

# WonderVerse: Video Geometric Correction and Enhancement for Extendable 3D Scene Generation

Hao Feng<sup>1\*</sup> Zhi Zuo<sup>1\*</sup> Jia-Hui Pan<sup>2</sup> Ka-Hei Hui<sup>3</sup> Qi Dou<sup>2</sup> Jingyu Hu<sup>2†</sup> Zhengzhe Liu<sup>1†</sup>  
<sup>1</sup>Lingnan University <sup>2</sup>The Chinese University of Hong Kong <sup>3</sup>AutoDesk Research

## Abstract

We introduce *WonderVerse*, a simple but effective framework for generating extendable 3D scenes. Unlike existing methods that rely on iterative depth estimation and image inpainting, often leading to geometric distortions and inconsistencies, *WonderVerse* leverages video generative foundation models to create highly immersive, geometrically coherent, extendable, and large-scale 3D environments. For enhanced scene completeness and user-driven expansion, an LLM-driven extension module enables semantic expansion and personalization of 3D scenes based on user intent. To bridge the video-3D domain gap and enhance video geometric fidelity, *WonderVerse* incorporates two key modules: a video geometric correction module, which filters geometrically inconsistent abnormal videos by analyzing camera trajectory smoothness, and a 3D geometric enhancement module, which enforces depth and normal consistency between the generated video and the reconstructed 3D scene. Extensive experiments on 3D scene generation demonstrate that our *WonderVerse*, with an elegant and simple pipeline, delivers extendable and highly realistic 3D scenes, markedly outperforming existing works that rely on more complex architectures. The code is available at <https://github.com/FH2070/WonderVerse>

**Keywords:** 3D Scene Generation, 3D Gaussian Splatting, Extendable Scene, Abnormal sequence detection.

## 1. Introduction

3D scene generation is a fundamental task in the vision community due to its wide-ranging applications in Virtual Reality, Mixed Reality (VR/MR), robotics, self-driving vehicles, and more. However, creating extendable and large-scale 3D scenes poses significant challenges. First, there is a lack of large-scale, high-quality datasets of 3D scenes. Second, generating realistic and geometrically accurate 3D scenes is inherently complex due to the unstructured nature of 3D space and the need to capture intricate object rela-

tionships and environmental context. Third, seamlessly extending 3D scenes while maintaining coherence remains a significant challenge.

To generate extendable 3D scenes, recent studies [19, 11, 59, 58] leverage 2D image generative models for 3D scene generation. These methods typically iteratively extend a given image to synthesize views from new camera poses. Typically, they first estimate a depth map from the input image, project it into 3D, and then use an image inpainting model to extend the scene from novel viewpoints. This iterative process builds and extends the 3D environment. However, the quality of scenes generated by these approaches remains limited. As shown in Figure 4, a primary issue is geometric distortion, arising from the scale ambiguity of single-view depth estimation and the lack of consistent camera parameters across iterations. Additionally, error accumulation is a concern due to the iterative application of depth estimation and inpainting during scene expansion. Furthermore, seams and discontinuities can appear between regions generated in different iterations. Overall, the iterative discrete-view-based generation pipeline inherently results in unsatisfactory scene quality.

Going beyond existing works that rely on image generation techniques, we propose a simple yet effective approach to leverage the recent advancements of LLM and video generative models [24, 8, 40, 60, 62, 53] to produce highly immersive and superior-quality 3D scenes without dedicated 3D datasets. Importantly, our contribution lies not in proposing new foundation models or reconstruction algorithms, but in demonstrating how recent video generative priors can be reliably harnessed for *extendable 3D scene generation*. Compared to image-based iterative pipelines, this design mitigates geometric distortion, error accumulation, and seam artifacts during iterative expansion. To achieve this, a natural initial thought is to combine video generative models with 3DGS (3D Gaussian Splatting) [23] for 3D scene generation. Nevertheless, this simple combination does not yield satisfactory results since the video generative model cannot faithfully capture the 3D geometric rules of the real world. First, current video generative models often lack the precise controllability required for directional extension in 3D space. Second, as indicated

\*Equal contribution.

†Corresponding authors.




Figure 1: WonderVerse can create large-scale, coherent, extendable, and high-quality 3D scenes from a text. It also allows users to interactively extend the 3D scene with text and can output the generated 3D scene with different representations, including a point cloud and a mesh.

by existing works [32, 7], AI-generated videos can be geometrically inconsistent across frames; thus, the contradictions among the frames can lead to the distortion of the created 3D scenes. Moreover, the inherent dimensionality reduction of projecting a 3D scene onto a 2D video inevitably leads to a loss of crucial 3D surface detail. This is because subtle geometric variations, which are essential for accurate 3D surface reconstruction, are often compressed or obscured during the video capture process, ultimately resulting in surface degradation in generated videos and inferior 3D surface quality.

To address these issues, we present WonderVerse, a simple but effective extendable 3D scene generative model by enhancing the geometric fidelity of video generative models. First, to make the generated scenes more complete, vivid, and adaptable to diverse user needs, we introduce an LLM-driven extension module. This component leverages large language models to semantically interpret and expand the generated 3D scenes, enabling the generation of richer contextual elements beyond the originally captured content, empowers users to flexibly customize or extend scenes based on high-level descriptions or creative intent, significantly enhancing interactivity and scalability of the generation pipeline. Second, we present an abnormal sequence detection module to filter out geometrically inconsistent generated videos. This design stems from our observation that COLMAP camera pose estimation effectively flags the geometric issue: inconsistent videos produce erratic camera poses and discontinuous trajectories. The module evaluates the continuity of the estimated camera trajectory to enhance

video geometric coherence. Third, we introduce a 3D geometric enhancement module to provide explicit depth and surface normal supervision during training. By enriching the generated videos with depth and normal maps, we encourage the resulting 3D scenes to align with the videos, enforcing consistency in both RGB values and depth/normal constraints. With consistent and geometrically enriched video sequences, we reconstruct the 3D scene. Note that our approach is compatible with various 3D reconstruction approaches for both efficient (DUS3R [41]) and high-quality (3DGS [23]) 3D scene generation.

