

# Generating 360° Video is What You Need For a 3D Scene

ZHAOYANG ZHANG\*, Yale University, USA and Adobe Research, USA

YANNICK HOLD-GEOFFROY, Adobe Research, USA

MILOŠ HAŠAN, Adobe Research, USA

ZIWEN CHEN, Oregon State University, USA

FUJUN LUAN, Adobe Research, USA

JULIE DORSEY, Yale University, USA

YIWEI HU, Adobe Research, USA



Fig. 1. We present *WorldPrompter*, a method that generates traversable environments from a text prompt alone (above). Our method first generates a spatially coherent 360° video capture (top) that travels through a static environment described by the prompt. Next, we elevate generated video capture into a 3D Gaussian Splat (middle), enabling free navigation within a 3D scene using a virtual perspective camera to generate novel views (bottom).

Generating 3D scenes is still a challenging task due to the lack of readily available scene data. Most existing methods only produce partial scenes and provide limited navigational freedom. We introduce a practical and scalable solution that uses 360° video as an intermediate scene representation, capturing the full-scene context and ensuring consistent visual content

\*Work partially done during internship at Adobe Research.

Authors' Contact Information: Zhaoyang Zhang, Yale University, New Haven, CT, USA and Adobe Research, San Jose, CA, USA, zhaoyang.zhang@yale.edu; Yannick Hold-Geoffroy, Adobe Research, San Jose, CA, USA, holdgeof@adobe.com; Miloš Hašan, Adobe Research, San Jose, CA, USA, milos.hasan@gmail.com; Ziwen Chen, Oregon State University, Corvallis, OR, USA, chenziw@oregonstate.edu; Fujun Luan, Adobe Research, San Jose, CA, USA, luanfj11@gmail.com; Julie Dorsey, Yale University, New Haven, CT, USA, julie.dorsey@yale.edu; Yiwei Hu, Adobe Research, San Jose, CA, USA, yiwhu@adobe.com.



This work is licensed under a Creative Commons Attribution 4.0 International License. SA Conference Papers '25, Hong Kong, Hong Kong © 2025 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-2137-3/25/12 <https://doi.org/10.1145/3757377.3763876>

throughout the generation. We propose *WorldPrompter*, a generative pipeline that synthesizes traversable 3D scenes from text prompts. *WorldPrompter* incorporates a conditional 360° panoramic video generator, capable of producing a 128-frame video that simulates a person walking through and capturing a virtual environment. The resulting video is then reconstructed as Gaussian splats by a fast feedforward 3D reconstructor, enabling a true walkable experience within the 3D scene. Experiments demonstrate that our panoramic video generation model, trained with a mix of image and video data, achieves convincing spatial and temporal consistency for static scenes. This is validated by an average COLMAP matching rate of 94.6%, allowing for high-quality panoramic Gaussian splat reconstruction and improved navigation throughout the scene. Qualitative and quantitative results also show it outperforms the state-of-the-art 360° video generators and 3D scene generation models.

CCS Concepts: • **Computing methodologies** → **Computer graphics**.

Additional Key Words and Phrases: Diffusion Models, 360° Video Generative Model, 3D Reconstruction

**ACM Reference Format:**

Zhaoyang Zhang, Yannick Hold-Geoffroy, Miloš Hašan, Ziwen Chen, Fujun Luan, Julie Dorsey, and Yiwei Hu. 2025. Generating 360° Video is What You Need For a 3D Scene. In *SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25)*, December 15–18, 2025, Hong Kong, Hong Kong. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3757377.3763876>

## 1 Introduction

Scene-level 3D generation has attracted increasing research interest in recent years. The limited availability of data generally hinders direct 3D scene generation [Liu et al. 2024a], making it difficult for methods to generalize to realistic scene synthesis effectively. To address this challenge, recent approaches typically rely on an image diffusion model to create realistic scene views. For example, methods like RealmDreamer [Shriram et al. 2025] and LucidDreamer [Chung et al. 2023] can generate Gaussian Splatting (GS) 3D scenes from text prompts. They rely on single-view diffusion models and monocular depth estimation to generate scenes progressively from one view to another by 2D inpainting. However, the 3D scenes created by these methods are confined to a limited viewable area and are not freely traversable. This constraint differentiates these approaches from a "complete" scene generation that offers a navigable virtual world, similar to what is found in interactive video games. Recent works [Yang et al. 2024; Zhou et al. 2024] attempt to bridge the gap between partial and full scene generation by leveraging panoramic images. However, these methods still rely on outpainting to fill in missing regions, which inherently limits the navigable range and often leads to distortions during viewpoint transitions.

More importantly, we observe that scenes generated via these outpainting-based approaches frequently suffer from distortion and inconsistencies in visual content across the scene. We hypothesize that these issues stem from the commonly used progressive outpainting approach, which independently synthesizes new content based on the previously generated content. However, this assumption is flawed, as visual content is globally interconnected and consistent rather than strictly sequential [Tian et al. 2024]. As a result, while their generated scenes may appear locally plausible, they exhibit significant global distortions (Fig. 7).

In this paper, instead of synthesizing a scene in a view-by-view fashion, we aim to generate a maximally navigable scene in a single pass, ensuring undistorted global visual content synthesis. Although scenes can be represented in various formats, such as meshes or point clouds, data in these formats is often limited and difficult to scale. Drawing inspiration from conventional 3D capture methods and recent advancements in video generation models ([Blattmann et al. 2023a; Polyak et al. 2024; Zheng et al. 2024]), we propose using videos as an intermediate representation to model complete 3D environments. Modeling scenes using video captures provides a scalable approach for creating training data, enabling generative models to achieve strong generalization capabilities. To validate this idea, we propose *WorldPrompter*, a generative pipeline that produces a navigable 3D scene from text prompts. Specifically, we train a diffusion model to generate videos traversing the environment based on text prompts. Moreover, instead of synthesizing captures from perspective cameras, we leverage panoramic cameras that directly

capture the entire surroundings of an environment, providing 360° viewpoints that are ideal for our application.

*WorldPrompter* is a two-stage pipeline composed of a 360° panoramic video generation model and a 3D reconstruction step to generate the final 3D Gaussians, a summary of which is shown in Fig. 1. The text-conditioned panoramic video generation model is trained to generate a 360° panoramic video that simulates a person walking through an environment while holding a panoramic camera to capture the scene. To train the video generator, we collected a dataset comprising approximately 1,700 panoramic videos and a panorama image dataset containing more than 200,000 high-quality panoramic in-the-wild captures. To enhance the generative capability of our method, we fine-tuned a Diffusion Transformer (DiT)-based video generator on our dataset to produce 10.6-second 360° panoramic videos (128 frames at 12 fps). We introduce a masked diffusion loss to ensure a clean panoramic capture sequence free from the person holding the camera, ready for subsequent reconstruction purposes. From the generated videos, we extract panoramic frames and project them from their original spherical coordinates onto the image plane to obtain perspective images from multiple viewpoints throughout the scene. The reconstruction step then builds the final 3D Gaussian Splat (3DGS) scene from these cropped images using a fast feed-forward reconstruction model [Ziwen et al. 2024]. The resulting scene allows users to navigate freely within a virtual environment, beyond the original video camera path.

We demonstrate that our 360° video generation model—trained on a mixture of image and video data—achieves state-of-the-art performance in temporal consistency and visual quality. The high consistency of the views, confirmed by a matching rate of up to 94.5% using COLMAP [Schönberger and Frahm 2016], ensures that the videos can be reliably used as proxies for 3D scene generation. This results in a traversable, 360° viewable environment with globally coherent visual content. We demonstrate that our generated scene allows longer navigation range and offers superior global visual coherence and quality compared to prior art. We highlight our contributions as follows:

- We propose leveraging panoramic capture videos as an intermediate 3D scene representation. We found it is a practical solution for scaling up training data for high-quality photo-realistic scene synthesis.
- We present *WorldPrompter*, a practical generative text-to-scene pipeline that synthesizes a traversable 3D scene with strong global visual coherence.
- We introduce a high-quality 360° panoramic video generation model, where the proposed mixed datasets and masked diffusion loss play a critical role in achieving artifact-free and generalizable static 360° scene generation.

