

DreamAnywhere: Object-Centric Panoramic 3D Scene Generation

Edoardo A. Dominici¹ Jozef Hladky¹ Floor Verhoeven¹ Lukas Radl²
 Thomas Deixelberger^{1,2} Stefan Ainetter¹ Philipp Drescher^{1,2} Stefan Hauswiesner¹
 Arno Coomans¹ Giacomo Nazzaro¹ Konstantinos Vardis¹ Markus Steinberger^{1, 2}

¹Huawei Technologies

²Graz University of Technology

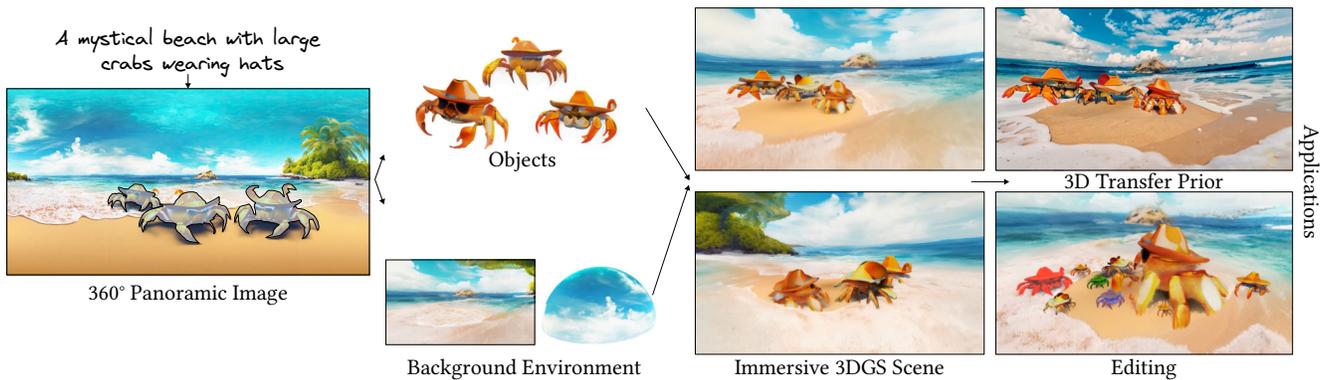


Figure 1. Starting from a text prompt, our method generates 3D scenes leveraging a 360° panoramic image as intermediate representation to extract, reconstruct and compose objects in the environment. Our scenes, represented as 3D Gaussian splats, support long-range exploration and have high structural coherence even under large camera offsets. This improves immersive scene generation for existing applications, enables intuitive object-level editing and makes them strong 3D priors for world-to-world transfer models [1].

Recent advances in text-to-3D scene generation have demonstrated significant potential to transform content creation across multiple industries. Although the research community has made impressive progress in addressing the challenges of this complex task, existing methods often generate environments that are only front-facing, lack visual fidelity, exhibit limited scene understanding, and are typically fine-tuned for either indoor or outdoor settings. In this work, we address these issues and propose *DreamAnywhere*, a modular system for the fast generation and prototyping of 3D scenes. Our system synthesizes a 360° panoramic image from text, decomposes it into background and objects, constructs a complete 3D representation through hybrid inpainting, and lifts object masks to detailed 3D objects that are placed in the virtual environment. *DreamAnywhere* supports immersive navigation and intuitive object-level editing, making it ideal for scene exploration, visual mock-ups, and rapid prototyping—all with minimal manual modeling. These features make our system particularly suitable for low-budget movie production, enabling quick iteration on scene layout and visual tone without the overhead of traditional 3D workflows. Our modular pipeline is highly customizable as

it allows components to be replaced independently. Compared to current state-of-the-art text- and image-based 3D scene generation approaches, *DreamAnywhere* shows significant improvements in coherence in novel view synthesis and achieves competitive image quality, demonstrating its effectiveness across diverse and challenging scenarios. A comprehensive user study demonstrates a clear preference for our method over existing approaches, validating both its technical robustness and practical usefulness.

1. Introduction

Generating large and immersive virtual 3D scenes from text remains a significant challenge with transformative applications in movie production, gaming, virtual reality, and simulation. Recent advances in generative modeling have led to impressive results but the 3D scenes are only coherent within small ranges of camera movement, resulting in incomplete structures and hallucinations when viewers attempt to explore beyond narrow boundaries.

