

LayerPano3D: Layered 3D Panorama for Hyper-Immersive Scene Generation

SHUAI YANG*, Shanghai Jiao Tong University, China and Shanghai Artificial Intelligence Laboratory, China
JING TAN*, The Chinese University of Hong Kong, Hong Kong
MENGCHEN ZHANG, Zhejiang University, China and Shanghai Artificial Intelligence Laboratory, China
TONG WU†, The Chinese University of Hong Kong, Hong Kong
GORDON WETZSTEIN, Stanford University, USA
ZIWEI LIU, Nanyang Technological University, Singapore
DAHUA LIN†, The Chinese University of Hong Kong, Hong Kong




Fig. 1. **Overview of LAYERPANO3D.** Guided by simple text prompts, LAYERPANO3D leverages multi-layered 3D panorama to create hyper-immersive panoramic scene with $360^\circ \times 180^\circ$ coverage, enabling free 3D exploration among complex scene hierarchies.

3D immersive scene generation is a challenging yet critical task in computer vision and graphics. A desired virtual 3D scene should 1) exhibit omnidirectional view consistency, and 2) allow for large-range exploration in complex scene hierarchies. Existing methods either rely on successive scene expansion via inpainting or employ panorama representation to represent large FOV scene environments. However, the generated scene suffers from semantic drift during expansion and is unable to handle occlusion among scene hierarchies. To tackle these challenges, we introduce **LAYERPANO3D**, a novel framework for full-view, explorable panoramic 3D scene generation from a single text prompt. Our key insight is to decompose a reference 2D panorama into multiple layers at different depth levels, where each layer

*Equal contribution.

†Corresponding Authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGGRAPH Conference Papers '25, Vancouver, BC, Canada

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1540-2/25/08

<https://doi.org/10.1145/3721238.3730643>

reveals the unseen space from the reference views via diffusion prior. LAYERPANO3D comprises multiple dedicated designs: 1) We introduce a new panorama dataset UPRIGHT360, comprising 9k high-quality and upright panorama images, and finetune the advanced Flux model on UPRIGHT360 for high-quality, upright and consistent panorama generation related tasks. 2) We pioneer the Layered 3D Panorama as underlying representation to manage complex scene hierarchies and lift it into 3D Gaussians to splat detailed 360-degree omnidirectional scenes with unconstrained viewing paths. Extensive experiments demonstrate that our framework generates state-of-the-art 3D panoramic scene in both full view consistency and immersive exploratory experience. We believe that LAYERPANO3D holds promise for advancing 3D panoramic scene creation with numerous applications. More examples please visit our webpage: ys-imtech.github.io/projects/LayerPano3D/

CCS Concepts: • **Computing methodologies** → **Computer vision; Neural networks.**

Additional Key Words and Phrases: 3D Scene Generation, Panorama Generation, Diffusion Models, Neural Rendering

ACM Reference Format:

Shuai Yang, Jing Tan, Mengchen Zhang, Tong Wu, Gordon Wetzstein, Ziwei Liu, and Dahua Lin. 2025. LayerPano3D: Layered 3D Panorama for Hyper-Immersive Scene Generation. In *Special Interest Group on Computer Graphics*

and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25), August 10–14, 2025, Vancouver, BC, Canada. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3721238.3730643>

1 Introduction

The development of spatial computing, including virtual and mixed reality systems, greatly enhances user engagement across various applications, and drives demand for explorable, high-quality 3D environments. We contend that a desired virtual 3D scene should 1) exhibit high-quality and consistency in appearance and geometry across the full $360^\circ \times 180^\circ$ view; 2) allow for exploration among complex scene hierarchies with clear parallax. In recent years, many approaches in 3D scene generation [Gao et al. 2024; Li et al. 2024; Zhang et al. 2023b] were proposed to address these needs.

One branch of works [Chung et al. 2023; Fridman et al. 2023; Höllein et al. 2023; Ouyang et al. 2023; Yu et al. 2023; Zhang et al. 2024a] seeks to create extensive scenes by leveraging a “navigate-and-imagine” strategy, which successively applies novel-view rendering and outpaints unseen areas to expand the scene. However, this type of approaches suffer from the semantic drift issue: long sequential scene expansion easily produces incoherent results as the out-paint artifacts accumulate through iterations, hampering the global consistency and harmony of the generated scene.

Another branch of methods [Chen et al. 2022; Stan et al. 2023; Tang et al. 2023b; Wang et al. 2022, 2023b; Zhang et al. 2024b] employs Equirectangular Panorama to represent 360° , large field of view (FOV) environments in 2D. However, the absence of large-scale panoramic datasets hinders the capability of panorama generation systems, resulting in low-resolution images with simple structures and sparse assets. Moreover, 2D panorama [Tang et al. 2023b; Wang et al. 2022; Zhang et al. 2024b] does not allow for flexible scene exploration. Even when lifted to a panoramic scene [Zhou et al. 2024b], the simple spherical structure fails to provide complex scene hierarchies with clear parallax, leading to occluded spaces that cause blurry renderings, ambiguity, and gaps in the generated 3D panorama. Some methods [Zhou et al. 2024a] typically use inpainting-based disocclusion strategy to fill in the unseen spaces, but they require specific, predefined rendering paths tailored for each scene, limiting the potential for flexible exploration.

To this end, we present LAYERPANO3D, a novel framework that leverages Multi-Layered 3D Panorama for explorable, full-view consistent scene generation from text prompts. The main idea is to create a Layered 3D Panorama by first generating a reference panorama and treating it as a multi-layered composition, where each layer depicts scene content at a specific depth level. In this regard, it allows us to create complex scene hierarchies by placing occluded assets in different depth layers at full appearance.

Our contributions are two-fold. **First**, to generate high-quality and coherent $360^\circ \times 180^\circ$ panoramas, we curate a new dataset, namely UPRIGHT360, consisting of 9k high-quality, upright panorama images, and finetune the advanced Flux [Labs 2023] model with panorama LoRA on it for panorama generation and inpainting. This feed-forward pipeline prevents semantic drifts during panorama generation, while ensuring a consistent horizon level across all views. **Second**, we introduce the Layered 3D Panorama representation as a general solution to handle occlusion for different types

of scenes with complex scene hierarchies, and lift it to 3D Gaussians [Kerbl et al. 2023] to enable large-range 3D exploration. By leveraging pre-trained panoptic segmentation prior and K-Means clustering, we streamline an automatic layer construction pipeline to decompose the reference panorama into different depth layers. The unseen space at each layer is synthesized with the Flux-based inpainting pipeline.