## 2 Related Work

### 2.1 Panoramic Image and Video Generation

In computer graphics, 360° panoramas—often called IBLs or environment maps—are widely used to represent distant lighting and environments [Debevec and Malik 2023]. Numerous studies estimate lighting as High Dynamic Range (HDR) panoramas from input photos [Dastjerdi et al. 2023; Weber et al. 2022; Zhan et al. 2021].

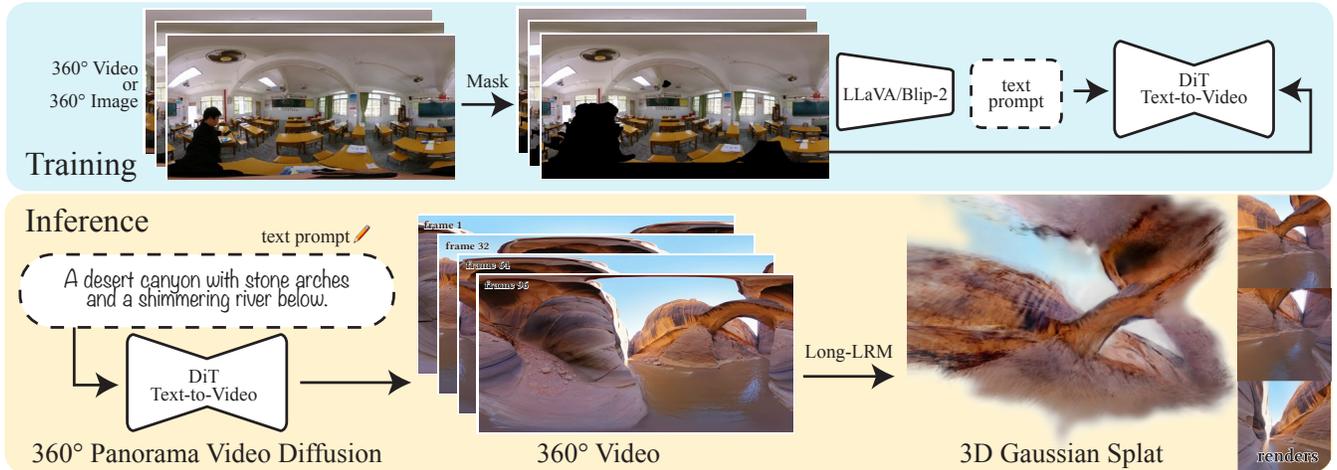


Fig. 2. **Overview of WorldPrompter.** (top) We train a text-to-video model on a mix of 360° videos and images depicting in-the-wild environments. As it is challenging to avoid the person and camera equipment being visible in the video capture, we mask these elements out of the frame using a pretrained image segmentation model [Adobe 2025], and obtain the prompts from the video frames using LLaVA [Liu et al. 2023b] for the videos, and BLIP-2 [Li et al. 2023a] for images. (bottom) At inference time, a user supplies a text prompt to our text-to-video model, which produces a “walk-through” video of the scene, which we reconstruct into a 3D Gaussian splat representation using Long-LRM [Ziwen et al. 2024].

However, the resolution of these panoramas is often limited, making them suitable for light representation but insufficient for 3D reconstruction.

Advancements in diffusion models have made high-resolution image synthesis possible [Esser et al. 2024; Rombach et al. 2022]. Panorama generation also benefits from the use of diffusion models. For example, Text2Light [Chen et al. 2022] generates 4K HDR panoramas via text-driven low-resolution generation and super-resolution inverse tone mapping. MVDiffusion [Tang et al. 2023] creates panoramas by stitching consistent multi-view images generated from text prompts using attention layers. PanFusion [Zhang et al. 2024b] employs a dual-branch diffusion model with equirectangular-Perspective Projection Attention (EPPA) to improve consistency and control across panorama and perspective domains. DiffPano [Ye et al. 2024a] fine-tunes Stable Diffusion [Rombach et al. 2022] using LoRA [Hu et al. 2021] to generate panorama images, and with a Spherical Epipolar-Aware Attention Module, it enables multi-view panorama generation.

Diffusion models have also demonstrated their potential in video generation [Blattmann et al. 2023a,b]. Diffusion Transformers [Peebles and Xie 2023] offer a more versatile method for modeling high-dimensional data, making them particularly well-suited for video generation [Brooks et al. 2024]. For 360° video generation, Wang et al. propose 360DVD for generating 360° panoramic videos from text prompts and motion conditions, where they leverage a lightweight 360-Adapter to finetune the Stable Diffusion [Rombach et al. 2022] text-to-image model. However, their model is limited to generating only 16 frames, whereas our 360° video generation model produces significantly more frames with superior temporal consistency.

## 2.2 3D Neural Representation and Reconstruction

3D reconstruction and representation is an extensively studied field. While mesh-based methods [Kazhdan et al. 2006] and point-based methods [Qi et al. 2017] have been extensively explored, they are limited in the quality of their results. In contrast, neural representations offer a flexible way to reconstruct 3D data. For instance, Neural Radiance Fields (NeRF) [Mildenhall et al. 2020] introduces a novel way to represent 3D scenes by encoding volumetric scene information as a continuous function parameterized by a neural network, which inspired many follow-up works [Barron et al. 2022, 2023; Müller et al. 2022] that optimize efficiency, quality, and rendering speed. However, the MLP-based paradigm requires training a distinct model per scene, making reconstruction time-consuming. Additionally, rendering views involves model queries, creating a bottleneck for real-time applications. 3D Gaussian Splatting (3DGS) [Kerbl et al. 2023] introduces a memory-efficient and low-overhead 3D implicit representation using Gaussians derived from Structure-from-Motion (SfM) points. By combining anisotropic covariance optimization with a fast visibility-aware rendering algorithm, 3DGS enables efficient scene representation, high-quality real-time rendering, and has quickly emerged as a new paradigm in neural reconstruction [Charatan et al. 2024; Gao et al. 2024; Jiang et al. 2024; Wu et al. 2024; Ye et al. 2024b]. Nevertheless, 3DGS still requires an optimization process, which can be costly for larger scenes. Recent advances such as Large Reconstruction Models (LRMs) [Wang et al. 2023; Wei et al. 2024; Xie et al. 2024; Zhang et al. 2024a; Ziwen et al. 2024] target scene reconstruction with a single feed-forward evaluation during inference. They achieve a reconstruction quality comparable to optimization-based methods while significantly improving their efficiency. In our scene generation pipeline, we also utilize this state-of-the-art feed-forward reconstruction approach for long-sequence reconstruction, achieving fast reconstruction.

### 2.3 3D Scene Generation

3D generation has emerged as a prominent focus in generative modeling. Given the scarcity of high-quality 3D data, recent works usually leverage intermediate representations, such as multi-view images [Hu et al. 2024; Li et al. 2023b; Liu et al. 2024b, 2023c; Shi et al. 2024; Voleti et al. 2024] and videos [Chen et al. 2024; Han et al. 2024; Parthasarathy et al. 2024], to improve the quality and efficiency of their reconstructions. Many studies have explored generating 3D objects in various formats, including Gaussian splats and meshes. For instance, Instant3D [Li et al. 2023b] accelerates 3D asset creation from text prompts by combining sparse-view generation with transformer-based reconstruction. Similarly, MVDream [Shi et al. 2024] improves consistency through multi-view diffusion by integrating 2D and 3D data. One-2-3-45 [Liu et al. 2024b] simplifies single-image 3D reconstruction, enabling view-consistent mesh generation.

Scene-level generation, however, is a more challenging task because it requires more observations to achieve comprehensive view coverage. Typically, 3D scene generation methods employ progressive pipelines [Chung et al. 2023; Höllein et al. 2023; Ouyang et al. 2023]: anchor cameras are first selected to generate text-aligned 2D images and corresponding depth maps, followed by a neural reconstruction process to build the initial 3D structure. This process is iteratively repeated to expand scene coverage by generating additional camera views and inpainting missing regions. However, such methods often suffer from prolonged generation times and incomplete coverage. Other studies have explored interactive and layout-guided approaches for 3D scene generation. Systems like WonderWorld [Yu et al. 2024a] provide user-friendly interfaces for scene customization, while Director3D [Li et al. 2024a] facilitates real-world scene creation with adaptive camera trajectories. Another potential limitation of this family of methods, as discussed in Sec. 1, is the lack of guaranteed global visual coherence, often resulting in distorted scene layouts and object shapes. Most recent literature has discovered the potential of video generation for consistent scene-level generation. VideoScene [Wang et al. 2025] leverages the temporal consistency in video diffusion models to generate novel views for 3D reconstruction by interpolating two input keyframes. Yu et al. and Wu et al. explore explicitly reconstructing a 3D point cloud as spatial memory for consistent video generation.