As illustrated in Figure 1, our WonderVerse can create an extendable and highly immersive 3D scene from a piece of text. Furthermore, both qualitative and quantitative experimental results demonstrate that WonderVerse’s neat and simple pipeline leads to state-of-the-art, extendable, and highly realistic 3D scene generation. Our approach demonstrates that sophisticated results can be attained through our highly elegant design. **Codes will be released upon publication.**

In summary, the key contributions of WonderVerse are:

- We propose a simple yet effective approach to leverage video generative models for the extendable 3D scene generation task, surpassing image-based iterative pipelines.
- We design a LLM-driven extension module to extend scenes based on high-level descriptions, significantly enhancing interactivity and scalability of the generation pipeline.

- We present two novel modules to enhance the geometric fidelity of AI-generated videos: a geometric correction module with an abnormal sequence detection module to ensure geometric consistency and a 3D geometric enhancement module to improve the 3D surface quality.
- WonderVerse achieves state-of-the-art extendable and highly-realistic 3D scene generation through a highly neat and simple pipeline, significantly outperforming existing methods with far more complex architectures.

## 2. Related Works

In this section, we will briefly introduce the development of video generation for 3D reconstruction, extendable 3D scene generation, and normal-depth estimation. And we also introduce novel view synthesis (NVS) in supplementary material.

### 2.1. Image and Video Generative Models for MVS and 3D Reconstruction

Recent progress in foundation models for image [36] and video [5] generation has driven the exploration of their use in 3D reconstruction. Image generative models can be fine-tuned with multi-view data for novel-view synthesis and 3D reconstruction of objects [28, 30] and scenes [16, 46]. This strategy has been extended to video foundation models, enabling video generation with controlled camera motion for 3D reconstruction [4, 40, 26, 39, 50, 17, 42]. Nevertheless, these methods primarily target object-level or limited-scale reconstruction and cannot directly adapt to our extendable 3D scene generation task. Our objective is to generate extendable 3D scenes, requiring models with the imaginative ability to extend 3D scenes beyond the initial input.

### 2.2. Extendable 3D Scene Generation

Another branch of works explore extendable 3D scene generation. Existing works like [48, 49, 61, 12, 27] only focus on urban environments and cannot be generalized to other scenes, restricting their semantic scope. For unbounded nature scenes, methods like [10, 6] require multi-view data, while others [47, 29, 43] rely on 3D scene datasets. To avoid such data dependencies, Text2Room [19], SceneScape [14], VividDream [25] LucidDreamer [11], WonderJourney [59], WonderWorld [58] generate scenes iteratively with depth estimation and inpainting. Among existing approaches, **these methods are the most relevant ones to our work.** However, their iterative pipelines inherently suffer from geometric distortion, discontinuity, and error accumulation during scene expansion, ultimately hampering scene quality and motivating our WonderVerse framework. In this work, we seek to overcome the above challenges and leverage video generative models to generate more immersive and expandable scenes.

### 2.3. Normal and Depth Priors for 3D Generation

Recent work in monocular depth estimation [51, 52, 57, 20, 22, 9, 21] has shown that pre-trained models can reliably predict depth priors from images or videos. Recently, the integration of geometric priors, i.e., normal and depth, for 3D generation has emerged as an important strategy to alleviate the inherent ambiguities in reconstructing 3D structures. Existing methods [55, 2] focus on extracting surface normals directly from images, providing complementary geometric details that improve downstream tasks such as 3D reconstruction and scene understanding. In particular, [33, 15] aim to unify these approaches by jointly predicting depth and surface normals. Moreover, recent methods [30, 56] explicitly leverage depth and normal maps as supervision signals during training, improving the fidelity of generated 3D assets. While methods like [30, 56] achieve impressive results in object-level 3D generation, they rely on ground truth depth and normal maps, which are expensive to obtain for 3D scenes. This problem hinders their direct extension to scene generation. To address this limitation, our approach leverages pre-trained models [52, 55] to extract robust depth and normal priors, which serve as foundational guidance for 3D scene generation while bypassing the need for ground-truth annotations.

## 3. Methods

In this section, we will introduce our *WonderVerse* which is illustrated in Figure 2. Given a text description as input, WonderVerse first generates a video through the text-guided video generation and extension module (Figure 2 (a)), which continuously presents the scene and extends it to the left and right sides. To ensure the geometric consistency of the video sequence, the Geometric Correction module (Figure 2 (b)) estimates the camera pose sequence, detects abnormal sequences by identifying pose discontinuities, and regenerates them. Simultaneously, the Geometric Enhancement module (Figure 2 (c)) predicts depth maps and normal maps for the video frames to provide important geometric prior information. Finally, the 3D Scene Reconstruction module (Figure 2 (d)) utilizes the geometrically corrected and enhanced video sequences for 3D scene reconstruction, thereby building the complete scene.

### 3.1. Video Generation and LLM Driven Extension

Most existing works on extendable 3D scene generation rely on iterative image generation. Such a design can hinder the quality of the generated results due to the significant domain gap between images and 3D representations. Inspired by the rapid development of video generative models like [24], we reconsider the problem by leveraging the superior generative capability of such models. Therefore, we propose to address the extendable 3D scene generation task through the lens of video generation.

