

# WonderWorld: *Interactive* 3D Scene Generation from a Single Image



Figure 1. Starting with a single image, we interactively generate an immersive, virtual world by extrapolating 3D scenes based on user input. We recommend seeing the interactive generation process at <https://WonderWorld-2024.github.io/>.

## Abstract

We present WonderWorld, a novel framework for interactive 3D scene extrapolation that enables users to explore and shape virtual environments based on a single input image and user-specified text. While significant improvements have been made to the visual quality of scene generation, existing methods are run offline, taking tens of minutes to hours to generate a scene. By leveraging Fast Gaussian Splatting and a guided diffusion-based depth estimation method, WonderWorld generates geometrically consistent extrapolation while significantly reducing computational time. Our framework generates connected and diverse 3D scenes in less than 10 seconds on a single A6000 GPU, enabling real-

time user interaction and exploration. We demonstrate the potential of WonderWorld for applications in virtual reality, gaming, and creative design, where users can quickly generate and navigate immersive, potentially infinite virtual worlds from a single image. Our approach represents a significant advancement in interactive 3D scene generation, opening up new possibilities for user-driven content creation and exploration in virtual environments. We will release full code and software for reproducibility. Project website: <https://WonderWorld-2024.github.io/>.

## 1. Introduction

Within the last year, 3D scene generation has surged in popularity, with many works successfully exploring strong gen-

\*Equal contribution.

erative image priors and improvements in monocular depth estimation. These works have substantially improved the visual quality, possible viewpoints, and diversity of generated scenes. However, all of this work has been done offline, where the user provides a single starting image or text prompt, and then the system, after tens of minutes to hours, returns a fixed 3D scene or a video of a specific camera path through the scene. While offline generation may work for small, discrete scenes or videos, this setup is problematic for many scene generation use cases. For example, in game development, world designers want to iteratively build 3D worlds step-by-step, having control over the generation process and being able to see intermediate steps with low latency. In VR and video games, users expect scalable, varied content that is larger and more diverse than the scenes currently generated. In the future, users may desire even more: a system that allows them to freely explore and shape a dynamically evolving, infinite virtual world. All of these motivate the problem of *interactive* 3D scene generation, where the user, with low latency, can control what the scene extrapolation should feature, e.g., through text prompts, and where that scene extrapolation should take place, e.g., through camera control.

To understand the technical problems that prevent interactivity, we examine several state-of-the-art 3D scene generation methods [7, 38, 39] and identify two major limitations. First, the scene generation is far too slow to be interactive. Each generated scene requires tens of minutes for the multiple passes of generative image inpainting and depth estimation. Second, the generated scenes have strong geometric distortion along the scene boundary, preventing extrapolation from the generated scene.

In this paper, we propose a framework named WonderWorld for interactive scene generation. Our input is a single image, and our output is a set of connected yet diverse 3D scenes. To address the speed issue, our core technique involves Fast Gaussian Surfels, whose optimization requires less than 1 second due to a principled, geometry-based initialization, and a layer-wise scene generation, where each scene is parsed for regions where disocclusion could happen and then pre-generating content to fill these special regions. To address the geometry distortion problem, we introduce a guided diffusion-based depth estimation method that ensures the alignment between the geometry of the extrapolated and existing scenes.

Using our framework, extrapolating or generating a scene takes less than 10 seconds on a single A6000 GPU. This breakthrough unlocks the potential for interactive scene generation, allowing users to extrapolate a single image into a vast and immersive virtual world. Our approach enables new possibilities for applications in virtual reality, gaming, and creative design, where users can quickly generate and explore diverse 3D scenes.

## 2. Related Work

**Perpetual view generation.** Early examples of scene generation focused on extending a single image into a perpetual video with a given camera trajectory: Infinite Images [19] used image stitching, while Infinite Nature [27] and its follow-up works [5, 26] used GANs specialized to nature images. Since the advent of generative diffusion models, subsequent work has expanded the scope and domain of this work. SceneScape [10] generates a domain-free perpetual scene using a single prompt to create a long mesh. WonderJourney [38] instead uses an LLM to generate diverse content and a point cloud representation for the world. While these later works use explicit 3D representations, their techniques are designed with their specific camera trajectory (forward / backward) in mind and can cause issues when facing generic camera trajectories, especially ones that pan sideways or have small rotations. These works also run offline and require tens of minutes to render a single path.

**Single 3D scene generation.** Recently, scene generation methods have focused primarily on a single, dense scene viewable from a local region of viewpoints, with many explicitly focusing on indoor scenes [2, 9, 16]. Text2NeRF [41] generates domain-free scenes but uses a NeRF representation and focuses on local camera trajectories. Recent methods such as LucidDreamer [7] and CAT3D [12] generate multi-view images of a scene, and RealmDreamer [34] distill multi-view image and depth priors to generate a 3D scene [22]. Another line of relevant work focuses on single-image 3D scene reconstruction by explicit pose-conditioning or training on scenes [6, 33, 36, 39]. While these approaches demonstrate significant improvements in the quality of 3D scene generation, they are offline processes generating a fixed scene that is then provided to the user. Since the scene is fixed, their methods do not allow user interaction, e.g., not enabling the user to choose what and where they want to see. We instead address the problem of *interactive* 3D scene generation, which requires significant improvements with respect to the runtime and the tackling of geometric distortion in extrapolation.

**Video generation.** Recent improvements in video generation [1, 3, 4, 24] have led to interest in whether these models can also be used as scene generators. Several works have attempted to add camera control to these models, allowing a user to “move” through the scene; e.g., both MotionCtrl [37] and CameraCtrl [13] train specialized modules to enable camera pose control using datasets with camera pose. While these techniques are promising, generated video has no guarantee regarding 3D consistency or adherence to the user-specified camera poses. In addition, current video generation models remain slow and thus not capable of enabling