Extensive experiments demonstrate the effectiveness of LAYERPANO3D in generating hyper-immersive layered panoramic scene from a single text prompt. LAYERPANO3D surpasses state-of-the-art methods in creating coherent, plausible, text-aligned 2D panorama and full-view consistent, explorable 3D panoramic environments. Furthermore, our framework does not require any scene-specific navigation paths, providing more user-friendly interface for non-experts. We believe that LAYERPANO3D effectively enhances the accessibility of full-view, explorable AIGC 3D environments for real-world applications.

2 Related Works

2.1 3D Scene Generation

Due to the recent success of diffusion models [Poole et al. 2022; Tan et al. 2024; Tang et al. 2023a; Wu et al. 2023], 3D scene generation has also achieved some development. Scenescape [Fridman et al. 2023] and DiffDreamer [Cai et al. 2023], for example, explore perpetual view generation through the incremental construction of 3D scenes. One major branch of work employ step-by-step inpainting from pre-defined trajectories. Text2Room [Höllein et al. 2023] creates room-scale 3D scenes based on text prompt, utilizing textured 3D meshes for scene representation. Similarly, LucidDreamer [Chung et al. 2023], WonderJourney [Yu et al. 2023] and Wonderworld [Yu et al. 2024] can generate domain-free 3D Gaussian splatting scenes from iterative inpainting. However, this line of work often suffer from the semantic drift issue, resulting in unrealistic scene from artifact accumulation and inconsistent semantics. While some other approaches [Cohen-Bar et al. 2023; Vilesov et al. 2023; Zhang et al. 2023b] endeavor to integrate objects with environments, they yield relatively low quality of comprehensive scene generation. Recently, our concurrent works, DreamScene360 [Zhou et al. 2024b] and HoloDreamer [Zhou et al. 2024a] also employ panorama as prior to construct panoramic scenes. However, they only achieve the $360^\circ \times 180^\circ$ field of view at a fixed viewpoint based on a single panorama of low-quality and simple structure, and do not support free roaming within the scene. In contrast, our framework leverages Multi-Layered 3D Panorama representation to construct high-quality, fully enclosed scenes that enable larger-range exploration paths in 3D scene.

2.2 Panorama Generation

Panorama generation methods are often based on GANs or diffusion models. Early in this field, with the different forms of deep generative neural networks, GAN-based panorama generation methods explore many paths to improve quality and diversity. Among them, Text2Light [Chen et al. 2022] focuses on HDR panoramic images by employing a text-conditioned global sampler alongside a structure-aware local sampler. However, training GANs is challenging and they encounter the issue of mode collapse. Recently, some studies

have utilized diffusion models to generate panoramas. MVDiffusion [Tang et al. 2023b] generates eight perspective views with multi-branch UNet but the resulting closed-loop panorama only captures the $360^\circ \times 90^\circ$ FOV. The image generated from MultiDiffusion [Bar-Tal et al. 2023] and Syncdiffusion [Lee et al. 2023] is more like a long-range image with wide horizontal angle as they do not integrate camera projection models. PanoDiff [Wang et al. 2023a] can generate 360° panorama from one or more unregistered Narrow Field-of-View (NFOV) images with pose estimation and controlling partial FOV LDM, while the quality and diversity of results are limited by the scarcity of panoramic image training data like most other methods [Li and Bansal 2023; Wang et al. 2023b; Wu et al. 2024]. In contrast, our model can generate Multi-Layered 3D Panorama for immersive, high-quality, and coherent scene generation from text prompts.

3 Method

The goal of our work is to create a panoramic scene guided by text prompts. This generated scene encompasses a complete $360^\circ \times 180^\circ$ field of view from various viewpoints within an extensive range in the scene, while allowing for immersive exploration along complex trajectories. LAYERPANO3D consists two stages. In **Stage I** (Sec. 3.1), we first generate the reference panorama from text prompt by finetuning flux [Labs 2023] with panorama LoRA in our UPRIGHT360 dataset. With the reference panorama, we construct our Layered 3D Panorama representation by iterative layer decomposition, completion and alignment process. In **Stage II** (Sec. 3.2), the Layered 3D Panorama is lifted to 3D Gaussians in a cascaded manner to enable large-range 3D exploration.

3.1 Multi-Layer Panorama Generation

We introduce the Layered 3D Panorama representation based on the following assumption: “an enclosed 3D scene contains a background and various assets positioned in front of it”. In this regard, using Layered 3D Panorama for 3D scene generation is a general approach to handle occlusion for various types of scene. To create a complete scene, we first generate a high-quality reference panorama from a single text prompt and decompose it into $N + 1$ layers along the depth dimension. As shown in Fig. 2, these layers, arranged from the farthest to the nearest, represent both the scene background (layer 0) and the layouts behind the observation point.

3.1.1 Reference Panorama Generation and UPRIGHT360. A good panorama captures rich scene details, creating an immersive and comprehensive view. This depth of detail fosters a stronger connection and deeper comprehension of the scene. Unlike the advancement in standard image generation, panorama generation still faces quality gaps. To ensure upright and geometrically consistent scenes, we curate a new dataset, namely UPRIGHT360, comprising high quality, upright panorama images, and finetune the Flux [Labs 2023] with LoRA on the UPRIGHT360 dataset. This lightweight training approach achieves optimal performance even with limited high-quality panorama data and can be directly extended to subsequent tasks, for example panorama inpainting for layer completion. For data curation, we first collect around 15k raw panorama images: 9684 panorama image samples are collected from Matterport3D [Chang

et al. 2017], 1824 images from the web, and 3592 synthetic panoramas generated from Blockadelabs. Based on this, we leverage GeoCalib [Veicht et al. 2024], a state-of-the-art single-perspective-image calibration method, to filter upright panoramas. Specifically, we employ the Equirectangular Projection (ERP), a mapping technique that projects a 3D sphere onto a 2D plane, to generate four perspective views from each panorama. These views are extracted at a fixed field of view (FOV) of 90° , an elevation of 0° , and four distinct azimuths ($0^\circ, 90^\circ, 180^\circ, 270^\circ$). We then use GeoCalib to calibrate the four views, computing their pitch and roll variances for filtering. Panoramas with variances exceeding 1.0 are classified as non-upright and excluded: $Var(pitch_{1:4}) > 1.0, Var(roll_{1:4}) > 1.0$, resulting in the creation of the final UPRIGHT360 dataset. This dataset comprises 9423 high-quality panoramas, rigorously filtered from the original collection. Building upon the UPRIGHT360 dataset, we fine-tune Flux to develop a panorama LoRA [Hu et al. 2021] for reference panorama generation.