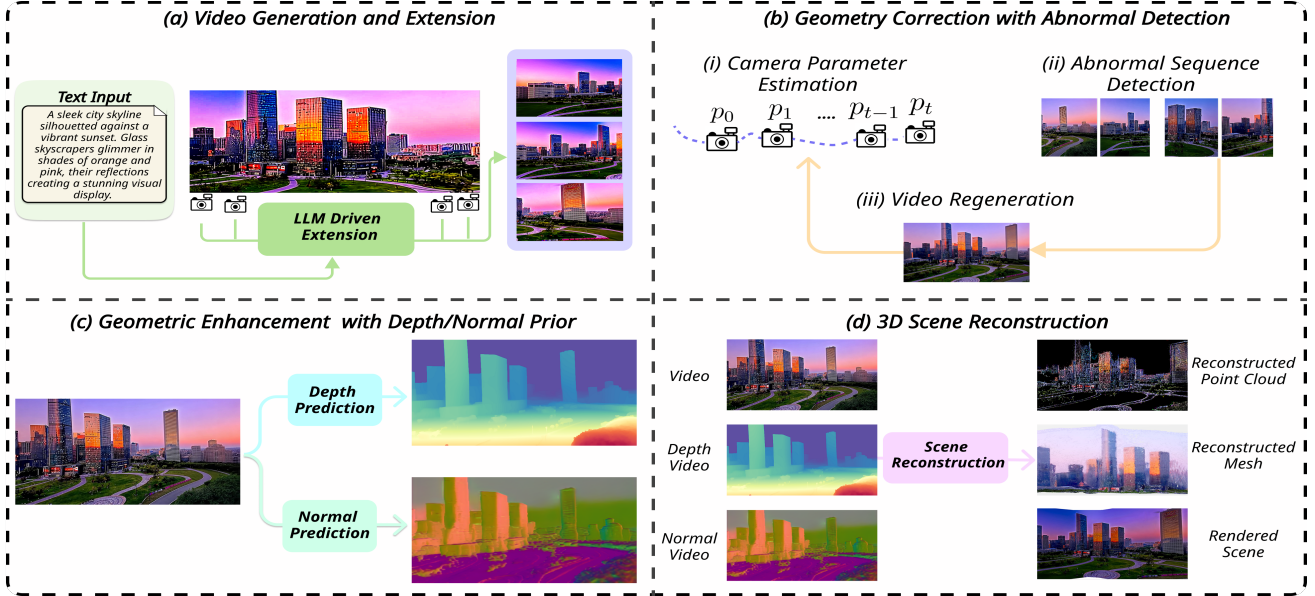


Figure 2: **Illustration of our WonderVerse.** This framework includes: (a) a text-guided video generation and LLM driven extension module that produces a video of a scene circularly in a continuous shot, followed by extensions in any direction; (b) geometry correction module with abnormal detection, designed to maintain cross-frame 3D geometric coherence in video sequences; (c) a geometric enhancement module that integrates pre-trained depth/normal prediction models to generate depth/normal maps, establishing robust geometric priors and improving the subsequent 3D scene reconstruction; and (d) a 3D scene reconstruction module that processes generated video sequences alongside predicted depth/normal maps as input, simultaneously producing both photorealistic renderings and geometrically consistent 3D reconstructions.

Given a text prompt  $P$ , we use text-guided video generation model [24] to obtain an initial video  $V_{\text{init}} \in \mathbb{R}^{T \times H \times W \times 3}$ , where  $H$  and  $W$  denote the height and width of each frame, respectively.  $T$  denotes the length of the video and 3 denotes RGB channels. To extend the scene, we further perform video-based scene extensions by extending the initial video. To maintain style consistency, we use the first frame  $V_{\text{init}}^{t=1}$  and the last frame  $V_{\text{init}}^{t=T}$  of the initial video as references and generate new videos to extend the scene in any direction. Here we use the left and right extensions as an example. We generate two new videos  $V_{\text{extend}}^{\text{left}} \in \mathbb{R}^{T' \times H \times W \times 3}$  and  $V_{\text{extend}}^{\text{right}} \in \mathbb{R}^{T' \times H \times W \times 3}$ . Then, we combine them into an extended video  $\{V_{\text{extend}}^{\text{left}}, V_{\text{init}}, V_{\text{extend}}^{\text{right}}\}$  with length  $T + 2T'$ . The video extension can be applied repeatedly  $n$  times, taking  $V_{\text{input}}$  as the initial video for the next iteration, further enriching the scene and yielding  $V_{\text{input}} = \{V_{\text{extend}}^{\text{left}_n}, \dots, V_{\text{init}}, \dots, V_{\text{extend}}^{\text{right}_{n-1}}, V_{\text{extend}}^{\text{right}_n}\} = \{V_1, V_2, \dots, V_{2n+1}\}$ . Figure 2 (a) illustrates two iterations of video extension. To achieve reasonable and natural scene expansion, we propose an LLM-driven extension module that generates relative camera trajectories based on user-provided prompts. Specifically, given a textual scene description, the module predicts plausible camera motions to explore newly imagined or extended areas of the scene. Furthermore, our system supports iterative language-based in-

teraction, allowing users to progressively expand the scene by issuing high-level natural language commands. In each iteration, the LLM generates a new camera trajectory segment conditioned on both the current scene context and the user’s intent, enabling flexible and controllable scene expansion. More details can be found in Experiments Section 4.

Despite the high quality of AI-generated videos, geometric inconsistencies across frames can still occur, which significantly degrades downstream 3D reconstruction. We empirically observe a strong correlation between visually evident geometric failures and discontinuous camera trajectories estimated by Structure-from-Motion (SfM). Motivated by this observation, we introduce an abnormal sequence detection module that acts as a lightweight *sampling/filtration* mechanism (rather than a theoretical guarantee of geometric correctness) to discard clearly degenerate video segments before reconstruction.

