is valuable, relying solely on them can lead to spatial inconsistencies. This is because models like DINOv2 might prioritize capturing semantic features over maintaining geometric coherence, potentially resulting in artifacts such as floaters. To mitigate this, we introduce a geometric regularization strategy designed to penalize discontinuities between pixels that exhibit inaccurate depth relationships. We employ a monocular depth estimator DPT [49] for this purpose [8]. Although it may not provide globally accurate depth, it allows us to ascertain the relative spatial relationships between pixels:

$$\mathcal{L}_{geo}(I_i, D_i) = 1 - \frac{\text{Cov}(D_i, \text{DPT}(I_i))}{\sqrt{\text{Var}(D_i)\text{Var}(\text{DPT}(I_i))}}, \quad (6)$$




Fig. 4: Visual Comparisons. We showcase 360° 3D scene generation. In each row, from left to right, displays novel views as the camera undergoes clockwise rotation in yaw, accompanied by slight random rotations in pitch and random translations. LucidDreamer [7] hallucinates novel views from a conditioned image (indicated by a red bounding box) but lacks global semantic, stylized, and geometric consistency. In contrast, our method provides complete 360° coverage without any blind spots (black areas in baseline results), and shows globally consistent semantics.

Here, I_i represents the rendered image at the i -th camera, and D_i signifies the rendered depth. Additionally, we incorporate an unsupervised local smoothness prior (TV loss) on the rendered depth at virtual views.

3.4 Optimization

The entire pipeline for transforming a panorama into 3D can be supervised end-to-end through a composite of all the loss functions.

$$\mathcal{L} = \mathcal{L}_{RGB} + \lambda_1 \cdot \mathcal{L}_{sem} + \lambda_2 \cdot \mathcal{L}_{geo} \quad (7)$$

Here, \mathcal{L}_{RGB} represents the photometric loss on projected perspective images, which consists of \mathcal{L}_1 and a D-SSIM term as described in [25]. \mathcal{L}_{sem} denotes semantic regularization, and \mathcal{L}_{geo} signifies geometric regularization. The weights λ_1 and λ_2 are set to 0.05 each, respectively.

Table 1: Quantitative comparisons between LucidDreamer and ours.

	CLIP Distance↓	Q-Align↑	NIQE↓	BRISQUE↓	Runtime
LucidDreamer [7]	0.8900	3.0566	6.2305	51.9764	6min.15sec.
Ours	0.8732	3.1094	4.9165	38.3911	7min.20sec.

4 Experiments

4.1 Experiment Setting

Implementation. Given the panorama generated from text, we resize it to 1024×2048 in order to produce a dense enough point cloud by projecting per-pixel panoramic depth values into 3D space along the per-pixel camera ray directions, where the directions can be obtained according to the coordinate transformation in spherical panoramic imaging [50]. Since the panoramic depth map is predicted from the optimized geometric field, the values in local patches are guaranteed to be consistent and smooth, thereby resulting in point cloud in 3D that accurately captures the smooth geometry at the objects’ surfaces. This setup enables effective 3D Gaussians initialization within a controllable, bounded area, eliminating redundant points in empty spaces. To circumvent the common floater issue in 3DGS rendering, we disable the densification process, enhancing the overall quality and consistency of the rendered scenes.

Baseline Methods. Our work tackles the challenging problem of unconstrained 360° generation, including both indoor and outdoor scenarios. However, the works utilizing a bounded NeRF representation using score distillation do not work very well in this case. Thus, the comparisons are conducted between DreamScene360 (ours) and the state-of-the-art LucidDreamer [7]. We use the open-source codebase of LucidDreamer, which starts from a single image and a text prompt. The framework constructs a global point cloud by progressive inpainting into 360° views and then distills a set of 3D Gaussians. In our experiments, we set the input image for the LucidDreamer to be generated from Stable Diffusion [51] v1.5 using the input text for a fair comparison.