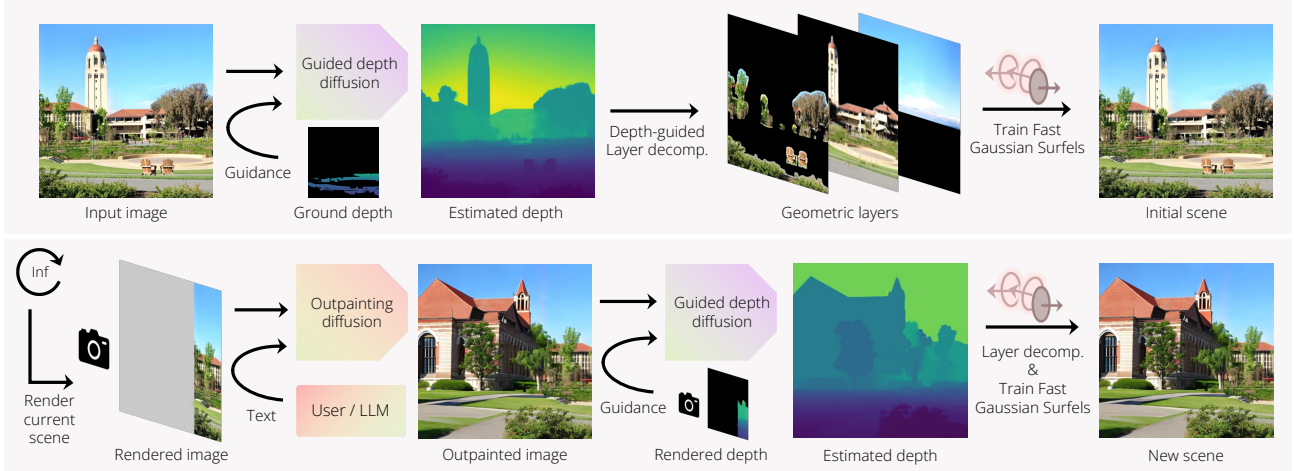


Figure 2. The proposed WonderWorld framework: Our system takes a single image as input and generates connected diverse 3D scenes to form a virtual world. Users can specify new scene contents and styles via text. Our system generates a single 3D scene in less than 10 seconds.

interactivity.

**Fast 3D scene representations** Substantial progress has been made in the last several years regarding the quality and speed of 3D representations; the seminal NeRF [29] paper was followed by Plenoxels [11], InstantNGP [30], and finally 3D Gaussian Splatting [22]. In the context of 3DGS, researchers also revisited the traditional idea of surfels [31, 35]: High-quality Surface Reconstruction using Gaussian Surfels [8] and 2D Gaussian splatting for geometrically accurate radiance fields [17]. While the main focus of these methods is improving reconstruction quality, our goal is to use surfels to speed up the optimization process since they can be initialized in a principled geometry-based way using depth and normals.

### 3. Approach

Our goal is to generate a set of diverse yet coherently connected 3D scenes, forming a potentially infinite virtual world. To this end, we propose WonderWorld, a framework that allows fast scene extrapolation and real-time rendering for an interactive visual experience.

**Overview** We show an illustration of our WonderWorld framework in Figure 2. The main idea is to start by generating a 3D scene from an input image and iteratively extending it by extrapolating existing scenes. Users can provide text to specify the scene contents to generate, or that can be left to a Large Language Model (LLM).

The major technical challenges include scene generation speed and geometric distortion in extrapolated scenes. To speed up scene generation, we adopt the traditional idea of surfels [31, 35], extend them to 3DGS, and show that

this extension allows a principled geometry-based initialization that significantly reduces the optimization time to  $< 1$  second. To deal with the disocclusion holes in generated scenes, we introduce a layer-wise scene generation strategy that is free from multi-view image generation. Thus, WonderWorld allows for fast scene generation within 10 seconds and real-time rendering, simultaneously on a single GPU. To address geometric distortion, we propose to utilize guided depth diffusion to generate geometry for the new scenes. Guided depth diffusion is robust and flexible, allowing the specification of various geometric constraints.

#### 3.1. Fast Gaussian Surfels

We introduce Fast Gaussian Surfels (FGS) to represent our generated 3D scenes. FGS can be seen as a lightweight version of 3DGS with every Gaussian kernel’s  $z$ -axis shrunk to zero. In particular, FGS consists of a set of Gaussian surfels, where each Gaussian surfel is represented by a set of parameters  $\{\mathbf{p}, \mathbf{q}, \mathbf{s}, o, \mathbf{c}\}$ , where  $\mathbf{p}$  denotes the 3D spatial position of the Gaussian kernel,  $\mathbf{q}$  denotes the orientation quaternion,  $\mathbf{s} = [s_x, s_y]$  denotes the scales for the  $x$ -axis and  $y$ -axis,  $o$  denotes the opacity, and  $\mathbf{c}$  denotes the RGB color. We assume Lambertian surfaces in generated scenes, and thus the 3-dimensional color  $\mathbf{c}$  is view-independent. The kernel of a Gaussian surfel is

$$G(\mathbf{x}) = e^{-\frac{1}{2}(\mathbf{x}-\mathbf{p})^T \Sigma^{-1}(\mathbf{x}-\mathbf{p})}, \quad (1)$$

where the covariance matrix  $\Sigma$  is constructed from the scales and the rotation matrix  $\mathbf{Q}$  that can be obtained from the quaternions  $\mathbf{q}$ . The covariance matrix is

$$\Sigma = \mathbf{Q} \text{diag}(s_x^2, s_y^2, 0) \mathbf{Q}^T. \quad (2)$$

The rasterization and alpha-blending rendering process remains the same as the 3D Gaussian Splatting (3DGS) [22].

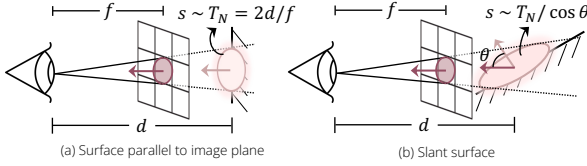


Figure 3. Scale initialization of FGS: We initialize the scales such that a surfel is proportional to the Nyquist interval to alleviate alias holes.

**Geometry-based initialization** The core idea of our fast optimization is that, since we generate our Fast Gaussian Surfels from a single-view image, we can assume that every single pixel in the image reveals a surfel in the underlying 3D scene. Therefore, a surfel’s parameters can be readily solved or approximated by leveraging the information of the corresponding pixel, rather than being randomly initialized and optimized. The optimization is thus simplified, accelerated, and properly regularized.