First, we perform *camera pose estimation* for each frame in the generated videos  $\{V_1, V_2, \dots, V_{2n+1}\}$  using COLMAP (SfM), producing an estimated camera trajectory per video segment. Second, we *detect abnormal sequences* by checking trajectory discontinuities: a frame is marked as discontinuous if its translation jump exceeds  $T_1$  or its rotation change exceeds  $T_2$  (relative to the previous frame). If

any discontinuity is detected, the corresponding video segment is treated as abnormal. Finally, we perform *geometric correction* via *full-segment regeneration*: once an abnormal segment is detected, we regenerate the *entire* video segment (instead of locally repairing a few frames or partially patching trajectories). Each regeneration is sampled independently with a new random seed, which avoids local patching artifacts and prevents error accumulation across retries.

In practice, this simple strategy is effective and does not lead to excessive retry loops: as reported in Sec. 3.2, the average number of additional trials is only 0.25. We also note a limitation: SfM-based pose estimation can miss subtle geometric warping or occasionally produce false positives; therefore, our module is not intended to “fix” the video generator, but to reliably filter out failure cases that would inevitably break the 3D reconstruction stage. An algorithm outlining the process is provided in the supplemental material.

### 3.2. Geometric Enhancement with Depth and Normal Priors

While reducing the 2D-to-3D domain gap compared to static images, videos still lack essential 3D geometric information like depth and surface normals. To address this issue, we introduce a video geometric enrichment module that predicts depth and normals from video frames to guide 3D reconstruction training. We utilize Depth Anything V2 [52] for accurate monocular depth estimation, since it is effective even in challenging scenes and offers strong geometric cues. Besides, to improve 3D detail and surface refinement, we apply the Dense Prediction Transformer (DPT) [34] with a hybrid backbone to predict per-frame surface normals, yielding smooth and consistent normals. These predicted depths and normals serve as explicit geometric supervision, encouraging geometrically consistent 3D scene reconstruction and enhancing surface quality and detail.

### 3.3. Geometric-Aware 3D Scene Reconstruction

With the geometrically validated and enhanced video sequences and their camera poses, we then optimize a 3DGS [23] representation of the scene. Please refer to the supplementary material for preliminaries of 3DGS.

**Depth from 3DGS** To incorporate the geometric-aware optimization, we estimate per-pixel depth rendered using the discrete volume rendering approximation, similar to color rendering of 3DGS:

$$\hat{D} = \sum_{i \in N} d_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (1)$$

where  $d_i$  is the  $i^{\text{th}}$  Gaussian depth coordinate in view space and  $\alpha$  is opacity. To resolve scale ambiguity in monocular depth estimation, we align its scale with that of the 3D

Gaussian Splatting (3DGS) depth, as computed by Equation 1. This alignment is achieved by solving for a per-image scale  $a$  and shift  $b$  using linear regression:

$$\hat{a}, \hat{b} = \operatorname{argmin}_{a,b} \sum_{ij} \|(a\hat{D}_{ij} + b) - D_{ij}\|_2^2, \quad (2)$$

where  $\hat{D}_{ij}$  and  $D_{ij}$  denote corresponding per-pixel depth values from 3DGS and monocular depth estimation.

**Surface Normal from 3DGS** For surface normal estimation, the minimal axis of each Gaussian serves as an approximation of the normal direction, which can be formulated as:

$$\hat{\mathbf{n}}_i = \operatorname{normalize}(\operatorname{argmin}(s_{i1}, s_{i2}, s_{i3})), \quad (3)$$

where  $s_{ij}$  indicates the scale in the  $j$  direction of Gaussian  $i$ .

Normals are then transformed into camera space using the current camera transform and alpha composited according to the rendering equation to provide a per-pixel normal estimate:

$$\hat{\mathbf{N}} = \sum_{i \in N} \hat{\mathbf{n}}_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j). \quad (4)$$

**Optimization** The optimization process is guided by a composite loss function that combines photometric and geometric cues. We adopt the original photometric loss [23], which minimizes the difference between the rendered color and the ground truth color of training images. In addition, we regularize the optimization process using the depth maps and surface normals obtained from the geometry enhancement module. The overall loss function is:

$$\mathcal{L} = \mathcal{L}_{\hat{C}} + \lambda_d \mathcal{L}_{\hat{D}} + \lambda_n \mathcal{L}_{\hat{\mathbf{N}}}, \quad (5)$$

where  $\lambda_d$  and  $\lambda_n$  are weighting factors to balance the contribution of each term.

To leverage the estimated depth, we minimize a depth loss term:

$$\mathcal{L}_{\hat{D}} = \frac{1}{|\hat{D}|} \sum \log(1 + \|\hat{D} - D\|_1), \quad (6)$$

where  $\hat{D}$  represents the rendered depth from the 3D Gaussians, and  $|\hat{D}|$  indicates the total number of pixels in  $\hat{D}$ . We found the logarithmic variant of the depth loss yields better results than linear or quadratic penalties, following [38].

We further incorporate a normal consistency term, minimizing the following loss function:

$$\mathcal{L}_{\hat{\mathbf{N}}} = \frac{1}{|\hat{\mathbf{N}}|} \sum \|\hat{\mathbf{N}} - \mathbf{N}\|_1 + \lambda_s \sum_{i,j} (|\nabla_i \hat{\mathbf{N}}_{i,j}| + |\nabla_j \hat{\mathbf{N}}_{i,j}|), \quad (7)$$

where  $\hat{N}$  represents the rendered surface normal and  $\lambda_s$  is a loss weight. The first term promotes alignment between the rendered and estimated normals, while the second enforces smoothness within the rendered normal field.

#### Compatibility on different 3D reconstruction methods

By optimizing the 3D Gaussians with both RGB, depth, and normal cues, WonderVerse achieves high-fidelity 3D scene reconstruction that is both photorealistic and geometrically accurate.

In contrast to prior works that rely on point clouds [59, 19, 11] or less common data formats such as layered images [58] as intermediate representations, WonderVerse directly generates images and corresponding camera poses. This design significantly enhances compatibility, allowing us to leverage well-established 3D reconstruction methods that natively accept image and pose inputs, such as the recent DUST3R [41]. This streamlined approach enables efficient 3D scene generation, with the potential for further improvements as reconstruction techniques advance.