The closest method to ours is DreamScene360 [Zhou et al. 2024], which generates 3D scenes from text prompts using a single panorama image, along with depth estimation and Gaussian Splats optimization. However, while the resulting scene supports 360° viewing, it suffers from a limited movable range. LayerPano3D [Yang et al. 2024] offers slightly increased navigational range by generating layered panoramas, but the content in the additional layers may suffer from distortion. ODIN [Wallingford et al. 2024] proposed learning 3D from 360° videos. However, their approach extracts only cropped views from videos to perform image-to-image style novel view synthesis, underutilizing the full potential of 360° videos. As a result, ODIN fails to model a complete 3D world and can only generate partial scenes. Beyond 3D, 4K4DGen [Li et al. 2024b], introduces panorama-to-4D generation but still remains limited in terms of traversable area.

In contrast to these approaches, our model introduces 360° videos as an explicit 3D scene proxy in the generation task, offering greater data scalability and ensuring strong visual coherence alongside comprehensive view coverage, ultimately enabling a realistic navigation experience.

## 3 WorldPrompter

### 3.1 Overview

We show that 360° video capture serves as an effective, all-you-need proxy for 3D scene generation, as validated by the *WorldPrompter* pipeline. It synthesizes a traversable, 360° viewable 3D scene in two stages: *generation* and *reconstruction*. The *generation* stage leverages a text-conditioned 360° panoramic video diffusion model to generate a 10.6-second panoramic capture video for a target scene, simulating someone holding the camera and walking through the environment. Later, the *reconstruction* stage runs COLMAP [Schönberger and Frahm 2016] to estimate precise camera poses from perspective crops of the generated videos and builds a navigable 3DGS using Long-LRM [Ziwen et al. 2024] with the calibrated camera poses. The generated panoramic video features rich details and geometrically consistent visual details over the entire scene, ensuring high-quality reconstruction and yielding appealing novel views in the traversable scene.

### 3.2 Video Generator for 360° Panoramic Captures

Training a generative model for 360° videos from scratch is challenging; therefore, we start with an existing video generator. We leverage a pretrained text-to-video diffusion model as backbone, and fine-tune it using panoramic data.

Specifically, our model’s backbone architecture  $f_{\theta}$  is based on Diffusion Transformers [Peebles and Xie 2023]. Typically, in latent DiT models, a video or an image is first encoded by a 3DVAE [Yu et al. 2024b], resulting a latent visual sequence  $x \in \mathbb{R}^{T \times 3 \times H \times W}$  ( $T = 1$  if it’s an image). The visual sequence is tokenized into patches, yielding a set of visual tokens  $\hat{x} \in \mathbb{R}^{K \times D}$ , where  $T$ ,  $H$ , and  $W$  correspond to the temporal and spatial dimensions of the video,  $K$  represents the total number of tokens, and  $D$  denotes the feature dimension. Being a text-conditioned model, it requires text prompts  $c$  that we encode using a pre-trained model [Radford et al. 2021] and map into the same feature space  $D$  by an embedding layer. As a decoder-only transformer architecture, the embedded text tokens are concatenated with the noised visual tokens. The transformer  $f_{\theta}(\hat{x}_t; c, t)$  operates on these inputs, denoising the visual tokens at each timestep  $t$ . The denoised tokens  $\hat{x}_0 \in \mathbb{R}^{K \times D}$  are subsequently decoded and reassembled into the visual data  $x_0 \in \mathbb{R}^{T \times 3 \times H \times W}$  via linear layers, and finally decoded as videos or images.

To preserve the generalization ability of the finetuned t2v model across diverse prompts, we curate two carefully selected datasets: (1) a 360° video dataset capturing static indoor environments, and (2) a 360° image dataset featuring a wide range of real-world scenes. We also incorporate a masked diffusion loss, which is key to ensuring artifact-free generation.

**3.2.1 360° video dataset.** We start by collecting a 360° video dataset with approximately 1,700 panoramic videos. Each video is captured

by a person holding a Ricoh Theta Z1 camera walking inside a nearly static indoor scene to ensure the quality of 3D scene reconstruction. Each video has around 50 seconds of capture, yielding roughly 1,500 frames. To build a dataset for training, we segment the longer video sequences into smaller chunks, typically a 128-frame video clip (5.3/10.6 seconds at 24/12fps). We use LLaVA [Liu et al. 2023a] to generate detailed captions describing each scene.

**3.2.2 360° image dataset.** However, as the 360° video data is primarily captured indoors to ensure a static environment with no moving objects during the capture process, this may bias the fine-tuned model, limiting its ability to generalize effectively to outdoor environments. To mitigate this issue, we gather another dataset with around 200,000 panoramic images taken with cameras on a tripod, without the photographer in the scene. We caption each image with Blip-2 [Li et al. 2023a] to describe the whole scene in detail. These stationary images, captured in-the-wild, represent a much broader variety of scenes, thereby enhancing the generalization capability of our video diffusion model.

**3.2.3 Mixed Training with Masked Loss.** We combine the two described datasets to train our model, enabling the text-to-360° video generator to produce both 360° videos and images, with images treated as single-frame videos. While panoramic images depict a clean environment without the photographer present, the 360° video capture often includes the photographer visible in the scene. This causes the trained model to produce camera operators in its outputs (see Fig. 4), complicating subsequent reconstruction steps. Hence, when training our diffusion models on video data, we introduce a *masked* diffusion loss, where the photographer is excluded from the loss computation.

To achieve this, we preprocess all 360° videos by masking out the photographer using a pretrained image segmentation model [Adobe 2025] on each frame in a video. For a 128-frame chunk, we merge all segmentation masks to minimize segmentation errors and ensure a clean region without the person. Additionally, we conservatively mask out the bottom region of the equirectangular panoramic frames to eliminate visible camera equipment and the photographer’s hand. Assuming the corresponding binary mask  $M$  is resized into the latent space by nearest neighbor interpolation, the masked diffusion loss is computed as

$$\mathbb{E}_{x \sim p_{\text{data}}, t \sim U(0,1)} \left[ \|M \odot (\epsilon_t - f_{\theta}(\hat{x}_t; c, t))\|^2 \right], \quad (1)$$

where  $\epsilon_t$  is the sampled noise at timestep  $t$ , and  $M$  has the same dimensions as  $\epsilon_t$ , where  $M[i, j] = 1$  includes the region in the loss calculation, and  $M[i, j] = 0$  excludes it (indicating the visible person).

Our video generator is finally trained on a mix of 1/3 image data and 2/3 video data, achieving an empirically reasonable balance. The video diffusion model generates a 128-frame video clip at 12 fps, equivalent to a 10.6-second 360° capture video at a native resolution of 352×704 pixels. We further apply a pretrained GAN-based video upsampler [Wang et al. 2018] to super-resolve our generated videos 4 times, resulting a final 360° video of 1416×2832 resolution.

### 3.3 Reconstruction from “Generated” Captures

Given the generated 360° capture video with 128 panoramic frames, we reconstruct a 3D Gaussian scene. For each panoramic frame, we randomly perform 3 perspective crops with a field of view of 120° and resolution of 512×512, resulting in 384 perspective images for pose estimation and reconstruction. Similar to standard scene reconstruction, we calibrate the camera poses of these perspective images using COLMAP [Schönberger and Frahm 2016]. The high average matching ratio 94.6% from COLMAP estimation demonstrates the strong view consistency of the generated videos from our text-to-360° video generator. We adopt Long-LRM, the state-of-the-art feed-forward 3DGS reconstructor [Ziwen et al. 2024] to reconstruct the final 3D scene using the estimated camera poses and generated views. Using Long-LRM allows us to achieve fast reconstruction compared to an optimization-based approach. We randomly subsample 32 random views along the original camera path as the inputs since the Long-LRM has a maximum limit for input views. The calibrated camera poses from COLMAP are converted into Plücker rays as the inputs for long-LRM.