Specifically, given an input image  $\mathbf{I}$  of  $H \times W$  pixels, we aim at generating  $H \times W$  surfels to represent the underlying 3D scene. The color  $\mathbf{c}$  of a surfel is initialized as the RGB values of the pixel. A surfel’s position  $\mathbf{p}$  can be estimated by back-projection:

$$\mathbf{p} = \mathbf{R}^{-1}(d \cdot \mathbf{K}^{-1}[u, v, 1]^T - \mathbf{T}), \quad (3)$$

where  $u, v$  denote the pixel coordinates, and  $\mathbf{K}, \mathbf{R}, \mathbf{T}$  denote the intrinsic matrix, rotation matrix, and translation vector of the current camera, respectively.  $d$  denotes the estimated depth of the pixel. We discuss depth estimation in more detail in Section 3.3.

To initialize the orientation of the surfel, note that the third column  $\mathbf{Q}_z$  of the rotation matrix  $\mathbf{Q} = [\mathbf{Q}_x, \mathbf{Q}_y, \mathbf{Q}_z]$  is the normal direction of the surfel. Thus, we can construct the rotation matrix  $\mathbf{Q}$ :

$$\mathbf{Q}_z = \mathbf{n}, \quad \mathbf{Q}_x = \frac{\mathbf{u} \times \mathbf{n}}{\|\mathbf{u} \times \mathbf{n}\|}, \quad \mathbf{Q}_y = \frac{\mathbf{n} \times \mathbf{Q}_x}{\|\mathbf{n} \times \mathbf{Q}_x\|}, \quad (4)$$

where  $\mathbf{u} = [0, 1, 0]^T$  denotes a unit up-vector,  $\mathbf{n} = \mathbf{R}^{-1}\mathbf{n}_c$  denotes an estimated normal of the pixel in world frame and  $\mathbf{n}_c$  denotes the camera-frame normal estimated from the image  $I$ .

When it comes to the scale  $s$ , we need to find an appropriate initialization such that it prevents aliasing, e.g., it should not lead to holes when moving closer to a scene. To achieve this, we consider the Nyquist interval for a surfel. Let the sampling interval of our image (i.e., pixel size) be 1. The Nyquist interval  $T_N$  for a surfel at distance  $d$  is then  $T_N = 2d/f$ , where  $f$  denotes the focal length. We want to set the scales of the surfel to be proportional to  $T_N$ , such that it approximately covers the interval  $T_N$  to minimize aliasing. Intuitively, this means that the surfels provide seamless coverage of the visible surface without significant overlap. We




Figure 4. Our FGS can alleviate alias holes when moving into a scene.




Figure 5. We use significant depth edges, i.e., edges in the depth map with large spatial gradient, to identify the foreground layer.

show an illustration in Figure 3 (a) for the case where the surface is parallel to the image plane. As shown in Figure 3 (b), if the surface is not parallel to the image plane, we need to add a cosine term to the scale. Thus, we initialize the scales as

$$s_x = kd/f_x \cos \theta_x, \quad s_y = kd/f_y \cos \theta_y, \quad (5)$$

where  $k = 1/\sqrt{2}$  denotes a hyperparameter,  $\cos \theta_x$  denotes the cosine between the image plane normal and the surfel normal projected to the  $XoZ$  plane. We show a comparison using our FGS scale initialization and Mip-Splatting [40] anti-aliasing 3D Gaussian initialization in Figure 4 after optimization. We observe that our scale initialization alleviates the alias holes when moving into the generated scene.

**Optimization** We use the same photometric loss function as 3DGS:  $\mathcal{L} = 0.8\mathcal{L}_1 + 0.2\mathcal{L}_{D-SSIM}$ . We optimize for the opacity, orientation, and scales, but not for colors and spatial positions. Our optimization includes 100 iterations and no densification process. In practice, we add a small number to the  $z$  axis instead of zero to allow higher representation capacity while leveraging our principled initialization.

### 3.2. Layer-wise Scene Generation

To fill the disocclusion holes in generated scenes, we introduce a layer-wise scene generation strategy. The main idea is to parse the geometric layer structure of the scene to discover the regions where significant disocclusion could occur, reveal these regions by removing the occluding contents, and generate contents to fill these regions. We refer to this process as depth-guided layer decomposition and show an example in the top row of Figure 2.




Figure 6. A major challenge in generating large scenes is the geometric distortion. Our guided depth diffusion significantly alleviates this problem.

In particular, we decompose an image into three layers from front to back: the foreground layer  $F$ , the background layer  $B$ , and the sky layer  $Y$ . Since disocclusion occurs at the depth edges, we separate these layers by finding depth edges. We compute a significant depth edge map by thresholding the spatial gradient magnitude of the estimated depth map. The foreground layer  $F$  is formed by finding semantic segments that contain significant depth edges. We slightly dilate the segments to ensure that they intersect depth edges when they indeed do. We show an example in Figure 5. As for the sky layer, we find it more robust to directly use semantic segmentation, as sky depth estimation is very challenging for depth estimators, as also observed in prior work [26, 38].

Given the layer segmentation, we first inpaint the sky layer by a diffusion model and use the inpainted sky image to train a FGS for it. Then, we inpaint the background layer and train a FGS on top of the frozen sky FGS. Finally, we train the foreground FGS on top of the frozen background and the sky FGS.

### 3.3. Guided Depth Diffusion

To generate an infinite world, we need to extrapolate existing scenes to unexplored space. One fundamental challenge is the geometric distortion during extrapolation, that is, the newly generated scene contents may have a significant geometric gap to the existing scene contents, so they appear to be broken when seen from a different viewpoint than the outpainting viewpoint. This is due to the inconsistency between the estimated depth and the existing geometry.

In particular, let  $\mathbf{D}_v$  of size  $H \times W$  be the depth map rendered from visible existing contents at an outpainting camera viewpoint with a binary mask  $\mathbf{M}_v \in \{0, 1\}^{H \times W}$  to indicate visible regions, and let  $\mathbf{D}_e$  be the estimated depth for an outpainted new image  $\mathbf{I}_e$ . Then we observe a strong discrepancy between  $\mathbf{D}_v \odot \mathbf{M}_v$  and  $\mathbf{D}_e \odot \mathbf{M}_v$  where  $\odot$  denotes the element-wise product. We show an example in Figure 6 to demonstrate this issue.

Simple post-processing heuristics, such as alignment by computing a global shift and scale [7] or fine-tuning the depth estimator to match the estimated depth with the existing geometry [38], do not suffice, as they do not reduce the

inherent ambiguity in the estimation of the new scene depth.