## 4. Experiments

### 4.1. Experimental Settings

**Implementation Details** Starting with a text prompt, WonderVerse generates an initial video and then extension videos. Taking the left and right extension as an example, the first and the last frames, along with a text prompt for each, are fed into the video generative model to create two extended videos, repeating this process twice. We set hyperparameters  $T_1$  and  $T_2$  as 5 and 0.5 through cross validation. Each text prompt starts with the template: “aerial shot, soft lighting, around left [or right], realistic, high-quality, displaying [scene description].” For initial scene generation, we used the Hunyuan video model [24], and for extensions, we adopted Gen-3 Alpha [37] due to its superior camera control capability. For the LLM-driven extension module, we employed GPT-4o mini [1] to generate camera trajectories conditioned on user-defined scene descriptions. The prompt used is as follows: “Provide the direction and angle of camera pose changes based on the given scene description given by the user. The answer only includes the corresponding angles of up, down, left, right, and the distance of forward and backward movement of four times extend, and it is required that the perspective expansion cannot be repeated. Just give the four answers, format of each answer is: up or down  $X$ , left or right  $X$ , forward or back  $X$ . ( $X$  is the angle or distance). If the extension prompt is given, just give one answer.”

**Evaluation metrics** Following recent work [58], we evaluate our scenes using CLIP-Score (CLIP-S) [18], Q-Align [45], and NIQE [31]. CLIP-S measures text-image alignment by comparing CLIP feature embeddings. Q-

Align, trained with human scoring patterns and instruction tuning, assesses the human-aligned generation quality. NIQE, a reference-free image quality metric, evaluates quality degradation by comparing spatial features to a pristine natural image model.

### 4.2. Our Results of Extendable 3D Scenes

In Figure 3, we present compelling examples of extendable indoor and outdoor 3D scenes generated by WonderVerse. Figure 3 (a, b) illustrate the remarkable realism and rich detail achieved in our created scenes, with flawless geometric integrity. Also, the scenes faithfully reflect the text description. This shows the strength of WonderVerse in making highly immersive and extendable 3D scenes. Besides, we also demonstrate that users can generate extendable 3D scenes iteratively using customized prompts in supplementary material, highlighting the user-controllability of WonderVerse.

Please also refer to the supplementary material for additional generative results across diverse scenes and styles of WonderVerse.

### 4.3. Comparison with Existing Works

**Baseline Approaches** We compare our WonderVerse with recent works WonderWorld [58] and LucidDreamer [11] both qualitatively and quantitatively. Since these two baselines need an image as input while our method only needs a text prompt as input, for a fair comparison, we take a random frame of our generated video and feed it into their models for 3D scene generation. Note that WonderWorld [58]’s layered image representation, including a sky layer, makes it less suited for indoor scene generation. Consequently, our primary comparative analysis focuses on outdoor scenes.

Our comparison focuses on *extendable 3D scene generation pipelines* that share the same task setting of *iterative scene expansion* (e.g., LucidDreamer [11] and WonderWorld [58]). Existing methods in this category are predominantly image-driven and rely on iterative image inpainting, making them the most relevant baselines under the same problem formulation. In contrast, video-based approaches (e.g., CameraCtrl [17], SV4D [50], GEN3C [35], GeoVideo [3]) mainly target controlled camera motion or bounded video-to-3D reconstruction along a fixed trajectory. WonderVerse addresses a different setting: text-driven and extendable scene generation beyond the original video span; therefore these methods are largely orthogonal to our objective and do not directly support extendable/interactive scene generation.

**Qualitative Comparison** Figure 4 visually demonstrates WonderVerse’s clear superiority over existing methods. As shown in Fig 4 (a), LucidDreamer’s generated scene is obscured by odd frame-like artifacts, and lacks de-



Figure 3: WonderVerse generates large-scale, extendable, coherent, and high-fidelity 3D scenes, both indoors and outdoors. Dashed lines show the camera’s direction during scene extension.




Figure 4: Qualitative comparison with existing works. Text prompt: A small house decorated for Christmas. We can see that LucidDreamer cannot generate a natural scene. Text2Room and WonderWorld suffer from geometric discontinuities, which lead to fragmented scenes and visually inconsistent outputs. In contrast, our WonderVerse can generate high-quality, geometric continuities and natural scenes.




Figure 5: Ablation studies on our geometry enhancement module.

tail. WonderWorld (Figure 4 (b, c)) suffers from geometric discontinuities, as indicated by the visibly stitched and unnatural scene. In contrast, WonderVerse (Figure 4 (d)) produces a satisfactory and geometrically sound 3D scene. The geometric integrity is maintained seamlessly across the scene, addressing the limitations apparent in both LucidDreamer’s quality and WonderWorld’s geometric consistency. Please refer to the supplementary material for additional comparison results with existing works.

**Quantitative Comparison.** We quantitatively compare WonderVerse to prior works LucidDreamer [11] and WonderWorld [58], with results shown in Table 1 for both outdoor and indoor scenes. Our method demonstrates state-of-the-art scene generation, outperforming existing methods in semantic alignment, structural consistency, and perceptual quality. Notably, for outdoor scenes, WonderVerse achieves the best CLIP-S score (0.9219), indicating superior semantic alignment, along with SOTA Q-align and NIQE scores,

demonstrating high generation quality. Similarly, for indoor scenes, WonderVerse significantly surpasses LucidDreamer across all metrics, achieving a new SOTA CLIP-S score of 0.9639, again affirming excellent text-scene semantic alignment and high-quality 3D indoor scene generation. Note that WonderWorld cannot generate indoor scenes due to its sky layer, so we do not compare with it in the indoor setting.