### 3.4 Implementation Details

We fine-tune a pretrained text-to-video diffusion model using mixed 360° image and video datasets. The fine-tuning is conducted on 32 NVIDIA A100 GPUs with a mixed fps training schedule (12fps and 24fps). The model was finetuned from a pretrained text-to-video checkpoint for 20,000 steps in around 200 hours. During inference, we generate videos at 12 fps to extend the traversal range while maintaining good temporal consistency comparable to 24 fps generation.

The inference of a 10.6-second video takes around 10 minutes with an unoptimized PyTorch runtime [Paszke et al. 2019] using 50 diffusion steps, together with the video upsampling process. The entire reconstruction takes around 5 minutes, including pose estimation and Long-LRM reconstruction. Our full 3D generation pipeline requires approximately 15 minutes to generate a detailed GS scene, with potential for further optimization to improve inference speed and efficiency for different components.

## 4 Results

By generating 360° video capture, we demonstrate that our pipeline achieves state-of-the-art quality in 3D scene generation (Sec. 4.2), offering significantly improved view coverage, traversability and coherent global visual content. This is primarily due to a strong 360° video generator for scenes (Sec. 4.1). To quantitatively evaluate our method, we prompt ChatGPT [OpenAI 2023] to generate a list of 320 descriptions of 3D scenes. Using this in-the-wild test set, we conduct systematic comparisons against existing methods, demonstrating superior video and 3D scene generation quality through both visual examples and quantitative metrics. More visual examples can be found in the supplementary material.

Additionally, we perform ablation studies in Sec. 4.3 to evaluate our design choices, including mixed data training and masked loss, to ensure the generation of high-quality videos suitable for artifact-free 3D scene reconstruction.

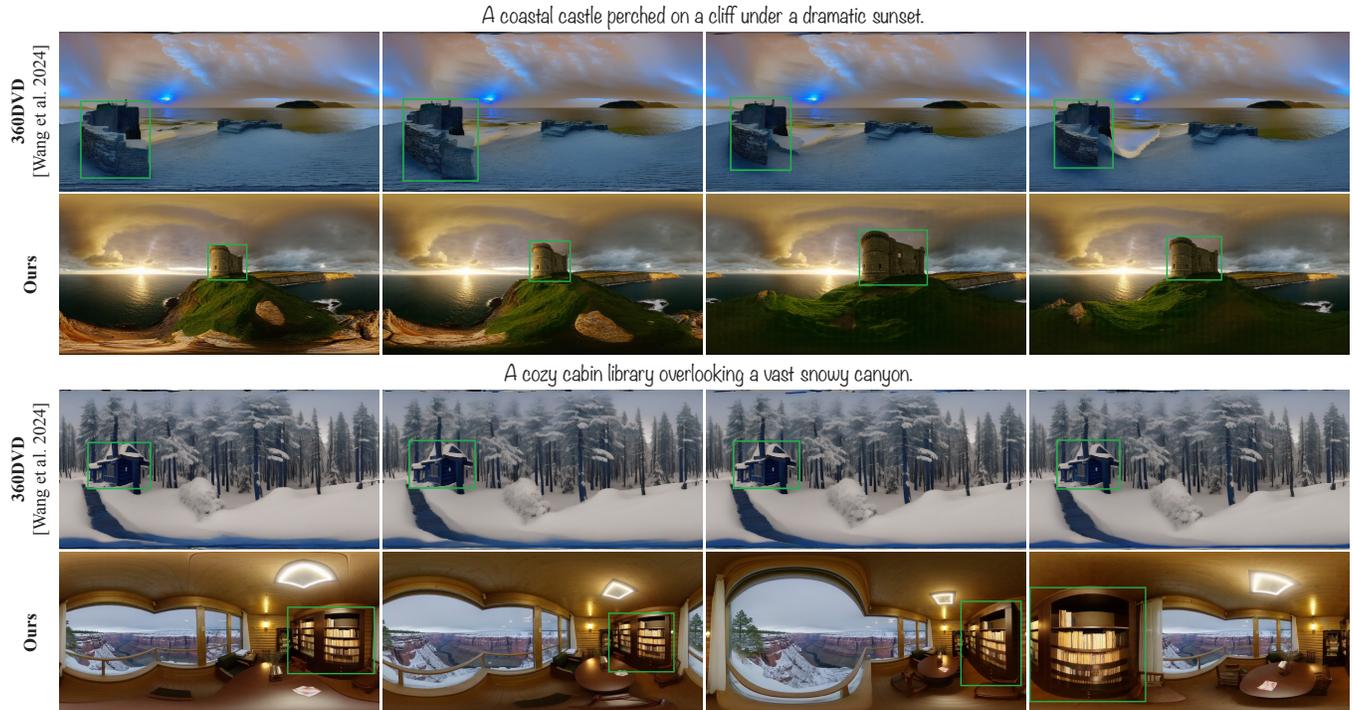


Fig. 3. **Comparison of 360° video generation results.** We present a comparison of 360° panoramic video generation results against 360DVD [Wang et al. 2024]. Our generated videos exhibit significantly better visual quality and prompt alignment. We manually mark an anchor region with a green box in the 360° image, highlighting the substantially longer traversal range of our generated videos. Stronger parallax is important for accurate 3D reconstruction.

Table 1. We quantitatively evaluate our method against the state-of-the-art (SOTA) 360° video generation model [Wang et al. 2024]. We report CLIP distance (CLIP) for prompt alignment quality, Q-Align score for image quality, VBench for aesthetic/imaging quality, and COLMAP Failure Rate (COLMAP-FR) along with Matching Rate (COLMAP-MR) for multi-view consistency. The metrics are computed on the generated 360° frames across 320 test text prompts. Quantitative metrics across multiple dimensions demonstrate that our generated 360° videos achieve high visual quality and can be reliably used for 3D scene representation and reconstruction.

	CLIP↓	Q-Align↑	VBench↑	COLMAP-FR↓	COLMAP-MR↑
360DVD	0.7643	0.6704	0.538/0.661	89.7%	41.8%
<b>Ours</b>	<b>0.7387</b>	<b>0.8334</b>	<b>0.615/0.717</b>	<b>0%</b>	<b>94.6%</b>

#### 4.1 Text-to-360°-Video Generation

A 360° video generator that accurately reproduces the distribution of real-world video captures is key to making it a sufficient proxy for 3D scene representation. Our generation model can generate spatial-temporal consistent 360° panoramic videos, exceeding the quality of the state-of-the-art panoramic video generation models. We compare our 360° video generation results with 360DVD [Wang et al. 2024] in Fig. 3. We showcase four evenly sampled frames comparing 360DVD and our 360° panorama video generation model. Our model generates videos that closely match the input text prompt and produce more realistic scene details compared to 360DVD. In Table 1, we evaluate generation quality across multiple axes, where

Table 2. We quantitatively evaluate our method against several text-to-scene models, including [Chung et al. 2023], [Zhou et al. 2024], and [Yang et al. 2024]. To assess prompt adherence and visual quality, we compute CLIP distance and Q-Align metrics on the generated 360° images or frames across 320 test prompts randomly generated by ChatGPT. Note that LucidDreamer is an image-to-scene generation model which requires both text and image as inputs, and we generate the initial image using ChatGPT; therefore CLIP and Q-Align score evaluating the panorama synthesis quality is not applicable for LucidDreamer. Our method outperforms all baselines across these evaluations. Furthermore, we conduct a 4AFC user study to assess three key aspects of scene quality: Realism/Completeness/Traversal Consistency and report the selection percentages.

	CLIP↓	Q-Align↑	Real. ↑	Complete. ↑	Consist. ↑
LucidDreamer	N/A	N/A	21.0%	10.0%	10.0%
DreamScene360	0.8116	0.8074	7.0%	6.0%	6.0%
LayerPano3D	0.7977	0.8312	34.0%	25.0%	29.0%
<b>Ours</b>	<b>0.7387</b>	<b>0.8334</b>	<b>38.0%</b>	<b>59.0%</b>	<b>55.0%</b>

our generator consistently outperforms previous SOTA in common image/video generation quality metrics e.g., CLIP [Radford et al. 2021], Q-align [Wu et al. 2023], VBench [Huang et al. 2024].