Metrics. Since there is no ground truth in the generated 3D scenes, we utilize CLIP [47] embedding distance, following previous works [58, 68], to measure the text-image alignment as a mechanism to quantify the novel view rendering quality. Specifically, for each method, we render images with camera rotations and translations to mimic the immersive trajectory inside the 3D scenes. Additionally, we utilize multiple non-reference image quality assessment metrics. NIQE [37] and BRISQUE [36] are widely adopted non-reference quality assessment methods in measuring in-the-wild image quality. QAlign [65] is the state-of-the-art method in quality assessment benchmarks, which adopts a large multi-modal model fine-tuned on available image quality assessment datasets. We utilize the “quality” mode of QAlign, which focuses on the perceptual quality of image contents.

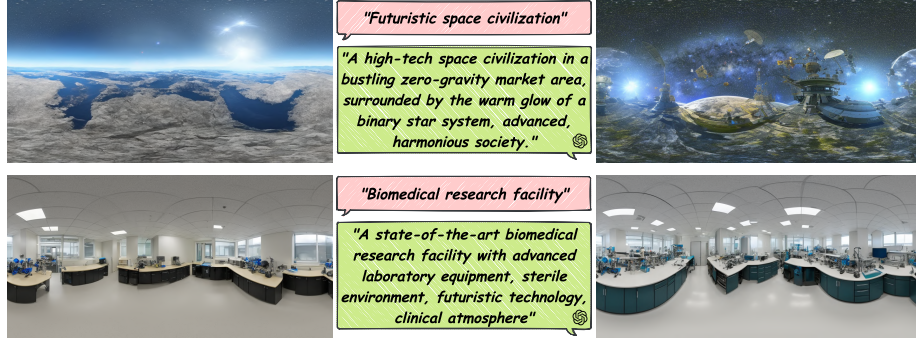


Fig. 5: Ablation of Self-Refinement. We demonstrate that the self-refinement process greatly enhances the image quality by improving the text prompt. As shown in each row, the image on the left is generated using a simple user prompt, while a prompt augmented by GPT-4V obtains the image on the right. We observe that after the multi-round self-refinement, GPT-4V selects the one panorama with better visual quality, which provides solid support for the immersive 3D scene we further generate.

4.2 Main Results

360° Scene Generation. As shown in Fig. 3, our method can generate diverse 3D scenes in different styles with distinct contents, while preserving high-fidelity novel-view rendering ability and realistic scene geometry. These results showcase our method’s generalization ability to diverse use cases, providing a user-friendly experience in realizing users’ imaginations.

Comparisons with Baseline Methods. We show visual comparisons against LucidDreamer [7] in Fig. 4. LucidDreamer involves progressive inpainting to find agreement between multiple synthesized images and tries to fuse a point cloud that can be later distilled into 3D Gaussians. Their pipeline, which inpaints each patch separately based on the same text prompt, tends to produce repetitive results especially when generating complex scenes. In comparison, our method delivers consistent results thanks to the intermediate panorama as a global 2D representation. We provide quantitative comparisons in Tab. 1, where the camera undergoes clockwise rotation in yaw, accompanied by slight random rotations in pitch and random translations to capture 4 novel views roughly representing front, back, left, and right, along with two views of up and down by 90° rotation in pitch, to mimic the exploration in an immersive 360° panoramic 3D scene.

For the *bedroom* text prompt, LucidDreamer begins with a view featuring a modern style bedroom with one bed but hallucinates the novel views into multiple bedrooms, while the style of the imagined bedrooms diverges from that of the starting view. In the case of the *Yosemite* text prompt, LucidDreamer merely replicates the waterfall seen in the initial view throughout. In conclusion, our results demonstrate global semantic, stylized, and geometric consistency, offering complete 360° coverage without any blind spots.




Fig. 6: Ablation of Optimization Loss. We demonstrate the impact of Semantic and Geometric losses on the synthesized virtual cameras. (a) Utilizing photometric loss on camera views from a rendered panorama induces artifacts when altering rendered camera positions. (b) Implementing Geometric and (c) Semantic regularizations effectively reduces the artifacts originating from invisible views. (d) Integrating both regularizations yields the most optimal outcomes.

4.3 Ablation Study

Self-refinement Process We further evaluate the importance of the self-refinement process. In this part, we mainly focus on the generated panorama since panorama provides a holistic view of the 3D scene and provides an upper bound for the visual quality. As can be seen in Fig. 5, we observe that using a simple text prompt usually delivers minimalist results with fewer details. With the help of GPT-4V, our self-refinement process enjoys prompt revision and quality assessment. This ability finally facilitates the selection of a more realistic and detailed panorama image among the draft image candidates. These functionalities are otherwise hard to achieve in previous baselines that do not have global 2D representations, and as a result, our results provide a much better visual appearance than baselines as shown in Fig. 4 and Tab. 1.