Table 1: Evaluation on novel view renderings.

Scene Type	Method	Reference	CLIP-S (↑)	Q-align (↑)	NIQE (↑)
Outdoor	LucidDreamer	arXiv24	0.4200	2.6704	0.7786
	WonderWorld	CVPR25	0.7817	2.8302	0.7127
	WonderVerse (ours)	-	<b>0.9219</b>	<b>3.2229</b>	<b>0.7806</b>
Indoor	LucidDreamer	arXiv24	0.4633	3.9870	0.6655
	WonderVerse (ours)	-	<b>0.9639</b>	<b>4.2680</b>	<b>0.7119</b>




Figure 6: Ablation study on our abnormal sequence detection module. (a) W/o abnormal sequence detection; (b) W abnormal sequence detection

#### 4.4. Ablation Studies

In this section, we examine our key modules, including abnormal sequence detection and geometric enhancement. In the supplementary material, we also examine scene extension module and compare different 3D representation methods 3DGS [23] and DUS3R [41] to demonstrate the compatibility of our approach.

##### 4.4.1 Abnormal Sequence Detection

In this section, we study the effectiveness of our abnormal sequence detection. As illustrated in Figure 6, without our detection algorithm, noticeable artifacts degrade the generated 3D scene due to the geometric inconsistency of the generated video. In contrast, when abnormal sequences are detected and addressed, the generative quality is much higher. The improvement is further demonstrated by quantitative evaluations, please refer to the supplementary material for more details.

##### 4.4.2 Geometric Enhancement

As illustrated in Figure 5, we present 3D mesh reconstructions with and without the geometry enhancement module.

Without depth and normal regularization, the tower exhibits considerable surface distortion, and the bowl displays noticeable artifacts. Conversely, incorporating geometric priors significantly improves surface quality; for instance, the tower’s intricate details are faithfully generated, and the glass bowl’s surface is much smoother. The normal maps further demonstrate that the geometry enhancement module enables a cleaner 3D surface and creates complex geometric details.

## 5. Limitations and Future Work

Our WonderVerse framework, while effective, has certain limitations. First, The quality of our simulated 3D environments is inherently limited by existing video generation models. Second, our abnormal sequence detection is a heuristic filtering mechanism based on SfM (COLMAP) pose smoothness. It can miss subtle geometric warping or trigger false positives under challenging conditions (e.g., repetitive textures or weak features). Finally, our model mainly focuses on static 3D scene generation due to the use of 3DGS and DUS3R for reconstruction. To overcome this, an exciting avenue for future research is to combine with recent advancements in dynamic scene representation [13, 54, 44] to pave the way for generating dynamic and animated 3D environments. We also provide representative failure cases and analysis in the supplementary material.

## 6. Conclusion

In this paper, we introduce WonderVerse, a simple yet effective framework for 3D scene generation. Unlike previous methods that rely on image-based generation schemes, we leverage the world-level priors embedded within video generative models to improve the generation quality. Our contribution lies in demonstrating how recent video generative priors can be reliably harnessed for extendable 3D scene generation, a setting where existing methods struggle. Our work bridges video generation and usable 3D scene reconstruction, enabling extensible scene generation. In particular, prior extendable scene generation pipelines are predominantly image-based and often suffer from geometric distortion, error accumulation, and seam artifacts during iterative expansion. Equipped with a video geometric correction module that detects abnormal videos due to extreme camera pose changes and a 3D geometric enhancement module that ensures depth and normal consistency, WonderVerse can generate extendable and highly realistic 3D scenes, markedly outperforming existing works with more complex architectures. Please refer to the supplementary material for discussions of limitations.

## Acknowledgement

This work has been supported by National Natural Science Foundation of China (Category C) fund code 62506149, Lingnan University StartUp Grant fund code: SUG-001/2526, and Faculty Research Grant fund code:106106 and 106119.