3.1.2 Layer Decomposition. As shown in Fig. 2, the reference panorama is decomposed by first identifying the scene assets and then cluster these assets in different layers according to depth. First, we employ an off-the-shelf panoptic segmentation model [Jain et al. 2023] pretrained on ADE20K [Zhou et al. 2017] to automatically find all scene assets visible in the reference panorama. A good layer decomposition requires that the layer assets share a similar depth level within layers and are distant from assets in other layers. In this sense, we assign each asset a depth value and apply K-Means to cluster these masks into different groups. Given the reference panorama depth map, the depth value for each asset mask is determined by calculating the 75th percentile of the depth values within the masked region. According to the depth values, the assets are clustered into N groups from layer 0 to $N - 1$ and are merged into layer masks to guide the subsequent layer completion.

3.1.3 Layer Completion. With the layer mask, we focus on completing the unseen content caused by asset occlusion. In order to synthesize background pixels instead of creating new elements, we directly utilize the above trained panorama lora and integrate it into the Flux-Fill model to accomplish domain transfer, thereby employing it as a panoramic canvas inpainter. Specifically, at each layer, our model takes the layer mask M_l , the reference panorama, and the “empty scene, nothing” [Zhang and Agrawala 2024] prompt as input, and output coherent content at the masked area. The inpainted panorama at layer l is denoted P_l and is used as supervision to the subsequent panoramic 3D Gaussian scene optimization. Note that, we additionally apply SAM [Kirillov et al. 2023] to extend the layer mask, based on the inpainted panorama from the previous layer, to eliminate unwanted new generations from inpainting.

Moreover, to enable large-range rendering in 3D, where observers can examine scenes from varying distances, the unprocessed textures of distant assets may appear blurred as the observer approaches. Therefore, distant layers require higher resolution to preserve texture details at different viewpoints. To address this, Super Resolution (SR) module [Yang et al. 2023] is employed to enhance the resolution of the layered panorama from layer 0 (background layer) to layer N (reference panorama), achieving a $2\times$ upscale in resolution. SR

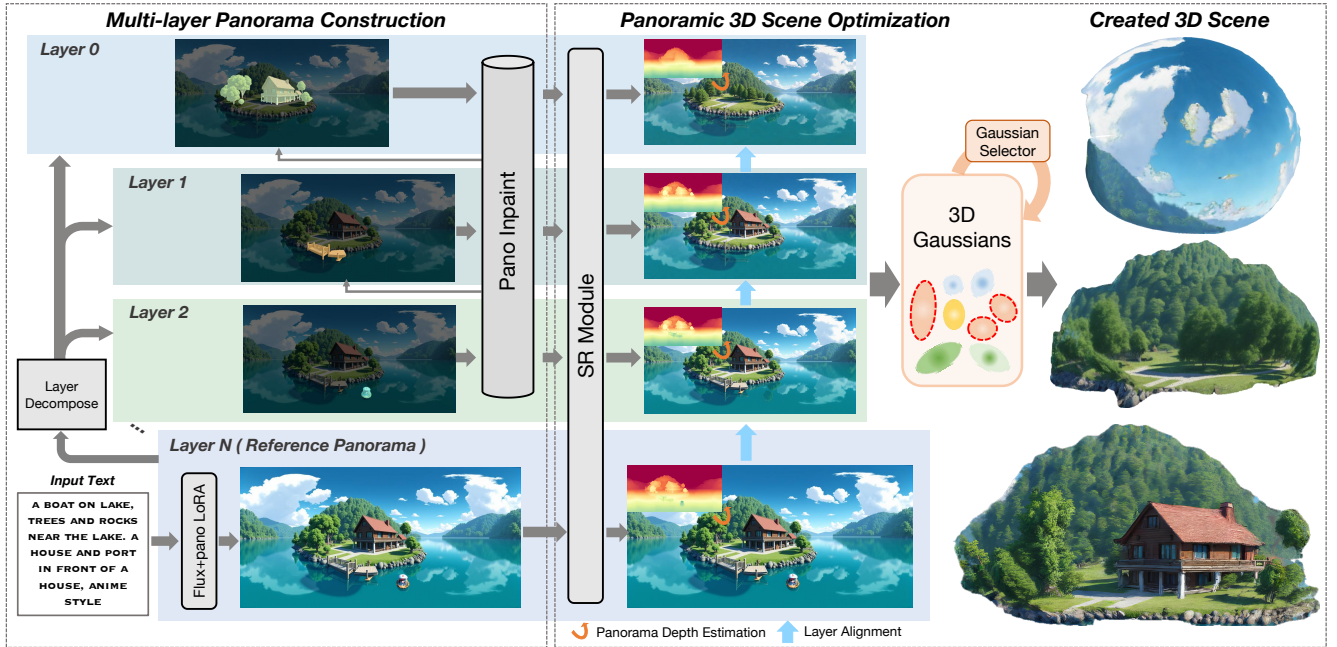


Fig. 2. **Pipeline Overview of LAYERPANO3D.** Our framework consists of two stages, namely multi-layer panorama construction and panoramic 3D scene optimization. LAYERPANO3D streamlines an automatic generation pipeline without manual efforts to design scene-specific navigation paths.

processing significantly improves the texture quality of distant objects, maintaining their visual clarity and texture details even when observed from a closer perspective.