Loss Function We investigate the importance of multiple loss functions we adopted in Fig. 6. As shown in the figure, especially in the highlighted region, our full configuration of loss functions delivers the best visual quality when rendered from a novel viewpoint. In comparison, removing the geometric regularizations or semantic constraints will result in unpleasant artifacts.

Initialization We showcase the importance of proper initialization in Fig. 7. As can be seen in the image, adopting our curated point initialization largely benefits the rendering quality of our 3D Panoramic Gaussian Splatting. An alternate solution involves adopting random initialization of the point cloud locations. This option, however, generates blurry results, mainly due to the challenge of learning proper point locations on the fly without ground truth geometry as supervision. Note that since our full model incorporates this high-quality point initialization, at the training stage, we can safely disable the adaptive density control of 3D Gaussians and, therefore, speed up convergence. Using random



(a) Random initialized geometry for optimizing 3D Gaussians.



(b) Initialized from globally aligned monocular geometry.

Fig. 7: Ablation Study on 3D Initialization. We present a comparative visualization of various initialization methods for 3D Panoramic Gaussian Splatting. In the absence of geometric priors of the scene (a), the optimized 3D Gaussian rendering yields plausible results in panoramic camera views. However, artifacts become evident when the rendered camera position is altered. To address this, we employ a monocular depth approach combined with learnable alignment factors and a geometric field (b), which ensures consistent alignment across multiple views.

initialization of 3D Gaussians, on the other hand, has to rely on adaptive density control introduced in 3D Gaussian Splatting paper [25] (*e.g.* densify, clone, prune) to move the 3D Gaussians to desired locations via backpropagation.

5 Conclusion

In this work, we present DreamScene360, a novel framework that unrestrictedly generates high-quality, immersive 3D scenes with complete 360° coverage in various styles from text inputs, regardless of their specificity level. Our pipeline leverages panorama as a middle ground, which provides us with a self-refinement opportunity by asking GPT-4V for quality assessment and prompt revision. Subsequently, we construct a geometric field that initializes 3D Gaussians. With the help of unsupervised geometric and semantic loss on virtual views, we lift the 2D panorama into panoramic 3D Gaussians. Experiments demonstrate our superiority against baseline methods in terms of global consistency and visual quality. However, our generation results are still limited by the default resolution (512×1024) of the pretrained text-to-panorama diffusion model. Moving forward, we will explore generating 3D scenes at higher resolutions and even 4D dynamic scenes [28] for a better, seamless, and immersive user experience.

Acknowledgements

We thank the Visual Machines Group (VMG) at UCLA, Visual Informatics Group at UT Austin (VITA), and ARL Intelligent Perception Branch for feedback and support. This project was supported by LUCI program under the Basic Research Office and partially supported by ARL grants W911NF20-2-0158 and W911NF-21-2-0104 under the cooperative A2I2 program. Z.W. is supported by the U.S. DEVCOM Army Research Laboratory and LUCI program, and an Army Young Investigator Award. A.K. is supported by a DARPA Young Faculty Award, NSF CAREER Award IIS-2046737, and Army Young Investigator Award.