## References

- [1] O. J. Achiam, S. Adler, and S. Agarwal. Gpt-4 technical report. 2023. **6**
- [2] G. Bae and A. J. Davison. Rethinking inductive biases for surface normal estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9535–9545, 2024. **3**
- [3] Y. Bai, S. Fang, C. Yu, F. Wang, and Q. Huang. Geovideo: Introducing geometric regularization into video generation model. *arXiv preprint arXiv:2512.03453*, 2025. **6**
- [4] O. Bar-Tal, H. Chefer, O. Tov, C. Herrmann, R. Paiss, S. Zada, A. Ephrat, J. Hur, G. Liu, A. Raj, et al. Lumiere: A space-time diffusion model for video generation. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–11, 2024. **3**
- [5] A. Blattmann, T. Dockhorn, S. Kulal, D. Mendelevitch, M. Kilian, D. Lorenz, Y. Levi, Z. English, V. Voleti, A. Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. **3**
- [6] L. Chai, R. Tucker, Z. Li, P. Isola, and N. Snavely. Persistent nature: A generative model of unbounded 3d worlds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20863–20874, 2023. **3**
- [7] C. Chang, Z. Liu, X. Lyu, and X. Qi. What matters in detecting ai-generated videos like sora? *arXiv preprint arXiv:2406.19568*, 2024. **2**
- [8] S. Chen, C. Ge, Y. Zhang, Y. Zhang, F. Zhu, H. Yang, H. Hao, H. Wu, Z. Lai, Y. Hu, T.-C. Lin, S. Zhang, F. Li, C. Li, X. Wang, Y. Peng, P. Sun, P. Luo, Y. Jiang, Z. Yuan, B. Peng, and X. Liu. Goku: Flow based video generative foundation models. *arXiv preprint arXiv:2502.04896*, 2025. **1**
- [9] S. Chen, H. Guo, S. Zhu, F. Zhang, Z. Huang, J. Feng, and B. Kang. Video depth anything: Consistent depth estimation for super-long videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025. **3**
- [10] Z. Chen, G. Wang, and Z. Liu. Scenedreamer: Unbounded 3d scene generation from 2d image collections. *IEEE transactions on pattern analysis and machine intelligence*, 45(12):15562–15576, 2023. **3**
- [11] J. Chung, S. Lee, H. Nam, J. Lee, and K. M. Lee. Lucidreamer: Domain-free generation of 3d gaussian splatting scenes, 2023. **1, 3, 6, 7**
- [12] J. Deng, W. Chai, J. Guo, Q. Huang, W. Hu, J.-N. Hwang, and G. Wang. Citygen: Infinite and controllable 3d city layout generation. *arXiv preprint arXiv:2312.01508*, 2023. **3**
- [13] Y. Duan, F. Wei, Q. Dai, Y. He, W. Chen, and B. Chen. 4d-rotor gaussian splatting: Towards efficient novel view synthesis for dynamic scenes. In *International Conference on Computer Graphics and Interactive Techniques*, 2024. **8**
- [14] R. Fridman, A. Abecasis, Y. Kasten, and T. Dekel. Scenescape: Text-driven consistent scene generation. *Advances in Neural Information Processing Systems*, 36:39897–39914, 2023. **3**
- [15] X. Fu, W. Yin, M. Hu, K. Wang, Y. Ma, P. Tan, S. Shen, D. Lin, and X. Long. Geowizard: Unleashing the diffusion priors for 3d geometry estimation from a single image. In *European Conference on Computer Vision*, pages 241–258. Springer, 2024. **3**
- [16] R. Gao, A. Holynski, P. Henzler, A. Brussee, R. Martin-Brualla, P. Srinivasan, J. T. Barron, and B. Poole. Cat3d: Create anything in 3d with multi-view diffusion models. *arXiv preprint arXiv:2405.10314*, 2024. **3**
- [17] H. He, Y. Xu, Y. Guo, G. Wetzstein, B. Dai, H. Li, and C. Yang. Cameractrl: Enabling camera control for text-to-video generation. *arXiv preprint arXiv:2404.02101*, 2024. **3, 6**
- [18] J. Hessel, A. Holtzman, M. Forbes, R. L. Bras, and Y. Choi. CLIPScore: a reference-free evaluation metric for image captioning. In *EMNLP*, 2021. **6**
- [19] L. Höllein, A. Cao, A. Owens, J. Johnson, and M. Nießner. Text2room: Extracting textured 3d meshes from 2d text-to-image models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7909–7920, 2023. **1, 3, 6**
- [20] M. Hu, W. Yin, C. Zhang, Z. Cai, X. Long, H. Chen, K. Wang, G. Yu, C. Shen, and S. Shen. Metric3d v2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. **3**
- [21] W. Hu, X. Gao, X. Li, S. Zhao, X. Cun, Y. Zhang, L. Quan, and Y. Shan. Depthcrafter: Generating consistent long depth sequences for open-world videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025. **3**
- [22] B. Ke, A. Obukhov, S. Huang, N. Metzger, R. C. Daudt, and K. Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9492–9502, 2024. **3**
- [23] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. **1, 2, 5, 8**
- [24] W. Kong. Hunyuanvideo: A systematic framework for large video generative models, 2024. **1, 3, 4, 6**
- [25] Y.-C. Lee, Y.-T. Chen, A. Wang, T.-H. Liao, B. Y. Feng, and J.-B. Huang. Vividdream: Generating 3d scene with ambient dynamics. *arXiv preprint arXiv:2405.20334*, 2024. **3**
- [26] H. Li, H. Shi, W. Zhang, W. Wu, Y. Liao, L. Wang, L.-h. Lee, and P. Y. Zhou. Dreamscene: 3d gaussian-based text-to-3d scene generation via formation pattern sampling. In *European Conference on Computer Vision*, pages 214–230. Springer, 2024. **3**
- [27] C. H. Lin, H.-Y. Lee, W. Menapace, M. Chai, A. Siarohin, M.-H. Yang, and S. Tulyakov. Infinicity: Infinite-scale city synthesis. In *Proceedings of the IEEE/CVF international*