3.1.4 Layer Alignment. Given the Layered 3D RGB Panorama $[P_l]_{l=0}^N$, we perform the depth prediction and alignment to ensure consistency in a shared space. To begin with, we apply the 360MonoDepth [Rey-Area et al. 2022], to first estimate the layer N (reference panorama) as the reference depth P_{depth}^N . Then, to align the layer depth in 3D space, we find it infeasible to simply compute a global shift and scale as in [Chung et al. 2023; Höllein et al. 2023] due to the nonlinear nature of ERP. Therefore, we leverage depth inpainting model \mathcal{F}_{depth} from [Liu et al. 2024] to directly restore depth values based on the inpainted RGB pixels. \mathcal{F}_{depth} harnesses strong generalizability from large-scale diffusion prior and synthesizes inpainted depth values at an aligned scale with the base depth. We start from reference panoramic depth P_{depth}^N to implement step-by-step restoration from layer $N - 1$ to layer 0:

$$P_{depth}^l = \mathcal{F}_{depth}(P_l, M_l \odot P_{depth}^{l+1}), \quad (1)$$

where, in layer l , the inpainted panorama P_l and masked depth map $M_l \odot P_{depth}^{l+1}$ are provided as inputs to \mathcal{F}_{depth} for restoration.

3.2 Panoramic 3D Gaussian Scene Optimization

3.2.1 3D Scene Initialization. To enable large-range 3D exploration, we lift the Layered 3D Panorama to 3D Gaussians [Kerbl et al. 2023], where the Gaussians are initialized from the layered 3D panoramic point clouds. Considering the intrinsic spherical structure of panorama, we can easily transform an equirectangular image $P \in \mathbb{R}^{H \times W \times 3}$ into 3D point cloud $S(\theta, \phi, P_{depth})$. Each pixel (u, v)




Fig. 3. **Illustration of the Gaussian Selector.** Given the new asset point cloud, the Gaussian Selector identifies the active Gaussians for next layer’s optimization.

is represented as a 3D point and the angles θ, ϕ are computed as $\theta = (2u/W - 1)\pi$, $\phi = (2v/H - 1)\pi/2$.

Then, the corresponding 3D coordinates (X, Y, Z) from the depth value $P_{depth}(\theta_u, \phi_v)$ are derived as follows:

$$\begin{aligned} X &= P_{depth}(\theta_u, \phi_v) \cos \phi_v \cos \theta_u, \\ Y &= P_{depth}(\theta_u, \phi_v) \sin \phi_v, \\ Z &= P_{depth}(\theta_u, \phi_v) \cos \phi_v \sin \theta_u. \end{aligned} \quad (2)$$

Based on this transformation, we can extract the point cloud for each layer panorama to initialize 3D Gaussians.

Drastic depth changes at layout edges introduce noisy stretched outliers that would turn into artifacts during scene refinement. Therefore, we propose an outlier removal module that specifically targets stretched point removal using heuristic point cloud filtering strategies. As stretched points are usually sparsely distributed in space, we design the point filtering strategy based on its distance from the neighbors. First, we filter out all points with the minimum distance to neighbors over threshold β_1 . Then, we eliminate points with very few neighbors. The idea is to calculate the number of neighbors of each point within a given radius and drop the points where their number of neighbors is below threshold β_2 . To speed

Table 1. **Quantitative comparison with SoTA methods on 2D Panorama Generation.** Bold indicates the best result.

Method	FID ↓	Aesthetic ↑	CLIP ↑	User Study (AUR) ↑
Text2light	286.90	4.57	18.69	1.34
Panfusion	283.80	4.78	21.22	2.38
Diffusion360	274.03	5.07	21.65	2.52
Ours	223.51	5.86	22.25	3.76

up the calculation, we map the points into 3D grids, then remove all points within grids that have less than β_2 number of neighbors.

3.2.2 3D Scene Refinement. During scene refinement, we devise two types of Gaussian training schemes for varying scene content: the *base Gaussian* for reconstructing the scene background and the *layer Gaussian* for optimizing scene layouts. Additionally, a Gaussian selector module is introduced between layer Gaussians to facilitate scene composition.

In scene refinement, the base Gaussian model is initialized on a whole of the background point cloud, and the layer Gaussian model initiates on and optimizes the foreground assets. In practice, we project the layer mask \mathcal{M}_l onto point clouds and use the masked points to initiate Gaussians. The optimized Gaussians from previous layers are frozen to avoid unwanted modification. In this way, the scene background is optimized once in the base Gaussian to reduce unnecessary computation and conflicts of Gaussians in subsequent layers.

We observe that the quality of the optimized scene is easily hampered by unaligned layers, and sometimes \mathcal{F}_{depth} does fail to produce perfectly aligned layer depths. Gaussians at layer l could span into unwanted depth levels and block assets in the subsequent layer, as illustrated in Fig. 3(a). To handle this issue, we introduce the Gaussian selector module to detect these conflicted Gaussians, reactivate them from frozen, and optimize them away from the blockage. First, the selector computes the distance vector from the camera center $\mathbf{o} = (0, 0, 0)$ to each new point \mathbf{p} , as in Fig. 3(b). The absolute distance from asset points \mathbf{p} and scene Gaussians \mathbf{g} to the camera is denoted as d_p and d_g respectively: $d_p = \|\mathbf{p} - \mathbf{o}\|_2$, $d_g = \|\mathbf{g} - \mathbf{o}\|_2$.

By examining all Gaussians that on the same ray with asset points but at a closer distance: $\mathbf{p} / d_p = \mathbf{g} / d_g$, $d_g < d_p$, we mark them as active Fig. 3(c). We also observe that despite the Gaussian center being behind the asset points, if the scaling of the Gaussian covers the point, then the Gaussian is still affecting the overall rendering quality. Hence, we update the gaussian distance as $d'_g = d'_g - \max\{s(\mathbf{g})\}$, where s is the scaling of the Gaussian \mathbf{g} . If $d_{g_1} > d_{p_1}$ and $d'_{g_1} < d_{p_1}$ as in Fig. 3(d), the referred Gaussian is also marked active. For efficient memory storage and fast look-up, we hash the distance vectors into a 3D grid. The mapping function from vector coordinates to grid indices writes: $f(\mathbf{p}) = \text{ceil}(\beta_3 \log(\mathbf{p} + 1))$.