Additionally, videos generated by our model cover a broader scene span compared to 360DVD, which exhibits only slight rotations and translations of the viewpoint as indicated in the marked green box (Fig. 3). More importantly, since the primary goal of 360° panoramic video synthesis is to generate a coherent 3D scene, the quality of



Fig. 4. We conduct an ablation study to evaluate the impact of the masked diffusion loss. This loss function effectively guides the model to generate only the pixels of the scene, eliminating the moving photographer present in the original video training data. As a result, it enables the generation of a complete 360° scene capture for reconstruction.

360DVD’s results prevents it from effectively accomplishing this task. Its outputs lack structural consistency across views, exhibiting continuous morphing and noticeable changes, which undermine the perception of a stable 3D environment.

To validate this statement, we run COLMAP on their generated panoramas using a similar cropping and pose estimation procedure as described in Sec. 3.3. As shown in Table 1, we find that the failure rate of COLMAP reaches 89.7% (COLMAP-FR), and even in cases where COLMAP succeeds, the match rate (COLMAP-MR) remains very low, with only 41.8% images being successfully matched, making their results unsuitable for 3D reconstruction. In contrast, our model preserves view-consistent local structures across frames. We do not encounter any failure cases when running COLMAP, and the average image match rate is as high as 94.6% across our test prompts. This is the key reason our 360° videos can be a reasonable 3D representation for reconstruction.

## 4.2 Text-to-Scene Generation

As shown in Fig. 6, by chaining our 360° video generator with the 3D reconstruction process, we enable the creation of a diverse set of 3D scenes, allowing users to navigate the environment from novel viewpoints and spatial locations. More results can be found in the supplementary material.

In Fig. 7, we compare our results to the previous state-of-the-art text-to-scene method that can also create 3D scenes. LucidDreamer [Chung et al. 2023], an outpainting-based scene generation method, synthesizes single-perspective views using an image diffusion model,



Fig. 5. We verify that incorporating additional panorama image data helps the model generalize to diverse environments. Without the image data, the generated 360° videos not only exhibit degraded visual appearance but also display messy and incorrect content at the bottom of the video. This issue arises because these pixels lack supervision due to the masked loss.

but suffers from global geometric distortions and lacks support for full-view navigational experiences. Other methods that construct 3D scenes from panoramic images—such as DreamScene360 [Zhou et al. 2024] and LayerPano3D [Yang et al. 2024]—provide 360° view coverage, but do not enable truly navigable scenes, as significant camera movement often breaks scene coherence. For all other methods, when synthesizing novel views (Novel View 1–4 in Fig. 7) from positions different from the initial viewpoint, the rendered images exhibit significant artifacts, including blurriness, missing regions, and structural distortions. Similarly, we report the CLIP distance and Q-Align metrics on the generated 360° images, demonstrating that our method generates a 3D environment that more accurately responds to user requests while achieving superior visual quality, even at the initial position.

We argue that quality assessment is best conducted within the context of a real navigation experience. To this end, we perform a user study in which participants are presented with navigation videos generated by different methods. For each approach, we render a short video clip simulating camera movement through the scene, and ask users to select the most preferred result. Participants are instructed to evaluate the methods using a four-alternative forced-choice (4AFC) paradigm for the following dimensions: (1) Realism: Realistic Visual (2) Completeness: No missing regions (3) Traversal Consistency: Stable, distortion-free geometry during navigation. We collect 20 responses and report the selection percentages in

Table 2. Our method received the majority of votes across all evaluation criteria, with particularly strong margins in Completeness and Traversal Consistency.

### 4.3 Ablation Study

Training a high-quality 360° video generator within our 3D generative pipeline is a key factor in the success of our method, relying on the use of mixed training data and the masked diffusion loss. We conduct ablation studies on these design choices to demonstrate their effectiveness.

**4.3.1 Mixed Training Data.** We conduct an experiment by excluding the additional image panorama data. As shown in Fig. 5, incorporating image data significantly enhances the generalization capability of the video generator across a wide variety of scenes, and improves visual quality. Notably, due to the masked diffusion loss training, ground-truth pixels in the bottom regions of the panoramic videos are unavailable. The image data compensates for this limitation, enabling the model to intelligently auto-inpaint the bottom regions, resulting in a natural and realistic appearance.

**4.3.2 Masked Diffusion Loss.** We also train a model using the full diffusion loss applied to all pixels. The generated videos in this case reveal the photographer and the camera pod, which consistently appear in the 360° video training data (Fig. 4). This obstructs a significant portion of the generated capture videos, making comprehensive scene reconstruction challenging. By applying the masked diffusion loss and incorporating the aforementioned image panorama data, we guide the model to generate clean 360° videos that can be directly used for perspective cropping and reconstruction.

### 4.4 Discussions and Limitations

The quality of our text-to-scene generation pipeline is constrained by both the length and resolution of the panoramic video generative model. On one hand, generating minute-long panoramic videos while maintaining visual consistency across extended sequences is essential for scaling to larger, more detailed 3D scenes. On the other hand, panoramic videos allocate more pixels to capture global 3D content and ensure spatial consistency, but this comes at the cost of fine-grained details compared to perspective videos at the same resolution. We address this limitation with an upsampler, though a more fundamental solution would be a high-resolution video generator.

## 5 Conclusion

In conclusion, we demonstrate that 360° video capture is an effective and practical representation for scene-level 3D generation, as validated by *WorldPrompter*, a generative pipeline capable of producing navigable 3D scenes from text prompts. By leveraging 360° panoramic video generation as an intermediate representation, combined with a fast 3D reconstruction module, the pipeline enables the synthesis of immersive environments with comprehensive view coverage and high visual fidelity. This scalable approach effectively

bridges the gap between text-to-3D generation and full-scene synthesis. We expect this novel idea highlights a promising research direction and has the potential to inspire future work in the field.

## References

- Adobe. 2025. Select Subject.
- Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. 2022. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. arXiv:2111.12077 [cs.CV] <https://arxiv.org/abs/2111.12077>
- Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. 2023. Zip-NeRF: Anti-Aliased Grid-Based Neural Radiance Fields. arXiv:2304.06706 [cs.CV] <https://arxiv.org/abs/2304.06706>
- Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, Varun Jampani, and Robin Rombach. 2023a. Stable Video Diffusion: Scaling Latent Video Diffusion Models to Large Datasets. arXiv:2311.15127 [cs.CV] <https://arxiv.org/abs/2311.15127>
- Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. 2023b. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 22563–22575.
- Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. 2024. Video generation models as world simulators. (2024). <https://openai.com/research/video-generation-models-as-world-simulators>
- David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. 2024. pixelSplat: 3D Gaussian Splats from Image Pairs for Scalable Generalizable 3D Reconstruction. arXiv:2312.12337 [cs.CV] <https://arxiv.org/abs/2312.12337>
- Zhaoxi Chen, Guangcong Wang, and Ziwei Liu. 2022. Text2Light: Zero-Shot Text-Driven HDR Panorama Generation. *ACM Transactions on Graphics (TOG)* 41, 6, Article 195 (2022), 16 pages.
- Zilong Chen, Yikai Wang, Feng Wang, Zhengyi Wang, and Huaping Liu. 2024. V3D: Video Diffusion Models are Effective 3D Generators. arXiv:2403.06738 [cs.CV] <https://arxiv.org/abs/2403.06738>
- Jaeyoung Chung, Suyoung Lee, Hyeongjin Nam, Jaerin Lee, and Kyoung Mu Lee. 2023. LucidDreamer: Domain-free Generation of 3D Gaussian Splatting Scenes. *arXiv preprint arXiv:2311.13384* (2023).
- Mohammad Reza Karimi Dastjerdi, Jonathan Eisenmann, Yannick Hold-Geoffroy, and Jean-François Lalonde. 2023. EverLight: Indoor-Outdoor Editable HDR Lighting Estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 7420–7429.
- Paul E Debevec and Jitendra Malik. 2023. Recovering high dynamic range radiance maps from photographs. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. 643–652.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*.
- Lin Gao, Jie Yang, Botao Zhang, Jiamu Sun, Yujie Yuan, Hongbo Fu, and Yu-Kun Lai. 2024. Real-time Large-scale Deformation of Gaussian Splatting. *ACM Transactions on Graphics (SIGGRAPH Asia 2024)* (2024).
- Junlin Han, Filippos Kokkinos, and Philip Torr. 2024. VFusion3D: Learning Scalable 3D Generative Models from Video Diffusion Models. arXiv:2403.12034 [cs.CV] <https://arxiv.org/abs/2403.12034>
- Lukas Höllein, Ang Cao, Andrew Owens, Justin Johnson, and Matthias Nießner. 2023. Text2room: Extracting textured 3d meshes from 2d text-to-image models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7909–7920.
- 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>
- Hanzhe Hu, Zhizhuo Zhou, Varun Jampani, and Shubham Tulsiani. 2024. MVD-Fusion: Single-view 3D via Depth-consistent Multi-view Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 9698–9707.
- Ziqi Huang, Yanan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, Yaohui Wang, Xinyuan Chen, Limin Wang, Dahua Lin, Yu Qiao, and Ziwei Liu. 2024. VBench: Comprehensive Benchmark Suite for Video Generative Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Kaiwen Jiang, Yang Fu, Mukund Varma T, Yash Belhe, Xiaolong Wang, Hao Su, and Ravi Ramamoorthi. 2024. A Construct-Optimize Approach to Sparse View Synthesis without Camera Pose. *SIGGRAPH* (2024).
- Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. 2006. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, Vol. 7.
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on*