To address this challenge, we adopt guidance for depth diffusion networks. The main idea is to formulate the depth estimation of an extrapolated scene as a conditional depth generation problem, i.e., sampling from a depth distribution  $p(\mathbf{D}_e \mid \mathbf{I}_e, \mathbf{D}_v, \mathbf{M}_v)$ , which explicitly takes the observed depth  $\mathbf{D}_v \odot \mathbf{M}_v$  as a conditioning guidance signal. We use a diffusion model because, unlike feedforward networks, they provide a natural way to sample from the depth posterior [14]. Our guided depth diffusion is based on a latent depth diffusion model [21]. In short, a latent depth diffusion model learns to generate a depth map by sampling from  $p(\mathbf{D} \mid \mathbf{I})$ , achieved via gradual denoising of a randomly initialized latent depth map  $\mathbf{d}_T$  with a learned denoiser U-Net  $\epsilon_\theta(\mathbf{d}_t, \mathbf{I}, t)$ , where  $t$  denotes a time step. The generated depth is given by a decoder  $\mathbf{D} = \text{Decoder}(\mathbf{d})$ . We show an illustration in Figure 7 (a).

From a score-based perspective [15], the denoiser  $\epsilon_\theta(\mathbf{d}_t, \mathbf{I}, t)$  predicts an update direction, and the latent depth generation process is done by recursively applying the updates:

$$\epsilon_t = \epsilon_\theta(\mathbf{d}_t, \mathbf{I}, t), \quad \mathbf{d}_{t-1} = \text{update}(\mathbf{d}_t, t, \epsilon_t). \quad (6)$$

We inject the visible depth as guidance by modifying the denoiser as

$$\hat{\epsilon}_t = \epsilon_\theta(\mathbf{d}_t, \mathbf{I}, t) - s_t \nabla_{\mathbf{d}_t} \|\mathbf{D}_{t-1} \odot \mathbf{M}_v - \mathbf{D}_v \odot \mathbf{M}_v\|^2, \quad (7)$$

where  $\hat{\epsilon}_t$  denotes the guided denoiser,  $\mathbf{D}_{t-1} = \text{Decoder}(\mathbf{d}_{t-1})$  denotes a pre-decoded depth map, and  $s_t$  denotes the guidance weight.

Our modification can be seen as composing two score functions to sample from the conditional distribution  $p(\mathbf{D}_e \mid \mathbf{I}_e, \mathbf{D}_v, \mathbf{M}_v)$ . This conditional distribution considers both the visible existing depth  $\mathbf{D}_v$  and the new scene geometry in  $\mathbf{I}_e$  simultaneously, leading to much smoother geometry extrapolation.

**Tackling ground plane distortion** We note that our guided depth diffusion formulation is highly flexible and allows us to specify different depth constraints. For example, another significant geometric distortion is that the ground plane is often curved. Thus, for all generated scenes, we add depth guidance for the ground plane by replacing the mask  $\mathbf{M}_v$  in Eq. 3.3 with a ground mask  $\mathbf{M}_g$  obtained from semantic segmentation, and replacing the depth of visible content  $\mathbf{D}_v$  with an analytically calculated flat ground depth  $\mathbf{D}_g$ .

## 4. Results

In this section, we present the results of our WonderWorld. As we are not aware of any baseline method that allows interactive scene generation, we focus on showcasing the

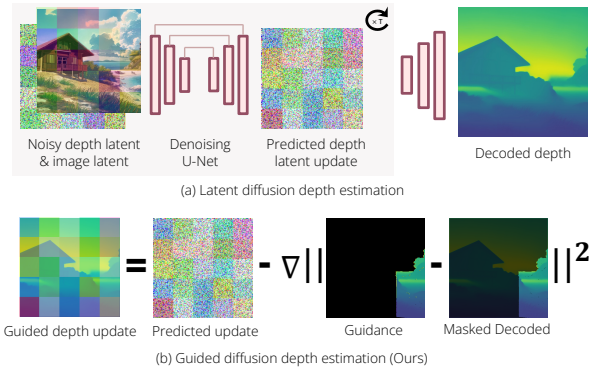


Figure 7. Illustration of guided diffusion depth estimation. Note, the colored patches indicate that depth is computed in latent space.

quality of generated large-scale 3D scenes. For this purpose, we consider open-source baselines and use their official code. We demonstrate examples of interactive scene generation in our video and strongly encourage the reader to view it first.

Our baseline methods include WonderJourney [38], a state-of-the-art perpetual view generation method, and LucidDreamer [7], a recent 3D scene generation method. WonderJourney takes a single image as input and generates a sequence of point clouds by outpainting images and unprojecting pixels. LucidDreamer takes a single image as input and synthesizes multi-view images from it to train a 3DGS. We use publicly available real and synthetic images in our examples.

#### 4.1. Implementation details

In our implementation, we use the Stable Diffusion Inpaint model [32] as our outpainting model. We also use it for inpainting the background and sky layers. We use OneFormer [18] to segment the sky, the ground, and foreground objects. In the initial scene, we generate the entire sky using SyncDiffusion [25] offline. We estimate normal using the Marigold normal estimator [21]. We use Marigold as our depth diffusion model. In our guided depth diffusion, we set the guidance weights  $s_t$  such that the norm of the guidance signal is proportional to the norm of the predicted update. We use the Euler scheduler [20] with 30 steps for our depth diffusion, where we apply our guidance in the last 8 steps. We post-process estimated depth using an efficient SAM [23, 28], similar to WonderJourney [38]. We also follow WonderJourney to use GPT4 to generate prompts when users do not provide text, and to enrich the prompt by adding plausible objects and background text according to the scene name. We will release full code and software for reproducibility.

#### 4.2. Qualitative results

We show a qualitative comparison using the same input image for our WonderWorld and the baseline methods in

Figure 10. Note that our WonderWorld results consist of 9 scenes, and the LucidDreamer result consists of one scene. WonderJourney only supports extracting 3D points between two consecutive scenes; we extend the code to support generating points for up to 4 scenes here.

From Figure 10, we observe that single 3D scene generation methods like LucidDreamer [7] do not extrapolate out of predefined scenes and suffer from severe geometric distortion at the boundaries of the generated scene. Although WonderJourney [38] allows the generation of multiple scenes that appear to be coherently connected in specific views, the geometric distortion is significant when rendered from a different camera angle. In contrast to baselines, our WonderWorld significantly alleviates geometric distortion, leading to a coherent large-scale 3D scene. We show more examples in Figure 8, Figure 12, and Figure 13].