Existing text-to-3D approaches leverage the richness and stylistic diversity of internet-scale diffusion models for pro-

gressive outpainting [18] or distillation sampling [25, 40], but struggle to capture global coherence, leading to increasing semantic drift and structural collapse due to the difficulty of integrating depth estimates into a single 3D scene representation. Video diffusion models have been combined with large reconstruction models [5, 60] but still face limitations in generating complete scenes observable from vastly different viewpoints maintaining coherency. Panoramic 360° images are an inherently good starting point for immersive viewing, in fact [8, 69] combine panoramic images with inpainting to enable parallax effects when moving away from the origin, further extended to multiple layers by [53]. Similarly, approaches designed for 2D fly-throughs [14, 27, 58, 59] can create visually convincing animations, but lack the global 3D structure required for full scene exploration.

Whilst large-scale and artistic 3D scene data remains scarce, there exists a wealth of 3D data for individual objects, offering a valuable resource for compositional approaches. Moreover, decomposing a 3D environment into individual entities offers important advantages for downstream applications [3, 47]. If the layout and description of objects in a scene are known or provided manually, immersive scenes can be generated [25]. However, optimizing scene layouts end-to-end continues to be a major challenge [11, 65].

Building on these insights, we present *DreamAnywhere*, a text-to-3D generation framework that leverages 360° panoramic image generation to drive scene style and geometry, while decomposing scenes into individual objects to preserve structural coherence and enable the generation of navigable 3D environments. We make the following contributions:

- A novel unified and modular system for generating 3D scenes with high 3D coherence and visual fidelity while facilitating interactivity.
- A panoramic image generation method that jointly fine-tunes a 360° image diffusion model with a perspective conditioning mechanism, improving out-of-domain sampling for creative applications.
- A high-quality object reconstruction method that first synthesizes a detailed reference image from multimodal cues—textual, geometric, and stylistic—and then leverages it to guide robust 3D model generation and alignment, even for low-quality and incomplete segmented inputs.
- A hybrid inpainting strategy that combines 2D and 3D techniques: large-scale holes resulting from object removal are inpainted in the 360° image for global coherence, while smaller disocclusions caused by the 3D projection are addressed with 3D inpainting.

To evaluate our approach, we conduct both quantitative comparisons and a controlled user study, demonstrating that our method achieves superior quality, especially under large view offsets. Furthermore, qualitative results and user feed-

back indicate that our approach yields much more coherent and immersive scenes. As shown in Figure 1, this holds even for diverse camera movements.

2. Related Work

Our 3D scene generation work lies at the intersection of image-based generation and compositional scene synthesis, bridging recent advances in both areas. For an in-depth discussion we refer readers to recent surveys on scene generation [24, 46] and object generation [24, 32].

2.1. Image-based Scene Generation

A successful line of work in image-based scene generation focuses on monolithic representations, where the entire scene is captured in a single unified 3D structure. A common technique involves iterative outpainting with image diffusion models to progressively extend the scene along a camera trajectory [14, 18, 26, 58, 59]. Although effective to some extent, these methods often suffer from the accumulation of visual artifacts and distortions across iterations, and offer limited flexibility for isolating and manipulating individual scene components.

An alternative set of techniques reconstructs 3D scenes from 2D images using inpainting-based methods, which lift partial views into 3D and fill occluded regions through inpainting [40, 64]. Although these reconstructions can appear photorealistic from specific angles, their quality tends to degrade when viewed from arbitrary viewpoints.

More recently, panorama-based approaches such as PanoDreamer [31] and PanoDiffusion [50] attempt to synthesize entire 360° panoramic images with depth information through in- or outpainting. Rather than iteratively expanding standard 2D images, other methods focus on creating 360° panorama images in a single step [8, 43, 45, 69]. Another area of research addresses the handling of occlusions via multi-view generation [57], layered approaches [53] or explicitly defining the corresponding entities [71]. Video diffusion models have strong temporal coherency but are still limited in generating 3D content that is consistent along long camera trajectories [5, 19, 36, 60, 62, 66]. While image-based methods offer strong visual quality and global coherence, they generally lack fine-grained control over individual scene elements and are limited in navigational flexibility, especially when compared to compositional scene generation approaches, which we discuss in the following section.