4 Experiments

4.1 Implementation Details

In the layered panorama construction stage, we train the panorama LoRA starting from Flux [Labs 2023] with a batch size of 1 and learning rate of 10^{-5} for 100K iterations on our curated UPRIGHT360 dataset. The training is done using 6 NVIDIA A100 GPUs for 3 days. For Layer Decomposition, we employ OneFormer [Jain et al. 2023] to obtain

the panoptic segmentation map for the reference panorama. Background categories are manually determined (i.e. sky, floor, ceiling, etc.) to filter out background components in asset masks. Generally, we cluster all asset masks into $N = 3$ layers via KNN and merge all masks within each layer to form a unified layer mask. With the obtained layer mask, we integrate the above trained panorama LoRA into FLux-Fill model and combine with LaMa [Suvorov et al. 2021] to achieve multi-layer completion and apply 360MonoDepth [Rey-Area et al. 2022] to predict the reference panorama depth. In the 3D panoramic scene optimization stage, we lift the panorama RGBD into 3D point clouds. For scene initialization, we set β_1 to 0.0001 and β_2 to 4 based on empirical practice. These point clouds are used to initialize the base Gaussian model and the layer Gaussian model. During the scene refinement stage, we optimize the base Gaussian model for 3,000 iterations, then the layer Gaussians each for 2,000 iterations. The training objective for base and layer Gaussian is the L1 loss and D-SSIM term between the ground-truth views and the rendered views. We use a single 80G A100 GPU for reconstruction and the reconstruction time for each layer is 1.5 minutes on average for 1024×1024 resolution inputs.

4.2 Comparison Methods.

To evaluate the performance of our approach in text-driven 3D panoramic scene generation domain, we compare with existing methods in two phases: **2D Panorama Generation** and **3D Panoramic Scene Reconstruction**. For 2D Panorama Generation, we compare the quality and creativity of 2D panorama with *Text2light* [Chen et al. 2022] (GAN-based HDR panorama generation), *Diffusion360* [Feng et al. 2023] (diffusion-based text-to-panorama generation) and *Panfusion* [Zhang et al. 2024b] (dual-branch diffusion-based generation). For 3D Panoramic Scene Reconstruction, we compare with *Text2Room* [Höllein et al. 2023] (iterative indoor scene expansion with textured mesh), *LucidDreamer* [Chung et al. 2023] (single-view scene generation with 3DGS), and *Dreamscene360* [Zhou et al. 2024b] (text-guided panoramic 3DGS scene generation).

4.3 Qualitative Comparison

4.3.1 2D Panorama Generation. We show some qualitative comparisons with several state-of-the-art panorama generation works in Fig. 6. Text2Light [Chen et al. 2022] struggles to effectively interpret text prompt due to being trained on a realistic HDRI dataset based on the VQGAN structure, and the components in the generated panorama are relatively simple. The results by PanFusion [Zhang et al. 2024b] are ambiguous and low in quality. While the instances generated by Diffusion360 [Feng et al. 2023] exhibit superior quality in comparison to the aforementioned methods, they lack intricate scene details and are prone to the generation of artifacts. In contrast, our method achieves the highest quality, presenting creative and reasonable generations.

4.3.2 3D Panoramic Scene Reconstruction. We present qualitative comparisons with Text2Room [Höllein et al. 2023], LucidDreamer [Chung et al. 2023], and DreamScene360 [Zhou et al. 2024b] across two dimensions. *First*, for full $360^\circ \times 180^\circ$ view consistency, we render multiple views from scene center point with single input image and text prompts. As shown in Fig. 4, LucidDreamer [Chung et al. 2023]

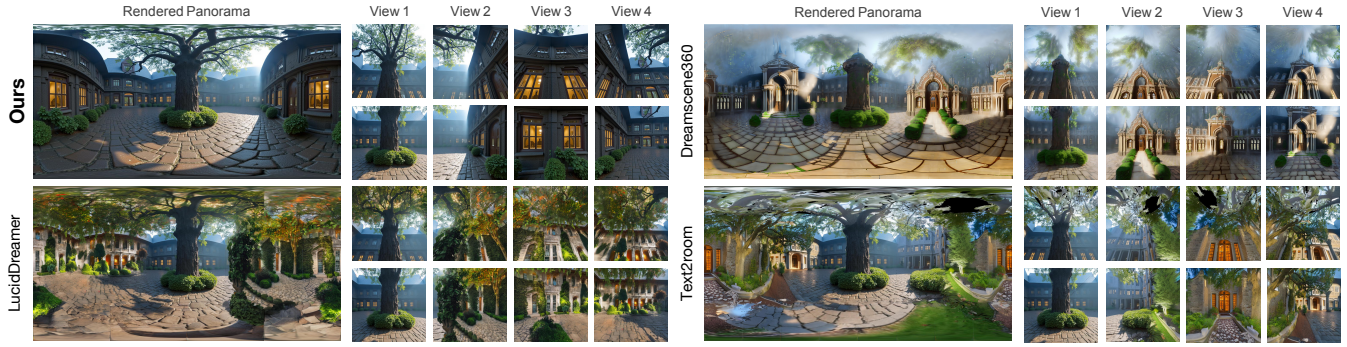


Fig. 4. **Qualitative comparisons in full $360^\circ \times 180^\circ$ Scene.** We compare the panorama and multiple views of the scene generated by four methods. LAYER-PANO3D exhibits consistent and rich details across full $360^\circ \times 180^\circ$ coverage, while other methods show obvious inconsistencies and disorganized patterns in regions that deviate from the input view.




Fig. 5. **Qualitative comparisons in Large-range Scene Exploration.** We show the novel view renderings along a zigzag trajectory to compare the capability of large-range scene exploration. Our method is able to maintain high-quality content rendering and does not show distortion or gaps in unseen space, which shows the ability of LAYER-PANO3D to create hyper-immersive panoramic scenes.

Table 2. **Quantitative comparison with SoTA methods on 3D Panoramic Scene.** LAYER-PANO3D achieves high-quality reconstruction and novel view synthesis while maintaining upright panoramic scene compared to other methods.

Method	Appearance					Geometry		User Study (AUR)	
	NIQE ↓	BRISQUE ↓	PSNR ↑	SSIM ↑	LPIPS ↓	Pitch-Mean ↓	Pitch-Var ↓	$360^\circ \times 180^\circ$ ↑	Free-path ↑
Text2room [Höllein et al. 2023]	5.231	46.127	30.126	0.882	0.038	2.029	1.724	1.69	2.31
LucidDreamer [Chung et al. 2023]	5.822	52.102	36.108	0.954	0.026	2.813	2.189	1.81	1.31
DreamScene360 [Zhou et al. 2024b]	5.051	39.891	30.056	0.958	0.062	1.328	2.018	2.86	2.77
Ours	4.023	38.287	42.057	0.986	0.015	0.732	0.032	3.64	3.61

and Text2room [Höllein et al. 2023] fail to cover the full $360^\circ \times 180^\circ$ view, resulting in semantic incoherence and artifacts due to their successive inpainting-based strategy. DreamScene360 [Zhou et al. 2024b] supports a $360^\circ \times 180^\circ$ view at a single fixed viewpoint, but the quality of the generated results is relatively low. In contrast,

our model excels in maintaining full $360^\circ \times 180^\circ$ view consistency while demonstrating superior content creativity. *Second*, to evaluate novel path rendering, we design a zigzag trajectory to guide the camera’s movement through the scene, with novel view renderings sampled along the trajectory for comparison. Fig. 5 shows 6 random



Fig. 6. **Qualitative comparisons in Panorama Generation.** LAYER-PANO3D demonstrates superior capability in generating high-quality outputs with precise alignment to text prompt, outperforming other methods in fidelity and input adherence.