- Graphics* 42, 4 (July 2023). <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>  
Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*. PMLR, 19730–19742.
- Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. 2023b. Instant3D: Fast Text-to-3D with Sparse-View Generation and Large Reconstruction Model. arXiv:2311.06214 [cs.CV] <https://arxiv.org/abs/2311.06214>
- Renjie Li, Panwang Pan, Bangbang Yang, Dejia Xu, Shijie Zhou, Xuanyang Zhang, Zeming Li, Achuta Kadambi, Zhangyang Wang, Zhengzhong Tu, and Zhiwen Fan. 2024b. 4K4DGen: Panoramic 4D Generation at 4K Resolution. arXiv:2406.13527 [cs.CV] <https://arxiv.org/abs/2406.13527>
- Xinyang Li, Zhangyu Lai, Linning Xu, Yansong Qu, Liujuan Cao, Shengchuan Zhang, Bo Dai, and Rongrong Ji. 2024a. Director3D: Real-world Camera Trajectory and 3D Scene Generation from Text. arXiv:2406.17601 (2024).
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023a. Improved Baselines with Visual Instruction Tuning.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual Instruction Tuning. In *NeurIPS*.
- Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. 2024b. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *Advances in Neural Information Processing Systems* 36 (2024).
- Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. 2023c. Zero-1-to-3: Zero-shot One Image to 3D Object. arXiv:2303.11328 [cs.CV] <https://arxiv.org/abs/2303.11328>
- Yuheng Liu, Xinke Li, Xueting Li, Lu Qi, Chongshou Li, and Ming-Hsuan Yang. 2024a. Pyramid Diffusion for Fine 3D Large Scene Generation. In *European Conference on Computer Vision (ECCV)*.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Trans. Graph.* 41, 4, Article 102 (July 2022), 15 pages. doi:10.1145/3528223.3530127
- OpenAI. 2023. ChatGPT: A Large Language Model. <https://chat.openai.com> Accessed: 2025-03-07.
- Hao Ouyang, Tiancheng Sun, Stephen Lombardi, and Kathryn Heal. 2023. Text2Immersion: Generative Immersive Scene with 3D Gaussians. arXiv (2023).
- Rishabh Parthasarathy, Zachary Anknor, and Aaron Gokaslan. 2024. Vid3D: Synthesis of Dynamic 3D Scenes using 2D Video Diffusion. arXiv:2406.11196 [cs.CV] <https://arxiv.org/abs/2406.11196>
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019).
- William Peebles and Saining Xie. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4195–4205.
- Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, David Yan, Dhruv Choudhary, Dingkang Wang, Geet Sethi, Guan Pang, Haoyu Ma, Ishan Misra, Ji Hou, Jialiang Wang, Kiran Jagadeesh, Kunpeng Li, Luxin Zhang, Mannat Singh, Mary Williamson, Matt Le, Matthew Yu, Mitesh Kumar Singh, Peizhao Zhang, Peter Vajda, Quentin Duval, Rohit Girdhar, Roshan Sumbaly, Sai Saketh Rambhatla, Sam Tsai, Samaneh Azadi, Samyak Datta, Sanyuan Chen, Sean Bell, Sharadh Ramaswamy, Shelly Sheynin, Siddharth Bhattacharya, Simran Motwani, Tao Xu, Tianhe Li, Tingbo Hou, Wei-Ning Hsu, Xi Yin, Xiaoliang Dai, Yaniv Taigman, Yaqiao Luo, Yen-Cheng Liu, Yi-Chiao Wu, Yue Zhao, Yuval Kirstain, Zecheng He, Zijian He, Albert Pumarola, Ali Thabet, Arsiom Sanakoyeu, Arun Mallya, Baishan Guo, Boris Araya, Breena Kerr, Carleigh Wood, Ce Liu, Cen Peng, Dimitry Vengertsev, Edgar Schonfeld, Elliot Blanchard, Felix Juefei-Xu, Fraylie Nord, Jeff Liang, John Hoffman, Jonas Kohler, Kaolin Fire, Karthik Sivakumar, Lawrence Chen, Licheng Yu, Luya Gao, Markos Georgopoulos, Rashel Moritz, Sara K. Sampson, Shikai Li, Simone Parmeggiani, Steve Fine, Tara Fowler, Vladan Petrovic, and Yuming Du. 2024. Movie Gen: A Cast of Media Foundation Models. arXiv:2410.13720 [cs.CV] <https://arxiv.org/abs/2410.13720>
- Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 652–660.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. arXiv:2112.10752 [cs.CV] <https://arxiv.org/abs/2112.10752>
- Johannes Lutz Schönberger and Jan-Michael Frahm. 2016. Structure-from-Motion Revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. 2024. MVDream: Multi-view Diffusion for 3D Generation. arXiv:2308.16512 [cs.CV] <https://arxiv.org/abs/2308.16512>
- Jaidev Shriram, Alex Trevithick, Lingjie Liu, and Ravi Ramamoorthi. 2025. Realm-Dreamer: Text-Driven 3D Scene Generation with Inpainting and Depth Diffusion. *International Conference on 3D Vision (3DV)*.
- Shitao Tang, Fuyang Zhang, Jiacheng Chen, Peng Wang, and Yasutaka Furukawa. 2023. MVDiffusion: Enabling Holistic Multi-view Image Generation with Correspondence-Aware Diffusion. arXiv (2023).
- Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. 2024. Visual Autoregressive Modeling: Scalable Image Generation via Next-Scale Prediction. In *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.), Vol. 37. Curran Associates, Inc., 84839–84865. [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/9a24e284b187f662681440ba15c416fb-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/9a24e284b187f662681440ba15c416fb-Paper-Conference.pdf)
- Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitry Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. 2024. SV3D: Novel Multi-view Synthesis and 3D Generation from a Single Image using Latent Video Diffusion. arXiv:2403.12008 [cs.CV] <https://arxiv.org/abs/2403.12008>
- Matthew Wallingford, Anand Bhattad, Aditya Kusupati, Vivek Ramanujan, Matt Deitke, Aniruddha Kembhavi, Roozbeh Mottaghi, Wei-Chiu Ma, and Ali Farhadi. 2024. From an Image to a Scene: Learning to Imagine the World from a Million 360° Videos. *Advances in Neural Information Processing Systems* 37 (2024), 17743–17760.
- Hanyang Wang, Fangfu Liu, Jiawei Chi, and Yueqi Duan. 2025. VideoScene: Distilling Video Diffusion Model to Generate 3D Scenes in One Step. arXiv:2504.01956 [cs.CV] <https://arxiv.org/abs/2504.01956>
- Peng Wang, Hao Tan, Sai Bi, Yinghao Xu, Fujun Luan, Kalyan Sunkavalli, Wenping Wang, Zexiang Xu, and Kai Zhang. 2023. PF-LRM: Pose-Free Large Reconstruction Model for Joint Pose and Shape Prediction. arXiv preprint arXiv:2311.12024 (2023).
- Qian Wang, Weiqi Li, Chong Mou, Xinhua Cheng, and Jian Zhang. 2024. 360DVD: Controllable Panorama Video Generation with 360-Degree Video Diffusion Model. arXiv preprint arXiv:2401.06578 (2024).
- Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. 2018. ESRGAN: Enhanced super-resolution generative adversarial networks. In *The European Conference on Computer Vision Workshops (ECCVW)*.
- Henrique Weber, Mathieu Garon, and Jean-François Lalonde. 2022. Editable Indoor Lighting Estimation. In *European Conference on Computer Vision (ECCV)*.
- Xinyue Wei, Kai Zhang, Sai Bi, Hao Tan, Fujun Luan, Valentin Deschaintre, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. 2024. MeshLRM: Large Reconstruction Model for High-Quality Mesh. arXiv preprint arXiv:2404.12385 (2024).
- Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Chunyi Li, Liang Liao, Annan Wang, Erli Zhang, Wenxiu Sun, Qiong Yan, Xiongkuo Min, Guangtai Zhai, and Weisi Lin. 2023. Q-Align: Teaching LLMs for Visual Scoring via Discrete Text-Defined Levels. arXiv preprint arXiv:2312.17090 (2023). Equal Contribution by Wu, Haoning and Zhang, Zicheng. Project Lead by Wu, Haoning. Corresponding Authors: Zhai, Guangtai and Lin, Weisi..
- Tong Wu, Shuai Yang, Ryan Po, Yinghao Xu, Ziwei Liu, Dahua Lin, and Gordon Wetzstein. 2025. Video World Models with Long-term Spatial Memory. arXiv:2506.05284 [cs.CV] <https://arxiv.org/abs/2506.05284>
- Tong Wu, Yu-Jie Yuan, Ling-Xiao Zhang, Jie Yang, Yan-Pei Cao, Ling-Qi Yan, and Lin Gao. 2024. Recent Advances in 3D Gaussian Splatting. arXiv:2403.11134 [cs.CV] <https://arxiv.org/abs/2403.11134>
- Desai Xie, Sai Bi, Zhixin Shu, Kai Zhang, Zexiang Xu, Yi Zhou, Sören Pirk, Arie Kaufman, Xin Sun, and Hao Tan. 2024. LRM-Zero: Training Large Reconstruction Models with Synthesized Data. arXiv:2406.09371 [cs.CV] <https://arxiv.org/abs/2406.09371>
- Shuai Yang, Jing Tan, Mengchen Zhang, Tong Wu, Yixuan Li, Gordon Wetzstein, Ziwei Liu, and Dahua Lin. 2024. LayerPano3D: Layered 3D Panorama for Hyper-Immersive Scene Generation. arXiv preprint arXiv:2408.13252 (2024).
- Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. 2024b. gsplat: An Open-Source Library for Gaussian Splatting. arXiv preprint arXiv:2409.06765 (2024). arXiv:2409.06765 [cs.CV] <https://arxiv.org/abs/2409.06765>
- Weicai Ye, Chenhao Ji, Zheng Chen, Junyao Gao, Xiaoshui Huang, Song-Hai Zhang, Wanli Ouyang, Tong He, Cairong Zhao, and Guofeng Zhang. 2024a. DiffPano: Scalable and Consistent Text to Panorama Generation with Spherical Epipolar-Aware Diffusion. arXiv:2410.24203 [cs.CV] <https://arxiv.org/abs/2410.24203>
- Hong-Xing Yu, Haoyi Duan, Charles Herrmann, William T. Freeman, and Jiajun Wu. 2024a. WonderWorld: Interactive 3D Scene Generation from a Single Image. arXiv:2406.09394 [cs.CV] <https://arxiv.org/abs/2406.09394>
- Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, et al. 2024b. LANGUAGE MODEL BEATS DIFFUSION-TOKENIZER IS KEY TO VISUAL GENERATION. In *12th International Conference on Learning Representations, ICLR 2024*.