Since WonderWorld allows for the choice of different text prompts to change the contents, the generated scenes can be diverse and different in each run. We show an example of diverse generation results from the same input image in Fig. 9. WonderWorld also allows a user to specify different styles in the same generated virtual world, e.g., Minecraft, painting, and Lego styles as shown in Fig. 11.

#### 4.3. Generation speed

Since our focus is on making 3D scene generation interactive, we report the scene generation time cost from starting generation until one can see the results. We show the scene generation time for a single scene in Table 1. From Table 1 we see that even the fastest existing method, WonderJourney, takes more than 700 seconds to generate a single scene, spending most of its time generating multiple views to fill in the holes between the existing scene and the newly generated scene. LucidDreamer generates a slightly extended scene from the input image and spends most of its time generating multiple views, aligning depth for these views, and training a 3DGS to fit these views. In general, prior approaches need to generate or distill multiple views and optimize their 3D scene representations for a significant amount of time. We accelerate the representation optimization by our FGS that benefits from a principled geometry-based initialization, and reduce the number of images needed by our layer-wise scene generation strategy. Together, these contribute to our fast scene generation. We show an analysis of our time cost in Table 2. Since diffusion model inference (outpainting, layer inpainting, depth, and normal estimation) takes the most time, our method will benefit from future advances in accelerating diffusion inference.

### 5. Conclusion

We introduce WonderWorld, a system for interactive 3D scene generation, featuring technical improvements that significantly speed up generation time and improve perfor-