- conference on computer vision*, pages 22808–22818, 2023. 3
- [28] R. Liu, R. Wu, B. Van Hoorick, P. Tokmakov, S. Zakharov, and C. Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9298–9309, 2023. 3
- [29] Y. Liu, X. Li, X. Li, L. Qi, C. Li, and M.-H. Yang. Pyramid diffusion for fine 3d large scene generation. In *European Conference on Computer Vision*, pages 71–87. Springer, 2024. 3
- [30] X. Long, Y.-C. Guo, C. Lin, Y. Liu, Z. Dou, L. Liu, Y. Ma, S.-H. Zhang, M. Habermann, C. Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. pages 9970–9980, 2024. 3
- [31] A. Mittal, R. Soundararajan, and A. C. Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2013. 6
- [32] Y. Qin, Z. Shi, J. Yu, X. Wang, E. Zhou, L. Li, Z. Yin, X. Liu, L. Sheng, J. Shao, et al. Worldsimbench: Towards video generation models as world simulators. *arXiv preprint arXiv:2410.18072*, 2024. 2
- [33] L. Qiu, G. Chen, X. Gu, Q. Zuo, M. Xu, Y. Wu, W. Yuan, Z. Dong, L. Bo, and X. Han. Richdreamer: A generalizable normal-depth diffusion model for detail richness in text-to-3d. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9914–9925, 2024. 3
- [34] R. Ranftl, A. Bochkovskiy, and V. Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12179–12188, 2021. 5
- [35] X. Ren, T. Shen, J. Huang, H. Ling, Y. Lu, M. Nimier-David, T. Müller, A. Keller, S. Fidler, and J. Gao. Gen3c: 3d-informed world-consistent video generation with precise camera control. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6121–6132, 2025. 6
- [36] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [37] Runway. <https://runwayml.com/research/introducing-gen-3-alpha>. 2024. 6
- [38] M. Turkulainen, X. Ren, I. Melekhov, O. Seiskari, E. Rahtu, and J. Kannala. Dn-splatter: Depth and normal priors for gaussian splatting and meshing. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2421–2431. IEEE, 2025. 5
- [39] V. Voleti, C.-H. Yao, M. Boss, A. Letts, D. Pankratz, D. Tochilkin, C. Laforte, R. Rombach, and V. Jampani. Sv3d: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. In *European Conference on Computer Vision*, pages 439–457. Springer, 2024. 3
- [40] C. Vondrick, H. Pirsiavash, and A. Torralba. Generating videos with scene dynamics. *Advances in neural information processing systems*, 29, 2016. 1, 3
- [41] S. Wang, V. Leroy, Y. Cabon, B. Chidlovskii, and J. Revaud. Dust3r: Geometric 3d vision made easy. In *CVPR*, 2024. 2, 6, 8
- [42] Z. Wang, Z. Yuan, X. Wang, Y. Li, T. Chen, M. Xia, P. Luo, and Y. Shan. Motionctrl: A unified and flexible motion controller for video generation. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 3
- [43] Y. Wei, M. R. Min, G. Vosselman, L. E. Li, and M. Y. Yang. Planner3d: Llm-enhanced graph prior meets 3d indoor scene explicit regularization. *arXiv preprint arXiv:2403.12848*, 2024. 3
- [44] G. Wu, T. Yi, J. Fang, L. Xie, X. Zhang, W. Wei, W. Liu, Q. Tian, and X. Wang. 4d gaussian splatting for real-time dynamic scene rendering. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20310–20320, 2023. 8
- [45] H. Wu, Z. Zhang, W. Zhang, C. Chen, C. Li, L. Liao, A. Wang, E. Zhang, W. Sun, Q. Yan, X. Min, G. Zhai, and W. Lin. 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. 6
- [46] R. Wu, B. Mildenhall, P. Henzler, K. Park, R. Gao, D. Watson, P. P. Srinivasan, D. Verbin, J. T. Barron, B. Poole, et al. Reconfusion: 3d reconstruction with diffusion priors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 21551–21561, 2024. 3
- [47] Z. Wu, Y. Li, H. Yan, T. Shang, W. Sun, S. Wang, R. Cui, W. Liu, H. Sato, H. Li, et al. Blockfusion: Expandable 3d scene generation using latent tri-plane extrapolation. *ACM Transactions on Graphics (TOG)*, 43(4):1–17, 2024. 3
- [48] H. Xie, Z. Chen, F. Hong, and Z. Liu. Citydreamer: Compositional generative model of unbounded 3d cities. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9666–9675, 2024. 3
- [49] H. Xie, Z. Chen, F. Hong, and Z. Liu. Gaussiancity: Generative gaussian splatting for unbounded 3d city generation. *arXiv preprint arXiv:2406.06526*, 2024. 3
- [50] Y. Xie, C.-H. Yao, V. Voleti, H. Jiang, and V. Jampani. Sv4d: Dynamic 3d content generation with multi-frame and multi-view consistency. *arXiv preprint arXiv:2407.17470*, 2024. 3, 6
- [51] L. Yang, B. Kang, Z. Huang, X. Xu, J. Feng, and H. Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10371–10381, 2024. 3
- [52] L. Yang, B. Kang, Z. Huang, Z. Zhao, X. Xu, J. Feng, and H. Zhao. Depth anything v2. *Advances in Neural Information Processing Systems*, 37:21875–21911, 2024. 3, 5
- [53] Z. Yang, J. Teng, W. Zheng, M. Ding, S. Huang, J. Xu, Y. Yang, W. Hong, X. Zhang, G. Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 1
- [54] Z. Yang, H. Yang, Z. Pan, and L. Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d

- gaussian splatting. In *International Conference on Learning Representations (ICLR)*, 2024. 8
- [55] C. Ye, L. Qiu, X. Gu, Q. Zuo, Y. Wu, Z. Dong, L. Bo, Y. Xiu, and X. Han. Stablenormal: Reducing diffusion variance for stable and sharp normal. *ACM Transactions on Graphics (TOG)*, 43(6):1–18, 2024. 3
- [56] C. Ye, Y. Wu, Z. Lu, J. Chang, X. Guo, J. Zhou, H. Zhao, and X. Han. Hi3dgen: High-fidelity 3d geometry generation from images via normal bridging. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3
- [57] W. Yin, C. Zhang, H. Chen, Z. Cai, G. Yu, K. Wang, X. Chen, and C. Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9043–9053, 2023. 3
- [58] H.-X. Yu, H. Duan, C. Herrmann, W. T. Freeman, and J. Wu. Wonderworld: Interactive 3d scene generation from a single image. *arXiv preprint arXiv:2406.09394*, 2024. 1, 3, 6, 7
- [59] H.-X. Yu, H. Duan, J. Hur, K. Sargent, M. Rubinstein, W. T. Freeman, F. Cole, D. Sun, N. Snavely, J. Wu, et al. Wonderjourney: Going from anywhere to everywhere. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6658–6667, 2024. 1, 3, 6
- [60] Y. Zeng, G. Wei, J. Zheng, J. Zou, Y. Wei, Y. Zhang, and H. Li. Make pixels dance: High-dynamic video generation, 2023. 1
- [61] S. Zhang, M. Zhou, Y. Wang, C. Luo, R. Wang, Y. Li, Z. Zhang, and J. Peng. Cityx: Controllable procedural content generation for unbounded 3d cities. *arXiv preprint arXiv:2407.17572*, 2024. 3
- [62] Z. Zheng, X. Peng, T. Yang, C. Shen, S. Li, H. Liu, Y. Zhou, T. Li, and Y. You. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*, 2024. 1