- Wangbo Yu, Jinbo Xing, Li Yuan, Wenbo Hu, Xiaoyu Li, Zhipeng Huang, Xiangjun Gao, Tien-Tsin Wong, Ying Shan, and Yonghong Tian. 2024c. ViewCrafter: Taming Video Diffusion Models for High-fidelity Novel View Synthesis. arXiv:2409.02048 [cs.CV] <https://arxiv.org/abs/2409.02048>
- Fangneng Zhan, Changgong Zhang, Yingchen Yu, Yuan Chang, Shijian Lu, Feiying Ma, and Xuansong Xie. 2021. EMLight: Lighting Estimation via Spherical Distribution Approximation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- 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*.
- Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. 2024a. GS-LRM: Large Reconstruction Model for 3D Gaussian Splatting. *European Conference on Computer Vision* (2024).
- Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. 2024. *Open-Sora: Democratizing Efficient Video Production for All*. <https://github.com/hpcaitech/Open-Sora>
- Shijie Zhou, Zhiwen Fan, Dejie Xu, Haoran Chang, Pradyumna Chari, Tejas Bharadwaj, Suyu You, Zhangyang Wang, and Achuta Kadambi. 2024. DreamScene360: Unconstrained Text-to-3D Scene Generation with Panoramic Gaussian Splatting. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part VI* (Milan, Italy). Springer-Verlag, Berlin, Heidelberg, 324–342. doi:10.1007/978-3-031-72658-3\_19
- Chen Ziwen, Hao Tan, Kai Zhang, Sai Bi, Fujun Luan, Yicong Hong, Li Fuxin, and Zexiang Xu. 2024. Long-LRM: Long-sequence Large Reconstruction Model for Wide-coverage Gaussian Splats. *arXiv preprint 2410.12781* (2024).