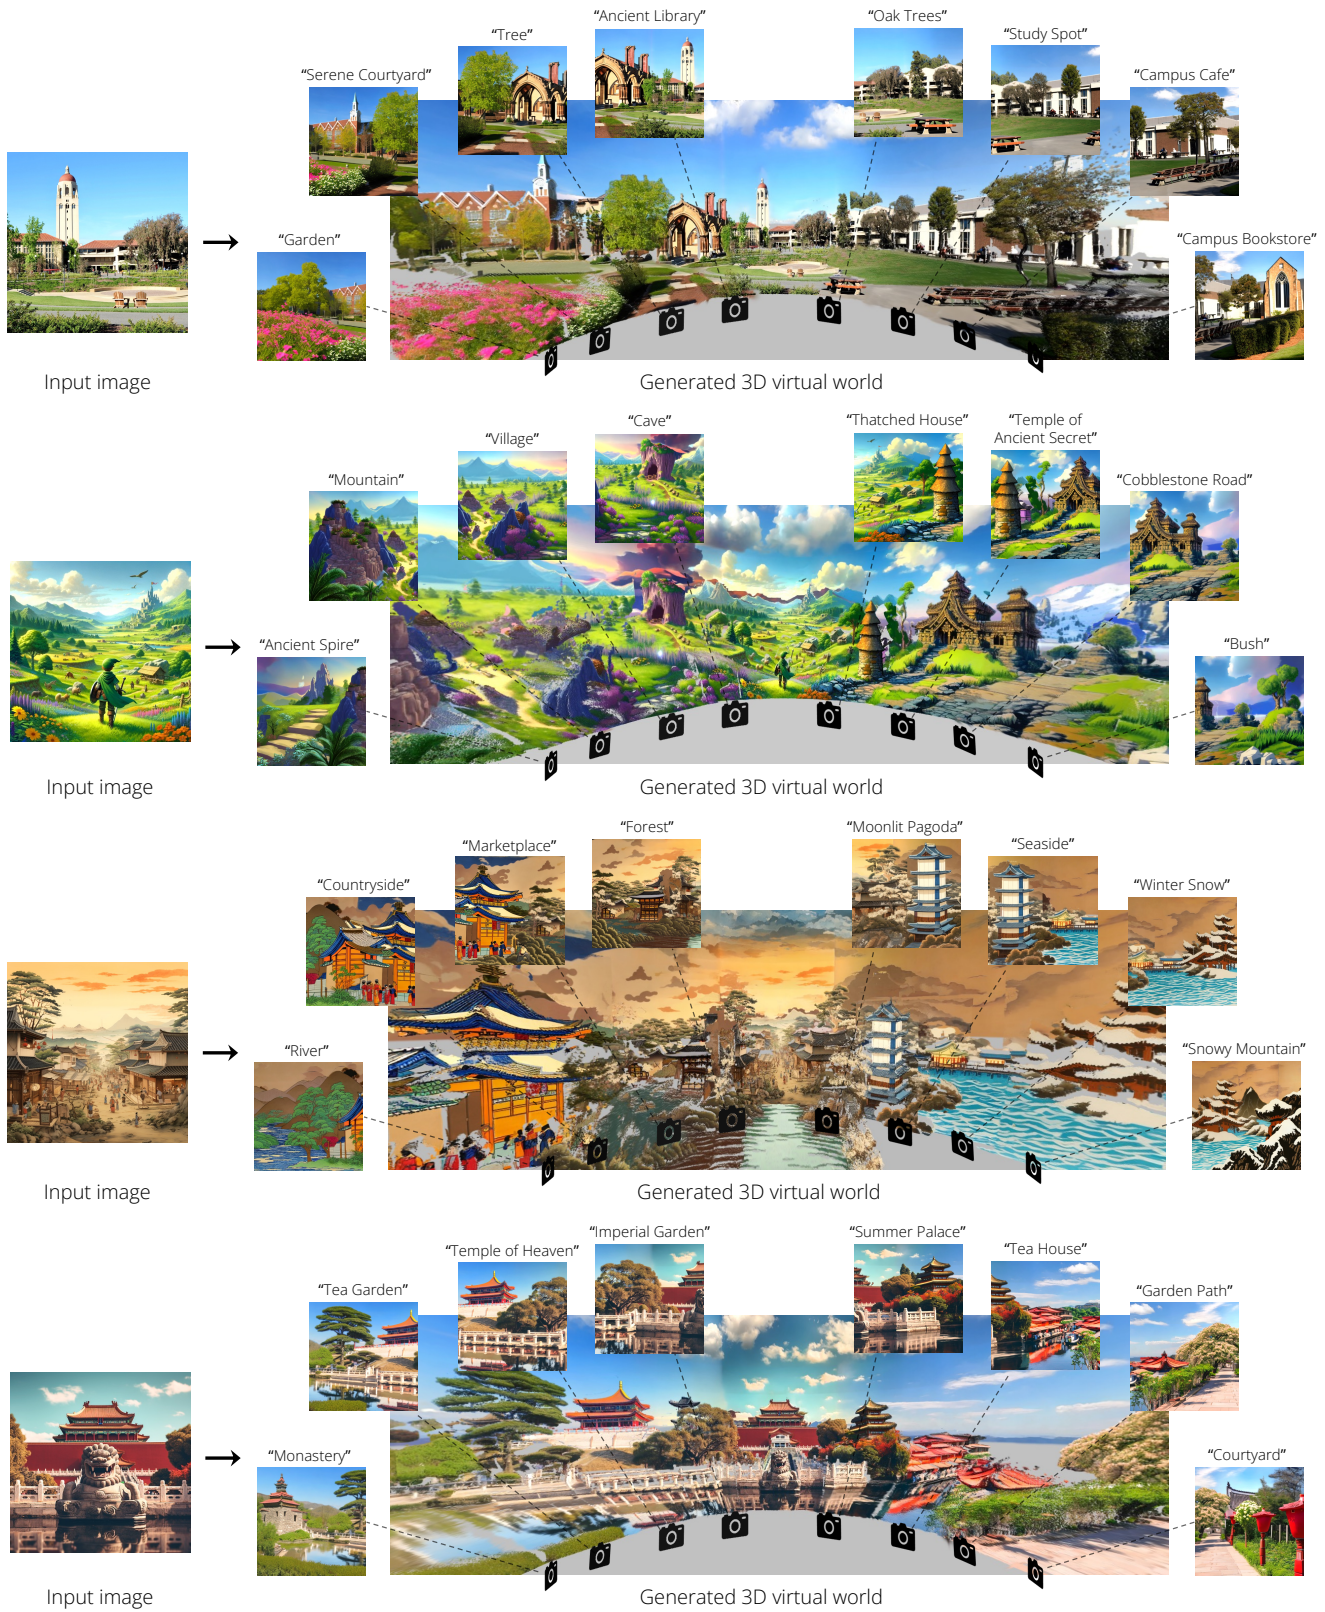


Figure 8. Qualitative examples. These examples are generated with scene contents automatically generated by the LLM.



Figure 9. Diverse generation: Our WonderWorld allows generating different virtual worlds from the same input image.

Table 1. Time costs for generating a single scene, on an A6000 GPU.

WonderJourney [38]	LucidDreamer [7]	WonderWorld (Ours)
749.5 seconds	798.1 seconds	9.5 seconds

Table 2. Time cost analysis for WonderWorld in generating a single scene, tested on an A6000 GPU.

Outpainting	Layer generation	Depth	Normal	FGS
2.1s	2.3s	2.5s	0.8s	1.9s

mance for large, diverse scenes. WonderWorld allows users to interactively generate and explore the parts of the scene they want with the content they request.

**Limitations** A limitation of WonderWorld is the low scene density, as each scene only has up to  $H \times W$  Gaussian surfels. Another limitation is the difficulty in handling detailed objects, such as trees, which can lead to inaccurate depth estimation, leaving “holes” or “floaters” when the viewpoint changes. We demonstrate a failure case in our video. Therefore, an exciting future direction is using WonderWorld to interactively prototype a coarse world structure, and then refining it to boost scene density, fill holes, and remove floaters

with slower single-scene multi-view diffusion models.

## References

- [1] Bar-Tal et al. 2024. Lumiere: A Space-Time Diffusion Model for Video Generation. *arXiv:2401.12945* (2024). 2
- [2] Miguel Angel Bautista, Pengsheng Guo, Samira Abnar, Walter Talbott, Alexander Toshev, Zhuoyuan Chen, Laurent Dinh, Shuangfei Zhai, Hanlin Goh, Daniel Ulbricht, et al. 2022. Gaudi: A neural architect for immersive 3d scene generation. *Advances in Neural Information Processing Systems* 35 (2022), 25102–25116. 2
- [3] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. 2023. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127* (2023). 2
- [4] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufeı 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> 2
- [5] Lucy Chai, Richard Tucker, Zhengqi Li, Phillip Isola, and Noah Snavely. 2023. Persistent Nature: A Generative Model of Unbounded 3D Worlds. In *Proceedings of the IEEE/CVF*





Figure 10. Baseline comparison. The inset with blur dashed bounding box is the input image.

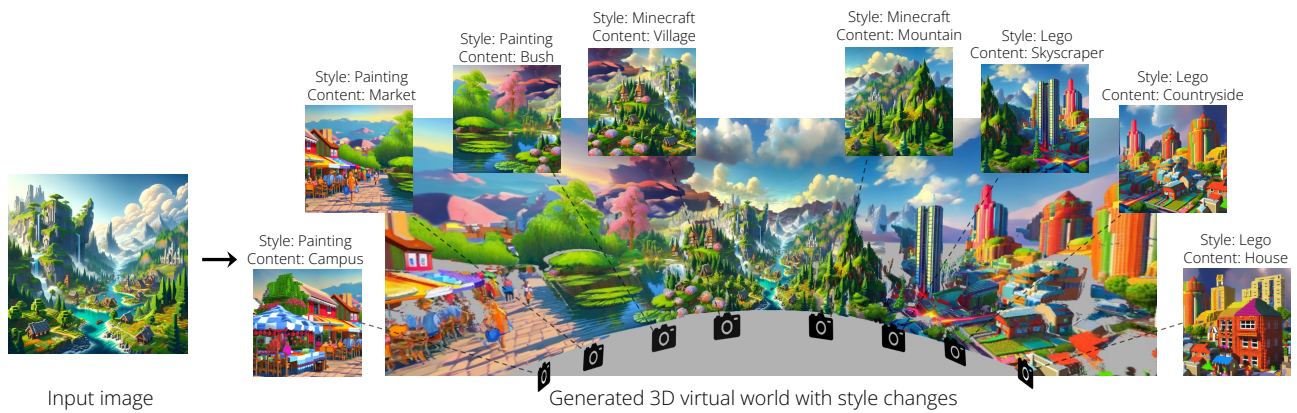


Figure 11. Our WonderWorld allows a user to specify different styles in the same generated virtual world, e.g., Minecraft, painting, and Lego styles.