References

1. Armandpour, M., Zheng, H., Sadeghian, A., Sadeghian, A., Zhou, M.: Re-imagine the negative prompt algorithm: Transform 2d diffusion into 3d, alleviate janus problem and beyond. arXiv preprint arXiv:2304.04968 (2023) [2](#)
2. Bai, J., Huang, L., Guo, J., Gong, W., Li, Y., Guo, Y.: 360-gs: Layout-guided panoramic gaussian splatting for indoor roaming. arXiv preprint arXiv:2402.00763 (2024) [4](#)
3. Bar-Tal, O., Yariv, L., Lipman, Y., Dekel, T.: Multidiffusion: Fusing diffusion paths for controlled image generation (2023) [4](#), [5](#)
4. Berger, M., Tagliasacchi, A., Seversky, L.M., Alliez, P., Levine, J.A., Sharf, A., Silva, C.T.: State of the art in surface reconstruction from point clouds. In: 35th Annual Conference of the European Association for Computer Graphics, Eurographics 2014-State of the Art Reports. No. CONF, The Eurographics Association (2014) [4](#)
5. Bhat, S.F., Mitra, N.J., Wonka, P.: Loosecontrol: Lifting controlnet for generalized depth conditioning. arXiv preprint arXiv:2312.03079 (2023) [3](#)
6. Bian, W., Wang, Z., Li, K., Bian, J.W., Prisacariu, V.A.: Nope-nerf: Optimising neural radiance field with no pose prior. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4160–4169 (2023) [7](#)
7. Chung, J., Lee, S., Nam, H., Lee, J., Lee, K.M.: Luciddreamer: Domain-free generation of 3d gaussian splatting scenes. arXiv preprint arXiv:2311.13384 (2023) [2](#), [4](#), [5](#), [6](#), [10](#), [11](#), [12](#)
8. Deng, C., Jiang, C., Qi, C.R., Yan, X., Zhou, Y., Guibas, L., Anguelov, D., et al.: Nerdi: Single-view nerf synthesis with language-guided diffusion as general image priors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20637–20647 (2023) [9](#)
9. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021) [3](#)
10. Fan, Z., Wang, K., Wen, K., Zhu, Z., Xu, D., Wang, Z.: Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. arXiv preprint arXiv:2311.17245 (2023) [4](#)
11. Fang, C., Hu, X., Luo, K., Tan, P.: Ctrl-room: Controllable text-to-3d room meshes generation with layout constraints. arXiv preprint arXiv:2310.03602 (2023) [4](#)
12. Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5501–5510 (2022) [4](#)

13. Gao, G., Liu, W., Chen, A., Geiger, A., Schölkopf, B.: Graphdreamer: Compositional 3d scene synthesis from scene graphs. *arXiv preprint arXiv:2312.00093* (2023) [4](#)
14. Geng, Z., Pokle, A., Kolter, J.Z.: One-step diffusion distillation via deep equilibrium models. *Advances in Neural Information Processing Systems* **36** (2024) [3](#)
15. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. *Communications of the ACM* **63**(11), 139–144 (2020) [3](#)
16. Hedman, P., Kopf, J.: Instant 3d photography. *ACM Transactions on Graphics (TOG)* **37**(4), 1–12 (2018) [7](#)
17. Hénaff, O.J., Koppula, S., Shelhamer, E., Zoran, D., Jaegle, A., Zisserman, A., Carreira, J., Arandjelović, R.: Object discovery and representation networks. *arXiv preprint arXiv:2203.08777* (2022) [9](#)
18. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020) [3](#)
19. Ho, J., Salimans, T.: Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598* (2022) [3](#)
20. Höllein, L., Cao, A., Owens, A., Johnson, J., Nießner, M.: Text2room: Extracting textured 3d meshes from 2d text-to-image models. *arXiv preprint arXiv:2303.11989* (2023) [2](#), [4](#)
21. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685* (2021) [5](#)
22. Jun, H., Nichol, A.: Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463* (2023) [4](#)
23. Karnewar, A., Vedaldi, A., Novotny, D., Mitra, N.J.: Holodiffusion: Training a 3d diffusion model using 2d images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18423–18433 (2023) [4](#)
24. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4401–4410 (2019) [3](#)
25. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)* **42**(4), 1–14 (2023) [4](#), [10](#), [14](#)
26. Kopanas, G., Philip, J., Leimkühler, T., Drettakis, G.: Point-based neural rendering with per-view optimization. In: *Computer Graphics Forum*. vol. 40, pp. 29–43. Wiley Online Library (2021) [7](#)
27. Lee, J.C., Rho, D., Sun, X., Ko, J.H., Park, E.: Compact 3d gaussian representation for radiance field. *arXiv preprint arXiv:2311.13681* (2023) [4](#)
28. Li, R., Pan, P., Yang, B., Xu, D., Zhou, S., Zhang, X., Li, Z., Kadambi, A., Wang, Z., Fan, Z.: 4k4dgen: Panoramic 4d generation at 4k resolution. *arXiv preprint arXiv:2406.13527* (2024) [14](#)
29. Li, W., Hosseini Jafari, O., Rother, C.: Deep object co-segmentation. In: *Asian Conference on Computer Vision*. pp. 638–653. Springer (2018) [9](#)
30. Lin, Y., Bai, H., Li, S., Lu, H., Lin, X., Xiong, H., Wang, L.: Componerf: Text-guided multi-object compositional nerf with editable 3d scene layout. *arXiv preprint arXiv:2303.13843* (2023) [4](#)
31. Lombardi, S., Simon, T., Saragih, J., Schwartz, G., Lehrmann, A., Sheikh, Y.: Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751* (2019) [4](#)