Fig. 7. **Ablation on the Gaussian Selector.** With the Gaussian Selector, the merged Gaussians are optimized to faithfully reconstruct the ground-truth panorama views.

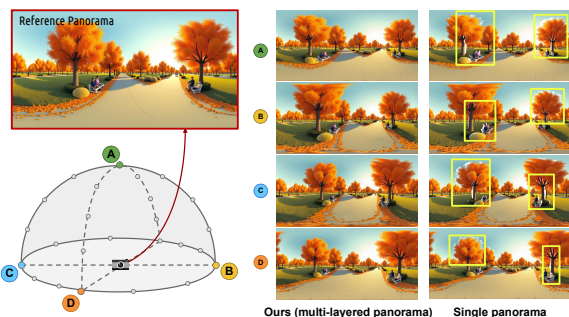


Fig. 8. **Analysis on Panorama Rendering at Off-center Viewpoints.** Compared with the single-layer variant, LAYERPANO3D render $360^\circ \times 180^\circ$ consistent panorama at various off-center viewpoints without holes or gaps from occlusion.

samples from this fixed flythrough trajectory. Compared with all three methods, our model achieves a more complete 3D scene with consistent textures and a reasonable geometric structure.

4.4 Quantitative Comparison

4.4.1 2D Panorama Generation. We adopt three metrics for quantitative comparisons: 1) *FID* [Heusel et al. 2017] evaluates both fidelity and diversity; 2) *Aesthetic* [Schuhmann et al. 2022] evaluates the aesthetics of panorama; 3) *CLIP* [Hessel et al. 2021] measures the compatibility of results with input prompts. Moreover, a user study is also conducted to further evaluate the quality of panoramas, where we project 4 views at a fixed FOV (90°) to the user for sorting. Here, we report the average Average User Ranking (AUR) [Zhang et al. 2023a], which is computed based on an integrated assessment of coherence, plausibility, aesthetics, and compatibility dimensions. As shown in Tab. 1, our method achieves the best scores among all quantitative metrics and human evaluation, demonstrating its fidelity, alignment with text and overall consistency.

4.4.2 3D Panoramic Scene Reconstruction. Following [Zhou et al. 2024b], we adopt non-reference image quality assessment metrics, *NIQE* [Mittal et al. 2012b] and *BRISQUE* [Mittal et al. 2012a], to evaluate novel view quality along scene navigation paths. We also follow [Zhou et al. 2024b] to measure the rendering quality with *PSNR*, *SSIM* and *LPIPS* [Zhang et al. 2018]. In terms of geometry evaluation, we render 4 orthogonal views (90° FOV, 0° elevation and $\{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$ azimuths) and predict the *Pitch-Mean* and *Pitch-Var* (mean and variance of the elevation angles) with [Veicht et al. 2024] to evaluate whether the scenes are upright. As shown in Tab. 2, our method surpasses the existing methods in both novel view quality metrics (*NIQE* and *BRISQUE*), 3D reconstruction metrics (*PSNR*, *SSIM*, *LPIPS*) and while ensuring the upright panoramic scene. Furthermore, we conduct another user study for 3D panoramic scene evaluation from two aspects: 1) $360^\circ \times 180^\circ$ view consistency and 2) novel path rendering quality. For the first aspect, we render 60 frames to cover 360-degree view at the 0° -degree and 45° -degree elevation respectively for evaluation. For the second aspect, we use the same trajectory as in Fig. 5 to render navigation videos for evaluation. We invite 52 users including graduate students that expertise in 3D and average users to rank the 40 results from 4 methods. The average ranking is shown in Tab. 2. Our LAYERPANO3D achieves the best performance in both $360^\circ \times 180^\circ$ view consistency and novel path rendering quality among all four approaches.

4.5 Analysis and Ablative Study

In this section and the supplementary material, we present analysis and ablations on the Gaussian selector (Sec. 4.5.1), multi-layer design (Sec. 4.5.2), view extrapolation, layer Gaussians, SR module, layer inpainting, and 3DGS optimization efficiency (all in supp.).

4.5.1 Ablation on Gaussian Selector. Our Gaussian selector is proposed to select the part of Gaussians that appears in the front of newly added scene assets. By selecting these Gaussians and re-activating them in the optimization, the model achieves accurate appearance and geometry at the current layer. As shown in Fig. 7, the leftmost column is the scene Gaussians at layer 0. When adding the building assets at the first layer, the sky Gaussians from the previous layer partially block the building assets (right column). After using the Gaussian selector to select and optimize the sky Gaussians, these Gaussians learn to either be translucent and pruned for low opacity or move to be a part of the building assets. Therefore in the middle column, we observe a consistent scene with no obvious blockage of the new building assets thanks to the Gaussian Selector.

4.5.2 Analysis on Panorama Renderings at Off-center Viewpoints. In Fig. 8, we demonstrate that LAYERPANO3D is robust to render consistent panorama images at various locations besides the original camera location in the center. We sample four camera locations on circular trajectories on the hemisphere centered at the origin and render 24 views at $(-45^\circ, 0^\circ, 45^\circ)$ elevation to compose new panorama images. By evaluating panorama renderings at new viewpoints, we show that our generated panoramic scene is $360^\circ \times 180^\circ$ consistent and enclosed, robust to various viewpoints at any angle. Compared to the single-layered 3D panorama, our multi-layered 3D panorama exhibits no gaps or holes from the scene occlusion,

demonstrating our capability for larger-range, complex 3D exploration in the generated scenes.