2.2. Compositional Scene Generation

In contrast to monolithic approaches, compositional scene generation focuses on constructing scenes from multiple, independently generated objects which can be edited or animated. Recent advances in this area predominantly leverage pre-trained image diffusion models for scene generation

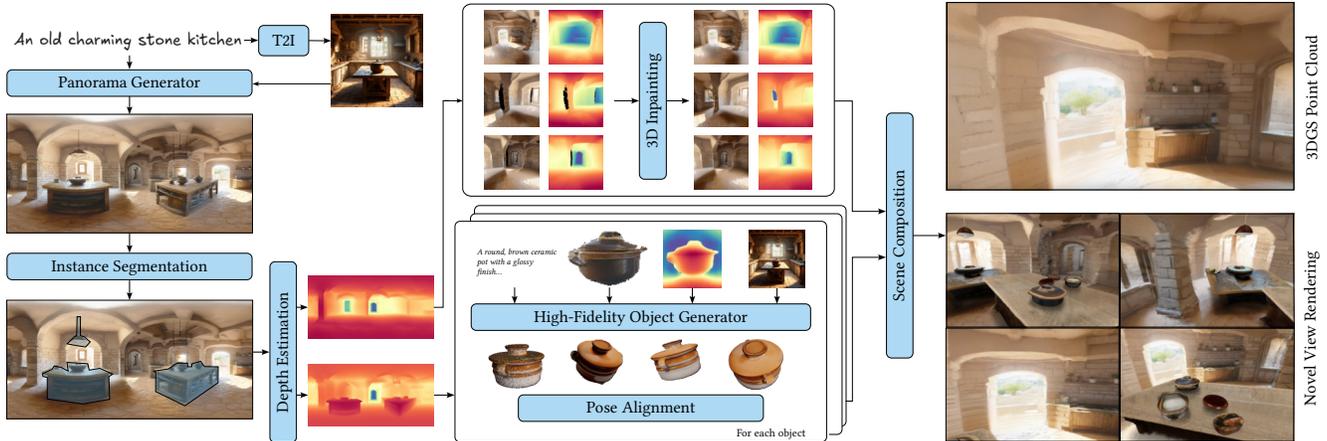


Figure 2. An overview of our architecture. We start by generating a 360° panorama, followed by instance segmentation to separate the foreground objects from the scene background. Segmented objects are lifted to 3D using semantic and geometric priors, while a 3D inpainting process converts the image to a 3DGS representation and fills disoccluded regions. Finally, the scene is composited, enabling navigation of the 3D environment at interactive rates.

[2, 6, 16, 20, 70]. The key differences between these methods lie in the strategies used to layout the scene.

Image-guided methods [2, 16, 20, 65, 70] use generated or real images to infer each object’s location and orientation in the scene. Similarly, distillation-based approaches [6, 11, 25] employ diffusion models to guide object placement. However, since these methods typically rely on single (non-panoramic) images, they are often limited in terms of scale and global coherence of the generated scenes.

To overcome these limitations, proxy-guided techniques introduce explicit control over scene structure, for example by letting users define the scene structure through various semantic annotations or by placing proxy geometries [9, 12, 38, 51, 54]. While these techniques provide great flexibility in terms of scale and coherence, they typically require extensive user input and manual scene design.

Overall, compositional approaches provide a powerful alternative to image-based scene generation, trading off the ease of automation for improved control, reusability, and interpretability of individual scene elements.

3. Method

The DreamAnywhere architecture proceeds through three stages as illustrated in Figure 2. In the **first stage**, a 360° panoramic image I_F is generated from a text prompt and stylized by an input perspective image. From I_F , foreground objects are identified, segmented, and erased, producing a *clean* background image I_B . Then, the depth of I_B is aligned to that of I_F , producing a single, unified depth map and laying the groundwork for subsequent 3D composition. The **second stage** focuses on generating high-fidelity 3D Gaussian Splatting (3DGS) models of the segmented objects and computing their transformations for accurate placement in

the 3D scene. In parallel, the **third stage** converts the background image I_B into a high-resolution 3DGS point cloud. During this process, disoccluded regions are inpainted to ensure visual consistency across viewpoints under novel camera offsets. Finally, the outputs of the second (objects) and third (background) stages are fused into a unified 3DGS scene, enabling immersive and interactive navigation. The remainder of this section presents all the components of our system, emphasizing the most essential technical and implementation details. Additional information and extended ablation studies regarding 2D and 3D inpainting, object and background generation, pose estimation and scene composition are available in the supplementary material.