32. Mao, W., Cao, Y.P., Liu, J.W., Xu, Z., Shou, M.Z.: Showroom3d: Text to high-quality 3d room generation using 3d priors. arXiv preprint arXiv:2312.13324 (2023) [4](#)
33. Meng, C., Rombach, R., Gao, R., Kingma, D., Ermon, S., Ho, J., Salimans, T.: On distillation of guided diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14297–14306 (2023) [3](#)
34. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European conference on computer vision. pp. 405–421. Springer (2020) [2](#)
35. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM **65**(1), 99–106 (2021) [4](#)
36. Mittal, A., Moorthy, A.K., Bovik, A.C.: No-reference image quality assessment in the spatial domain. IEEE Transactions on image processing **21**(12), 4695–4708 (2012) [11](#)
37. Mittal, A., Soundararajan, R., Bovik, A.C.: Making a “completely blind” image quality analyzer. IEEE Signal processing letters **20**(3), 209–212 (2012) [11](#)
38. Morgenstern, W., Barthel, F., Hilsmann, A., Eisert, P.: Compact 3d scene representation via self-organizing gaussian grids. arXiv preprint arXiv:2312.13299 (2023) [4](#)
39. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. arXiv preprint arXiv:2201.05989 (2022) [4](#)
40. Navaneet, K., Meibodi, K.P., Koohpayegani, S.A., Pirsavash, H.: Compact3d: Compressing gaussian splat radiance field models with vector quantization. arXiv preprint arXiv:2311.18159 (2023) [4](#)
41. Nguyen-Phuoc, T., Li, C., Theis, L., Richardt, C., Yang, Y.L.: Hologan: Unsupervised learning of 3d representations from natural images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7588–7597 (2019) [4](#)
42. Niedermayr, S., Stumpfegger, J., Westermann, R.: Compressed 3d gaussian splatting for accelerated novel view synthesis. arXiv preprint arXiv:2401.02436 (2023) [4](#)
43. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023) [3](#), [9](#)
44. Ouyang, H., Heal, K., Lombardi, S., Sun, T.: Text2immersion: Generative immersive scene with 3d gaussians. arXiv preprint arXiv:2312.09242 (2023) [4](#)
45. Po, R., Wetzstein, G.: Compositional 3d scene generation using locally conditioned diffusion. arXiv preprint arXiv:2303.12218 (2023) [4](#)
46. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988 (2022) [2](#), [4](#), [6](#)
47. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021) [11](#)
48. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 **1**(2), 3 (2022) [2](#)
49. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12179–12188 (2021) [3](#), [6](#), [9](#)