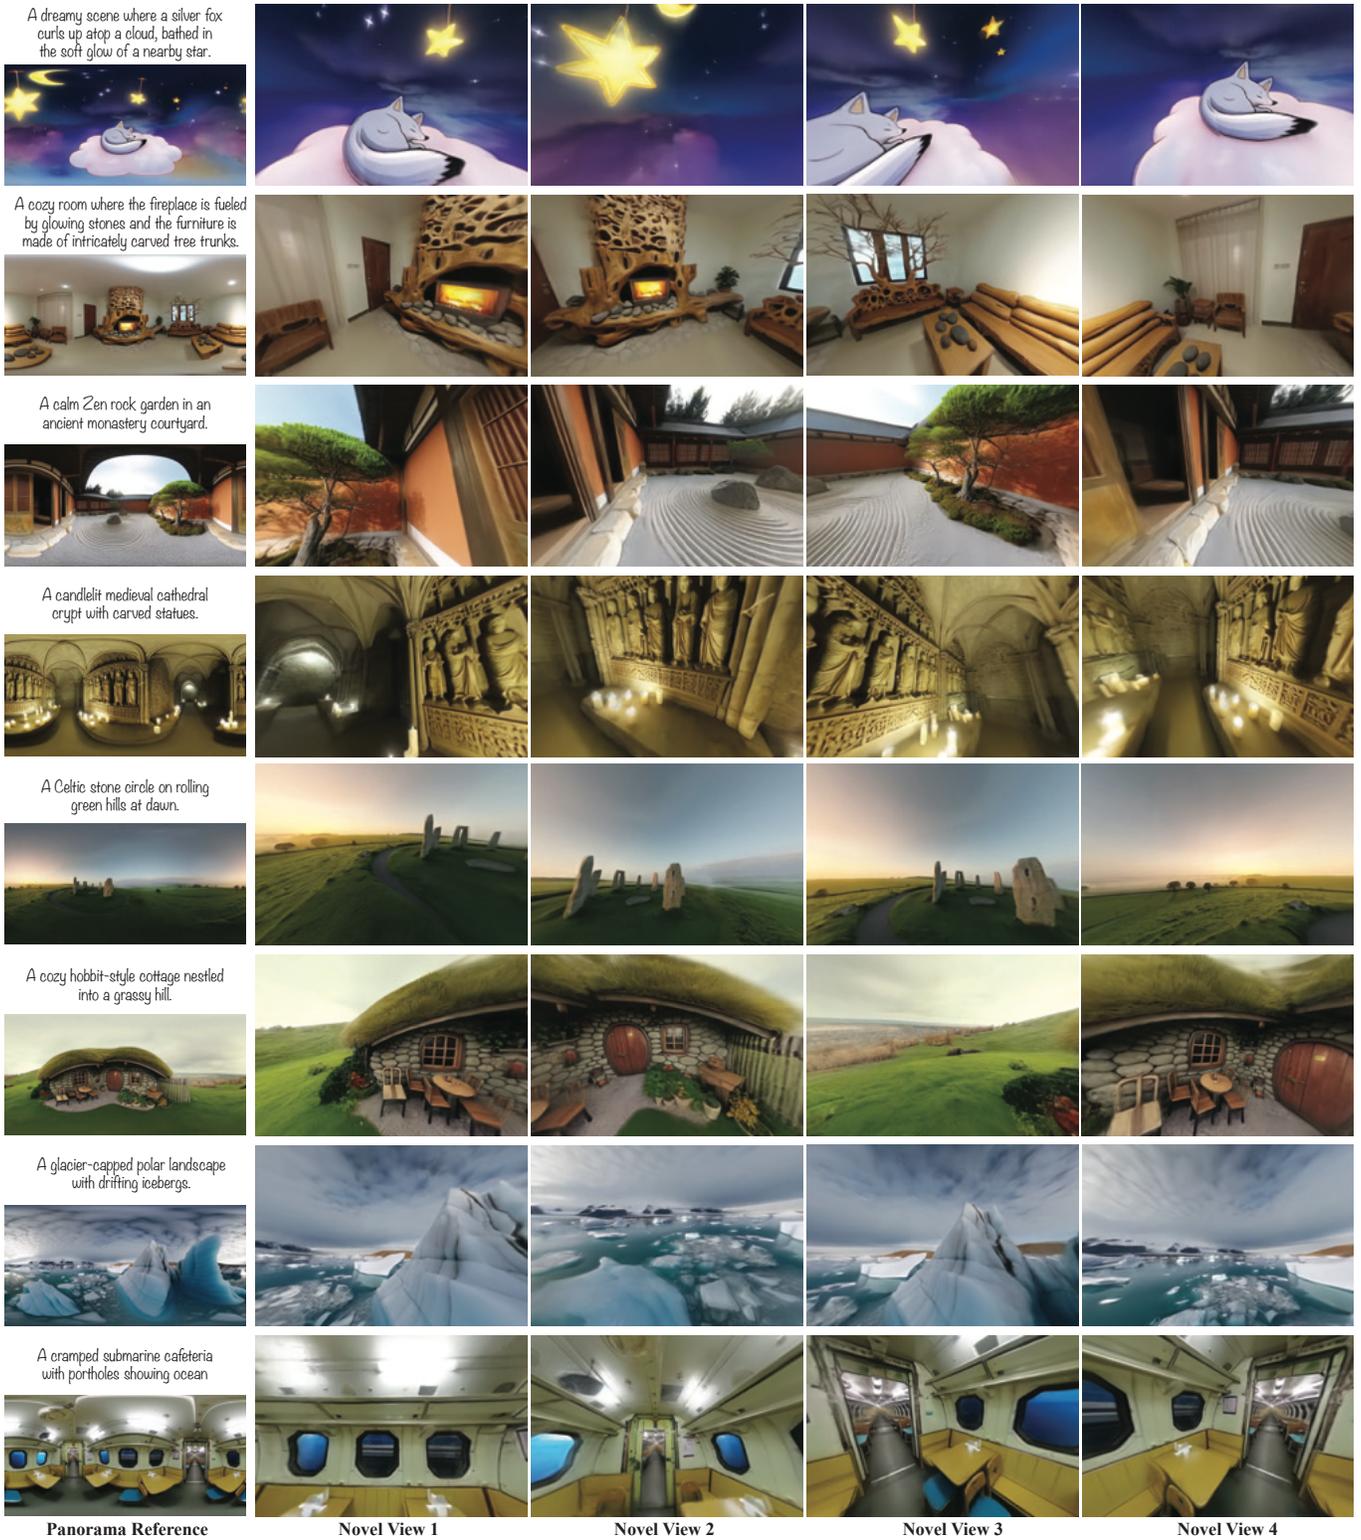


Fig. 6. **Scene generation results.** We present our scene generation results by showing (1) the first 360° image from our generated panoramic video as reference, and (2) novel video renderings (Novel View 1-4) from different spatial locations and viewpoints, demonstrating that our approach effectively synthesizes a traversable, 360° viewable scene. More rendered videos can be found in the supplementary documents.

A cottage in the woods where the walls are made of interwoven tree roots, and the floor is a carpet of moss.



A minimalist studio apartment with a bed, desk, and kitchenette.

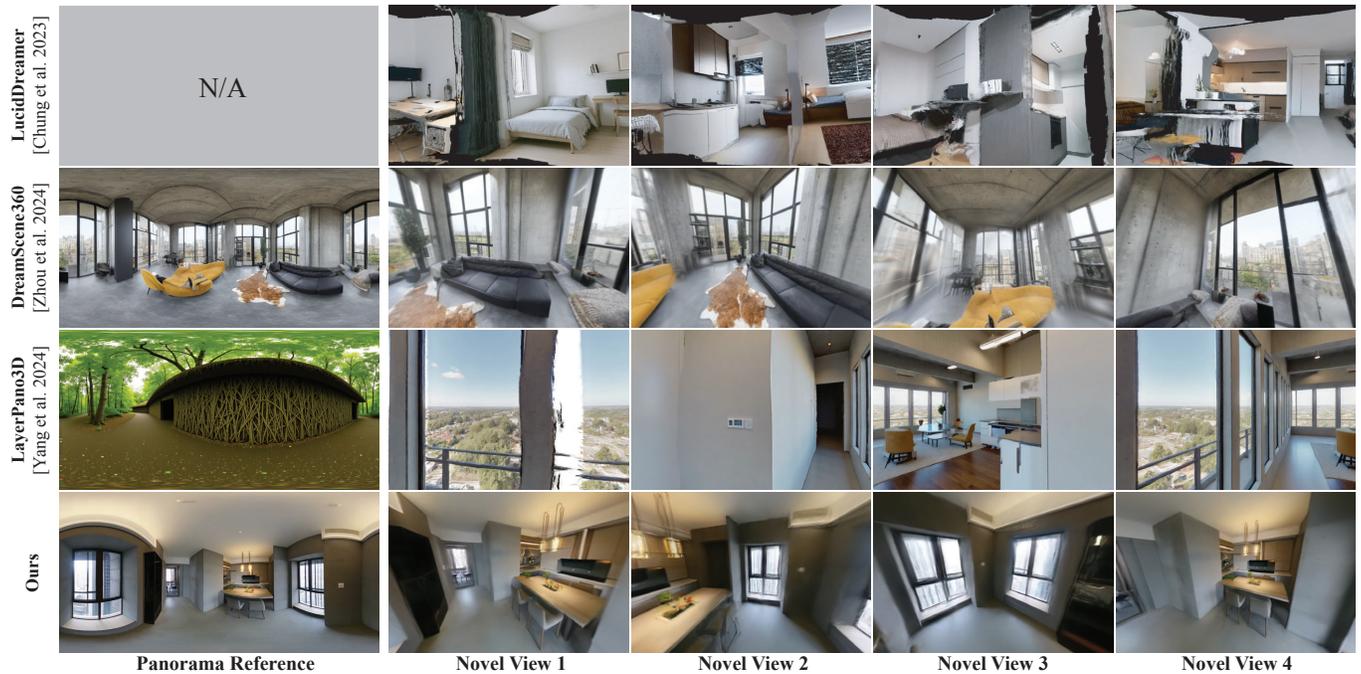


Fig. 7. **Comparison of scene generation results.** We compare our method with LucidDreamer [Chung et al. 2023], DreamScene360 [Zhou et al. 2024] and LayerPano3D [Yang et al. 2024] for text-based scene generation. LucidDreamer generates scenes progressively from a perspective images, synthesizing views one at a time. DreamScene360 and LayerPano3D both use a single panorama as input to an image diffusion model, shown in the first column for reference. For our method, we present the first frame of the generated 360° panoramic video as the reference image. While prior methods can produce visually reasonable results at the initial viewpoint, their range of motion is limited. When synthesizing novel viewpoints (Novel View 1–4) that deviate significantly from the starting position, their reconstructed Gaussian Splatting (GS) scenes exhibit severe artifacts, including blurriness, missing regions, and distorted geometries. As a result, global structural consistency cannot be held. In contrast, our method produces scenes that remain well-aligned with the input prompt and maintain high visual quality even at viewpoints far from the initial position.