5 Conclusion

In this paper, we propose LAYERPANO3D, a novel framework that generates hyper-immersive panoramic scene from a single text prompt. Our key contributions are two-fold. First, we propose the text-guided anchor view synthesis pipeline to generate detailed and consistent reference panorama. Second, we pioneer the Layered 3D Panorama representation to show complex scene hierarchies at multiple depth layers, and lift it to Gaussians to enable large-range 3D exploration. Extensive experiments show the effectiveness of LAYERPANO3D in generating $360^\circ \times 180^\circ$ consistent panorama at various viewpoints and enabling immersive roaming in 3D space. We believe that LAYERPANO3D holds promise to advance high-quality, explorable 3D scene creation in both academia and industry.

5.1 Limitations and Future Works.

LAYERPANO3D leverages good pre-trained prior to construct panoramic 3D scene, i.e., panoramic depth prior for 3D lifting. Therefore, the created scene might contain artifacts from inaccurate depth estimation. With advancements in more robust panorama depth estimation, we hope to create high-quality panoramic 3D scenes with finer asset geometry.

Acknowledgments

This study is supported by Shanghai AI Laboratory and the Ministry of Education, Singapore, under its MOE AcRF Tier 2 (MOE-T2EP20221-0012, MOE-T2EP20223-0002), and under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).