50. Rey-Area, M., Yuan, M., Richardt, C.: 360monodepth: High-resolution 360deg monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3762–3772 (2022) 6, 7, 11
51. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022) 2, 3, 11
52. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S.K.S., Gontijo-Lopes, R., Ayan, B.K., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M.: Photorealistic text-to-image diffusion models with deep language understanding. In: Oh, A.H., Agarwal, A., Belgrave, D., Cho, K. (eds.) Advances in Neural Information Processing Systems (2022), <https://openreview.net/forum?id=08Ykn512A1> 2
53. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. Advances in Neural Information Processing Systems 35, 36479–36494 (2022) 3
54. Salimans, T., Ho, J.: Progressive distillation for fast sampling of diffusion models. arXiv preprint arXiv:2202.00512 (2022) 3
55. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4104–4113 (2016) 7
56. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020) 3
57. Song, L., Cao, L., Xu, H., Kang, K., Tang, F., Yuan, J., Zhao, Y.: Roomdreamer: Text-driven 3d indoor scene synthesis with coherent geometry and texture. arXiv preprint arXiv:2305.11337 (2023) 4
58. Tang, J., Ren, J., Zhou, H., Liu, Z., Zeng, G.: Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. arXiv preprint arXiv:2309.16653 (2023) 11
59. Tumanyan, N., Bar-Tal, O., Bagon, S., Dekel, T.: Splicing vit features for semantic appearance transfer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10748–10757 (2022) 9
60. Vilesov, A., Chari, P., Kadambi, A.: Cg3d: Compositional generation for text-to-3d via gaussian splatting. arXiv preprint arXiv:2311.17907 (2023) 4
61. Wang, G., Wang, P., Chen, Z., Wang, W., Loy, C.C., Liu, Z.: Perf: Panoramic neural radiance field from a single panorama. arXiv preprint arXiv:2310.16831 (2023) 7
62. Wang, H., Xiang, X., Fan, Y., Xue, J.H.: Customizing 360-degree panoramas through text-to-image diffusion models. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 4933–4943 (2024) 3, 5
63. Wang, H., Du, X., Li, J., Yeh, R.A., Shakhnarovich, G.: Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12619–12629 (2023) 4
64. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. arXiv preprint arXiv:2305.16213 (2023) 4, 5, 6
65. Wu, H., Zhang, Z., Zhang, W., Chen, C., Liao, L., Li, C., Gao, Y., Wang, A., Zhang, E., Sun, W., et al.: Q-align: Teaching lmms for visual scoring via discrete text-defined levels. arXiv preprint arXiv:2312.17090 (2023) 11

66. Xiong, H., Muttukuru, S., Upadhyay, R., Chari, P., Kadambi, A.: SparseSegs: Real-time 360 $\{\backslash\deg\}$ sparse view synthesis using gaussian splatting. arXiv preprint arXiv:2312.00206 (2023) [4](#)
67. Xu, D., Jiang, Y., Wang, P., Fan, Z., Shi, H., Wang, Z.: Sinnerf: Training neural radiance fields on complex scenes from a single image. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII. pp. 736–753. Springer (2022) [9](#)
68. Xu, D., Jiang, Y., Wang, P., Fan, Z., Wang, Y., Wang, Z.: Neurallift-360: Lifting an in-the-wild 2d photo to a 3d object with 360deg views. arXiv preprint arXiv:2211.16431 (2022) [4](#), [11](#)
69. Yang, Z., Wang, J., Li, L., Lin, K., Lin, C.C., Liu, Z., Wang, L.: Idea2img: Iterative self-refinement with gpt-4v (ision) for automatic image design and generation. arXiv preprint arXiv:2310.08541 (2023) [6](#)
70. Yu, H.X., Duan, H., Hur, J., Sargent, K., Rubinstein, M., Freeman, W.T., Cole, F., Sun, D., Snavely, N., Wu, J., et al.: Wonderjourney: Going from anywhere to everywhere. arXiv preprint arXiv:2312.03884 (2023) [4](#)
71. Zhang, C., Chen, T.: Efficient feature extraction for 2d/3d objects in mesh representation. In: Proceedings 2001 International Conference on Image Processing (Cat. No. 01CH37205). vol. 3, pp. 935–938. IEEE (2001) [4](#)
72. Zhang, J., Li, X., Wan, Z., Wang, C., Liao, J.: Text2nerf: Text-driven 3d scene generation with neural radiance fields. IEEE Transactions on Visualization and Computer Graphics (2024) [4](#)
73. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023) [3](#)
74. Zheng, G., Zhou, X., Li, X., Qi, Z., Shan, Y., Li, X.: Layoutdiffusion: Controllable diffusion model for layout-to-image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22490–22499 (2023) [3](#)
75. Zhou, S., Chang, H., Jiang, S., Fan, Z., Zhu, Z., Xu, D., Chari, P., You, S., Wang, Z., Kadambi, A.: Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21676–21685 (2024) [4](#)
76. Zhu, Z., Fan, Z., Jiang, Y., Wang, Z.: Fsgs: Real-time few-shot view synthesis using gaussian splatting. arXiv preprint arXiv:2312.00451 (2023) [4](#)