Figure 12. Qualitative examples. These examples are generated with scene contents automatically generated by the LLM.

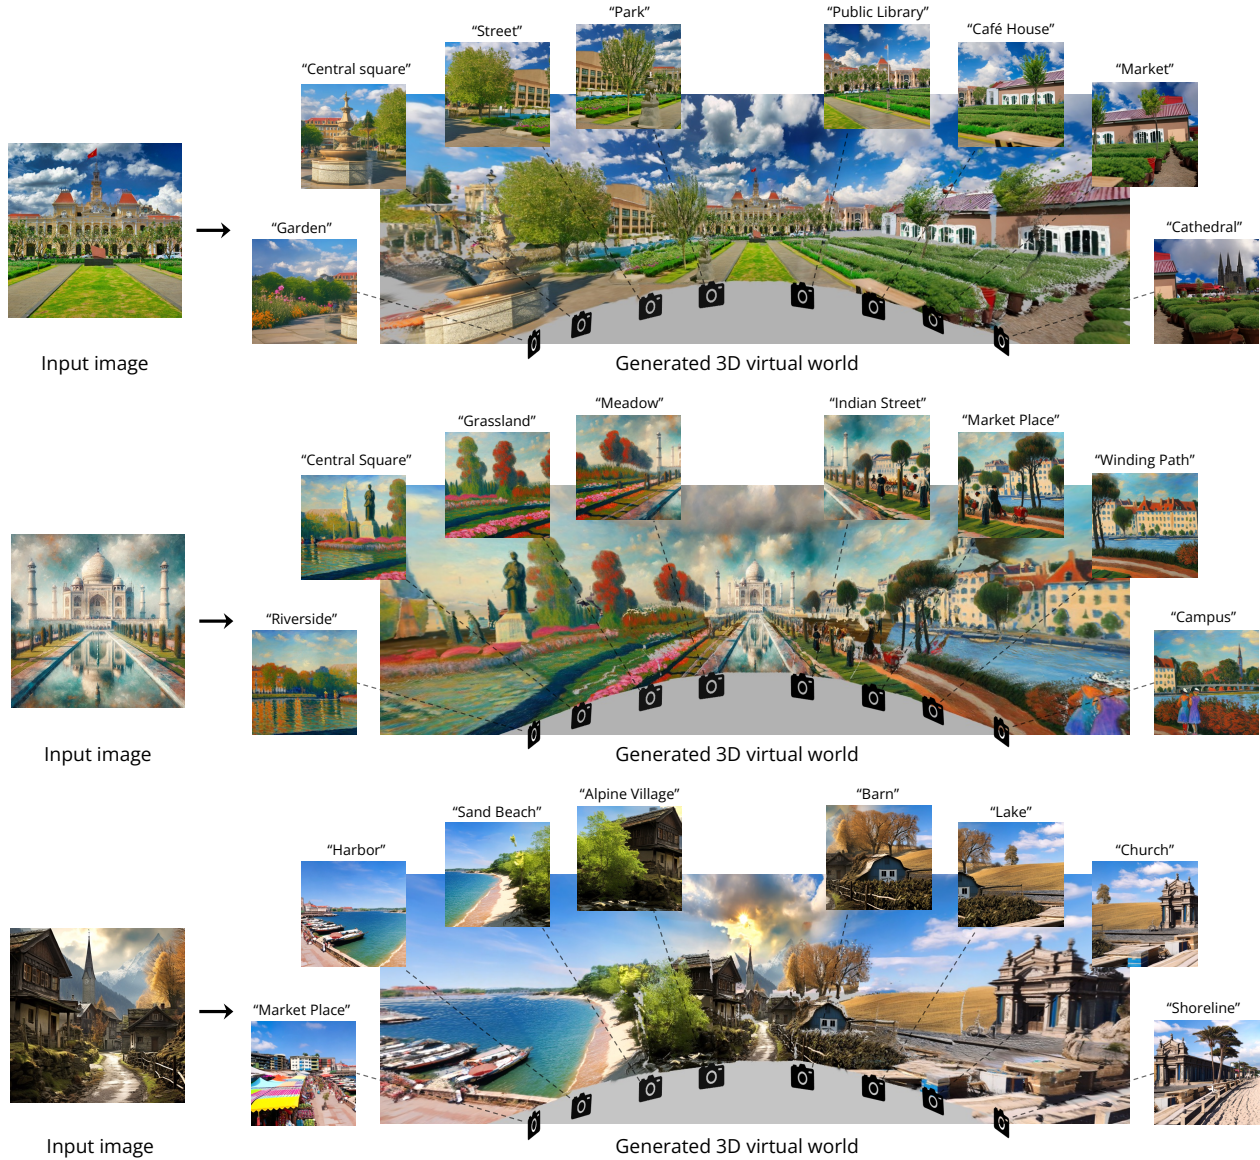


Figure 13. Qualitative examples. These examples are generated with scene contents automatically generated by the LLM.

- Conference on Computer Vision and Pattern Recognition.* 20863–20874. [2](#)
- [6] Eric R Chan, Koki Nagano, Matthew A Chan, Alexander W Bergman, Jeong Joon Park, Axel Levy, Miika Aittala, Shalini De Mello, Tero Karras, and Gordon Wetzstein. 2023. GeNVs: Generative novel view synthesis with 3D-aware diffusion models. [2](#)
- [7] Jaeyoung Chung, Suyoung Lee, Hyeongjin Nam, Jaerin Lee, and Kyoung Mu Lee. 2023. Lucidreamer: Domain-free generation of 3d gaussian splatting scenes. *arXiv preprint arXiv:2311.13384* (2023). [2](#), [5](#), [6](#), [8](#)
- [8] Pinxuan Dai, Jiamin Xu, Wenxiang Xie, Xinguo Liu, Huamin Wang, and Weiwei Xu. 2024. High-quality Surface Reconstruction using Gaussian Surfels. *arXiv preprint arXiv:2404.17774* (2024). [3](#)
- [9] Terrance DeVries, Miguel Angel Bautista, Nitish Srivastava, Graham W. Taylor, and Joshua M. Susskind. 2021. Unconstrained Scene Generation with Locally Conditioned Radiance Fields. In *ICCV*. [2](#)
- [10] Rafail Fridman, Amit Abecasis, Yoni Kasten, and Tali Dekel. 2023. Scenescape: Text-driven consistent scene generation. *arXiv preprint arXiv:2302.01133* (2023). [2](#)
- [11] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. 2022. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5501–5510. [3](#)
- [12] Ruiqi Gao, Aleksander Holynski, Philipp Henzler, Arthur