References

- Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. 2023. MultiDiffusion: Fusing Diffusion Paths for Controlled Image Generation. arXiv:2302.08113 [cs.CV]
- Shengqu Cai, Eric Ryan Chan, Songyou Peng, Mohamad Shahbazi, Anton Obukhov, Luc Van Gool, and Gordon Wetzstein. 2023. DiffDreamer: Towards Consistent Unsupervised Single-view Scene Extrapolation with Conditional Diffusion Models. In *ICCV. IEEE*, 2139–2150.
- Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. 2017. Matterport3D: Learning from RGB-D Data in Indoor Environments. *International Conference on 3D Vision (3DV)* (2017).
- Zhaoxi Chen, Guangcong Wang, and Ziwei Liu. 2022. Text2light: Zero-shot text-driven hdr panorama generation. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–16.
- Jaeyoung Chung, Suyoung Lee, Hyeonjin Nam, Jaerin Lee, and Kyoung Mu Lee. 2023. LucidDreamer: Domain-free Generation of 3D Gaussian Splatting Scenes. *CoRR* abs/2311.13384 (2023).
- Dana Cohen-Bar, Elad Richardson, Gal Metzger, Raja Giryes, and Daniel Cohen-Or. 2023. Set-the-Scene: Global-Local Training for Generating Controllable NeRF Scenes. arXiv:2303.13450 [cs.CV]
- Mengyang Feng, Jinlin Liu, Miaomiao Cui, and Xuansong Xie. 2023. Diffusion360: Seamless 360 Degree Panoramic Image Generation based on Diffusion Models. arXiv:2311.13141 [cs.CV]
- Rafail Fridman, Amit Abecasis, Yoni Kasten, and Tali Dekel. 2023. SceneScape: Text-Driven Consistent Scene Generation. arXiv:2302.01133 [cs.CV]
- Ruiqi Gao, Aleksander Holynski, Philipp Henzler, Arthur Brussee, Ricardo Martin-Brualla, Pratul Srinivasan, Jonathan T Barron, and Ben Poole. 2024. Cat3d: Create anything in 3d with multi-view diffusion models. *arXiv preprint arXiv:2405.10314* (2024).
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. In *EMNLP (1)*. Association for Computational Linguistics, 7514–7528.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *NIPS*, 6626–6637.
- Lukas Höllein, Ang Cao, Andrew Owens, Justin Johnson, and Matthias Nießner. 2023. Text2room: Extracting textured 3d meshes from 2d text-to-image models. *arXiv preprint arXiv:2303.11989* (2023).
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. LoRA: Low-Rank Adaptation of Large Language Models. arXiv:2106.09685 [cs.CL] <https://arxiv.org/abs/2106.09685>
- Jitesh Jain, Jiachen Li, Mang Tik Chiu, Ali Hassani, Nikita Orlov, and Humphrey Shi. 2023. Oneformer: One transformer to rule universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2989–2998.
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.* 42, 4 (2023), 139:1–139:14.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. 2023. Segment anything. *arXiv preprint arXiv:2304.02643* (2023).
- Black Forest Labs. 2023. FLUX. <https://github.com/black-forest-labs/flux>.
- Yuseung Lee, Kunho Kim, Hyunjin Kim, and Minhyuk Sung. 2023. SyncDiffusion: Coherent Montage via Synchronized Joint Diffusions. arXiv:2306.05178 [cs.CV]
- Haoran Li, Haolin Shi, Wenli Zhang, Wenjun Wu, Yong Liao, Lin Wang, Lik hang Lee, and Pengyuan Zhou. 2024. DreamScene: 3D Gaussian-based Text-to-3D Scene Generation via Formation Pattern Sampling. arXiv:2404.03575 [cs.CV]
- Jialu Li and Mohit Bansal. 2023. PanoGen: Text-Conditioned Panoramic Environment Generation for Vision-and-Language Navigation. arXiv:2305.19195 [cs.CV]
- Zhiheng Liu, Hao Ouyang, Qiuyu Wang, Ka Leong Cheng, Jie Xiao, Kai Zhu, Nan Xue, Yu Liu, Yujun Shen, and Yang Cao. 2024. InFusion: Inpainting 3D Gaussians via Learning Depth Completion from Diffusion Prior. *arXiv preprint arXiv:2404.11613* (2024).
- Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. 2012a. No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing* 21, 12 (2012), 4695–4708.
- Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. 2012b. Making a “completely blind” image quality analyzer. *IEEE Signal processing letters* 20, 3 (2012), 209–212.
- Hao Ouyang, Kathryn Heal, Stephen Lombardi, and Tiancheng Sun. 2023. Text2Immersion: Generative Immersive Scene with 3D Gaussians. *arXiv preprint arXiv:2312.09242* (2023).
- Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. 2022. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988* (2022).
- Manuel Rey-Area, Mingze Yuan, and Christian Richardt. 2022. 360monodepth: High-resolution 360deg monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3762–3772.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. 2022. LAION-5B: An open large-scale dataset for training next generation image-text models. In *NeurIPS*.
- Gabriela Ben Melech Stan, Diana Wofk, Estelle Aflalo, Shao-Yen Tseng, Zhipeng Cai, Michael Paulitsch, and Vasudev Lal. 2023. LDM3D-VR: Latent Diffusion Model for 3D VR. arXiv:2311.03226 [cs.CV]
- Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. 2021. Resolution-robust Large Mask Inpainting with Fourier Convolutions. arXiv:2109.07161 [cs.CV]
- Jing Tan, Shuai Yang, Tong Wu, Jingwen He, Yuwei Guo, Ziwei Liu, and Dahua Lin. 2024. Imagine360: Immersive 360 Video Generation from Perspective Anchor. *arXiv preprint arXiv:2412.03552* (2024).
- Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. 2023a. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653* (2023).
- Shitao Tang, Fuyang Zhang, Jiacheng Chen, Peng Wang, and Yasutaka Furukawa. 2023b. MVDiffusion: Enabling Holistic Multi-view Image Generation with Correspondence-Aware Diffusion. arXiv:2307.01097 [cs.CV]
- Alexander Veitch, Paul-Edouard Sarlin, Philipp Lindenberger, and Marc Pollefeys. 2024. GeoCalib: Learning Single-image Calibration with Geometric Optimization. arXiv:2409.06704 [cs.CV] <https://arxiv.org/abs/2409.06704>
- Alexander Vilessov, Pradyumna Chari, and Achuta Kadambi. 2023. CG3D: Compositional Generation for Text-to-3D via Gaussian Splatting. arXiv:2311.17907 [cs.CV]
- Guangcong Wang, Yinuo Yang, Chen Change Loy, and Ziwei Liu. 2022. StyleLight: HDR Panorama Generation for Lighting Estimation and Editing. arXiv:2207.14811 [cs.CV]
- Hai Wang, Xiaoyi Xiang, Yuchen Fan, and Jing-Hao Xue. 2023b. Customizing 360-Degree Panoramas through Text-to-Image Diffusion Models. arXiv:2310.18840 [cs.CV]
- Jionghao Wang, Ziyu Chen, Jun Ling, Rong Xie, and Li Song. 2023a. 360-degree panorama generation from few unregistered nfov images. *arXiv preprint arXiv:2308.14686* (2023).
- Tong Wu, Zhibing Li, Shuai Yang, Pan Zhang, Xingang Pan, Jiaqi Wang, Dahua Lin, and Ziwei Liu. 2023. HyperDreamer: Hyper-Realistic 3D Content Generation and Editing from a Single Image. In *SIGGRAPH Asia*. ACM, 53:1–53:10.
- Tianhao Wu, Chuanxia Zheng, and Tat-Jen Cham. 2024. PanoDiffusion: 360-degree Panorama Outpainting via Diffusion. arXiv:2307.03177 [cs.CV]
- Tao Yang, Peiran Ren, Xuansong Xie, and Lei Zhang. 2023. Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. *arXiv preprint arXiv:2308.14469* (2023).
- Hong-Xing Yu, Haoyi Duan, Junhwa Hur, Kyle Sargent, Michael Rubinstein, William T. Freeman, Forrester Cole, Deqing Sun, Noah Snavely, Jiajun Wu, and Charles Herrmann. 2023. WonderJourney: Going from Anywhere to Everywhere. *CoRR* abs/2312.03884 (2023).
- Hong-Xing Yu, Haoyi Duan, Charles Herrmann, William T Freeman, and Jiajun Wu. 2024. Wonderworld: Interactive 3d scene generation from a single image. *arXiv preprint arXiv:2406.09394* (2024).
- Cheng Zhang, Qianyi Wu, Camilo Cruz Gambardella, Xiaoshui Huang, Dinh Phung, Wanli Ouyang, and Jianfei Cai. 2024b. Taming Stable Diffusion for Text to 360 Panorama Image Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6347–6357.
- Jingbo Zhang, Xiaoyu Li, Ziyu Wan, Can Wang, and Jing Liao. 2024a. Text2NeRF: Text-Driven 3D Scene Generation with Neural Radiance Fields. arXiv:2305.11588 [cs.CV]
- Lvmin Zhang and Maneesh Agrawala. 2024. Transparent Image Layer Diffusion using Latent Transparency. arXiv:2402.17113 [cs.CV] <https://arxiv.org/abs/2402.17113>
- Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023a. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3836–3847.
- Qihang Zhang, Chaoyang Wang, Aliaksandr Siarohin, Peiye Zhuang, Yinghao Xu, Ceyuan Yang, Dahua Lin, Bolei Zhou, Sergey Tulyakov, and Hsin-Ying Lee. 2023b. SceneWiz3D: Towards Text-guided 3D Scene Composition. arXiv:2312.08885 [cs.CV]
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. 2017. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 633–641.
- Haiyang Zhou, Xinhua Cheng, Wangbo Yu, Yonghong Tian, and Li Yuan. 2024a. HoloDreamer: Holistic 3D Panoramic World Generation from Text Descriptions. *arXiv preprint arXiv:2407.15187* (2024).
- Shijie Zhou, Zhiwen Fan, DeJia Xu, Haoran Chang, Pradyumna Chari, Tejas Bharadwaj, Suyu You, Zhangyang Wang, and Achuta Kadambi. 2024b. DreamScene360: Unconstrained Text-to-3D Scene Generation with Panoramic Gaussian Splatting. *arXiv preprint arXiv:2404.06903* (2024).


Created Scene (Extracted)



Layer 2 (Reference Panorama)






Layer 1






Layer 0





MAGICIAN'S CABIN IN A TRANQUIL FOREST, BESIDE A FLOWING RIVER.





CHERRY BLOSSOMS IN FULL BLOOM, PETALS FALLING BY A TRANQUIL RIVER, VIBRANT COLORS.



SANDY BEACH, LARGE DRIFTWOOD IN THE FOREGROUND, CALM SEA BEYOND, REALISM STYLE



ANCIENT STONE MONOLITHS STANDING IN A GRASSY FIELD UNDER A CLOUDY SKY.



MAGICAL ACADEMY COURTYARD WITH FLOATING ORBS OF LIGHT, ANCIENT STONE BUILDINGS, AND A LARGE TREE IN THE CENTER, MYSTICAL AND ENCHANTING.

Fig. 9. Additional results of LAYERPANO3D on Diverse Generation. LAYERPANO3D generates various hyper-immersive scene with consistent and rich details across full $360^\circ \times 180^\circ$ coverage.