- Brussee, Ricardo Martin-Brualla, Pratul Srinivasan, Jonathan T Barron, and Ben Poole. 2024. CAT3D: Create Anything in 3D with Multi-View Diffusion Models. *arXiv preprint arXiv:2405.10314* (2024). 2
- [13] Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, and Ceyuan Yang. 2024. CameraCtrl: Enabling Camera Control for Text-to-Video Generation. *arXiv preprint arXiv:2404.02101* (2024). 2
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems* 33 (2020), 6840–6851. 5
- [15] Jonathan Ho and Tim Salimans. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598* (2022). 5
- [16] Lukas Höllein, Ang Cao, Andrew Owens, Justin Johnson, and Matthias Nießner. 2023. Text2room: Extracting textured 3d meshes from 2d text-to-image models. *arXiv preprint arXiv:2303.11989* (2023). 2
- [17] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2024. 2d gaussian splatting for geometrically accurate radiance fields. *arXiv preprint arXiv:2403.17888* (2024). 3
- [18] Jitesh Jain, Jiachen Li, Mang Tik Chiu, Ali Hassani, Nikita Orlov, and Humphrey Shi. 2023. Oneformer: One transformer to rule universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2989–2998. 6
- [19] Biliana Kaneva, Josef Sivic, Antonio Torralba, Shai Avidan, and William T Freeman. 2010. Infinite images: Creating and exploring a large photorealistic virtual space. *Proc. IEEE* 98, 8 (2010), 1391–1407. 2
- [20] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. 2022. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems* 35 (2022), 26565–26577. 6
- [21] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. 2023. Repurposing diffusion-based image generators for monocular depth estimation. *arXiv preprint arXiv:2312.02145* (2023). 5, 6
- [22] 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 (2023), 1–14. 2, 3
- [23] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollar, and Ross Girshick. 2023. Segment Anything. In *ICCV*. 4015–4026. 6
- [24] Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Rachel Hornung, Hartwig Adam, Hassan Akbari, Yair Alon, Vighnesh Birodkar, et al. 2023. Videopoet: A large language model for zero-shot video generation. *arXiv preprint arXiv:2312.14125* (2023). 2
- [25] Yuseung Lee, Kunho Kim, Hyunjin Kim, and Minhyuk Sung. 2023. Syncdiffusion: Coherent montage via synchronized joint diffusions. *Advances in Neural Information Processing Systems* 36 (2023), 50648–50660. 6
- [26] Zhengqi Li, Qianqian Wang, Noah Snavely, and Angjoo Kanazawa. 2022. Infinitenature-zero: Learning perpetual view generation of natural scenes from single images. In *European Conference on Computer Vision*. Springer, 515–534. 2, 5
- [27] Andrew Liu, Richard Tucker, Varun Jampani, Ameesh Makadia, Noah Snavely, and Angjoo Kanazawa. 2021. Infinite nature: Perpetual view generation of natural scenes from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14458–14467. 2
- [28] Xinyu Liu, Houwen Peng, Ningxin Zheng, Yuqing Yang, Han Hu, and Yixuan Yuan. 2023. Efficientvit: Memory efficient vision transformer with cascaded group attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14420–14430. 6
- [29] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (2021), 99–106. 3
- [30] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)* 41, 4 (2022), 1–15. 3
- [31] Hanspeter Pfister, Matthias Zwicker, Jeroen Van Baar, and Markus Gross. 2000. Surfels: Surface elements as rendering primitives. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. 335–342. 3
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10684–10695. 6
- [33] Kyle Sargent, Zizhang Li, Tanmay Shah, Charles Herrmann, Hong-Xing Yu, Yunzhi Zhang, Eric Ryan Chan, Dmitry Lagun, Li Fei-Fei, Deqing Sun, et al. 2023. ZeroNVS: Zero-Shot 360-Degree View Synthesis from a Single Real Image. *arXiv preprint arXiv:2310.17994* (2023). 2
- [34] Jaidev Shriram, Alex Trevithick, Lingjie Liu, and Ravi Ramamoorthi. 2024. RealmDreamer: Text-Driven 3D Scene Generation with Inpainting and Depth Diffusion. *arXiv preprint arXiv:2404.07199* (2024). 2
- [35] Richard Szeliski and David Tonnesen. 1992. Surface modeling with oriented particle systems. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*. 185–194. 3
- [36] Ayush Tewari, Tianwei Yin, George Cazenavette, Semon Rezhikov, Joshua B Tenenbaum, Frédo Durand, William T Freeman, and Vincent Sitzmann. 2023. Diffusion with Forward Models: Solving Stochastic Inverse Problems Without Direct Supervision. *arXiv preprint arXiv:2306.11719* (2023). 2
- [37] Zhouxia Wang, Ziyang Yuan, Xintao Wang, Tianshui Chen, Menghan Xia, Ping Luo, and Ying Shan. 2023. Motionctrl: A unified and flexible motion controller for video generation. *arXiv preprint arXiv:2312.03641* (2023). 2
- [38] Hong-Xing Yu, Haoyi Duan, Junhwa Hur, Kyle Sargent, Michael Rubinstein, William T Freeman, Forrester Cole, De-

- qing Sun, Noah Snaveley, Jiajun Wu, et al. 2024. WonderJourney: Going from Anywhere to Everywhere. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2, 5, 6, 8
- [39] Jason J Yu, Fereshteh Forghani, Konstantinos G Derpanis, and Marcus A Brubaker. 2023. Long-Term Photometric Consistent Novel View Synthesis with Diffusion Models. *arXiv preprint arXiv:2304.10700* (2023). 2
- [40] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. 2023. Mip-splatting: Alias-free 3d gaussian splatting. *arXiv preprint arXiv:2311.16493* (2023). 4
- [41] Jingbo Zhang, Xiaoyu Li, Ziyu Wan, Can Wang, and Jing Liao. 2024. Text2nerf: Text-driven 3d scene generation with neural radiance fields. *IEEE Transactions on Visualization and Computer Graphics* (2024). 2