3.1. Image Generation

3.1.1. 360° Panorama Diffusion

The initial stage of our pipeline generates the 360° panorama image I_F from a user-provided text prompt. Existing text to panorama diffusion models are typically LoRA-fine-tuned on small, less-diverse equirectangular datasets, limiting their ability to generalize to out-of-distribution styles/compositions despite these capabilities existing in the pre-trained base model. To improve generalization and style fidelity, we guide the generation process using a perspective image derived from the same prompt, providing soft conditioning without enforcing pixel-level alignment. We achieve this by using an IP-Adapter-style mechanism [56] that introduces separate cross-attention layers in all transformer blocks of the diffusion model. We observe that simply adding this mechanism post-hoc to a panoramic LoRA model is ineffective due to a mismatch in perspective (for IP-Adapter) and panoramic (for LoRA) training distributions as shown in the ablation in Figure 8. Instead, we *jointly*

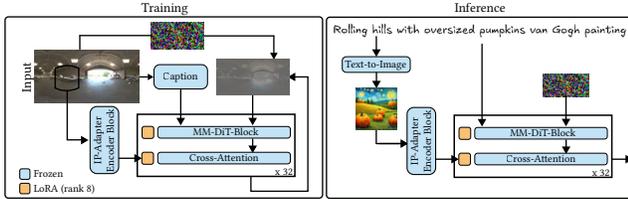


Figure 3. Our image generation pipeline uses a diffusion model to denoise an equirectangular image [13, 53], but additionally combines perspective image conditioning through a decoupled cross-attention layer [56]. We train an equirectangular LoRA jointly with the perspective pre-trained cross attention layer, by encoding random perspective crops of the panorama, showing stronger generalization capabilities to out-of-domain panorama sampling.

fine-tune the panoramic LoRA with the IP-adapter using random perspective renders of the equirectangular image (Figure 3), enabling effective style transfer from perspective to panoramic images.

We base our model on Stable-Diffusion-3.5-Large and fine-tune a rank-8 LoRA on top of all attention blocks on a mix of indoor [4] and outdoor [33] images. We train the model for 10k iterations with a learning rate of 10^{-5} and a batch size of 4.

3.1.2. Instance Segmentation

To identify suitable candidates for object reconstruction in the 360° images, we employ the off-the-shelf open-set segmentation pipeline *Grounded-SAM* [35]. To improve robustness in stylized or out-of-distribution scenes, we replace the original class identification stage (RAM [67]) with a larger vision-language model (GPT-4V), improving category recognition under diverse visual conditions. Given the limitations of current object generators, we exclude categories that are not good candidates for object reconstruction and are typically considered background elements (such as grass, walls, and floors). To further ensure high-quality segmentation, we sort all detected objects by their predicted IoU confidence and retain only the highest-scoring masks. For each retained object, the pipeline outputs a class identifier, a textual description, and a fine-grained segmentation mask, which are later used in Section 3.2.

3.1.3. 2D Inpainting

We generate the background panorama I_B by erasing the foreground objects identified in the full panorama using the union of their segmentation masks. To synthesize plausible new content in the masked regions, we employ a diffusion-based inpainting model, Dreamshaper-8-Inpainting [29], chosen for its ability to synthesize meaningful background content. We reduce hallucinations and artifacts in a two-stage process. First, we pre-inpaint the masked regions using LaMA [44] to obtain a coherent object-free initialization. Second we feed this through a custom LoRA fine-tuned

on panoramic images from the Structured3D dataset [68] for indoor defurnishing (in a similar spirit to [41]). To further improve image quality, we apply standard techniques: the segmentation masks are slightly dilated, wrap-around consistency is enforced via horizontal padding (256 pixels), and Poisson blending seamlessly integrates the original panorama with the newly inpainted content. Although we experimented with repurposing the panoramic LoRA for inpainting, we found that a dedicated, specialized pipeline significantly reduces hallucinations without requiring manual intervention, especially in indoor scenes. The resulting empty panorama image is then lifted to 3D to form the background point cloud as described in Section 3.3.

3.1.4. Depth Estimation

We employ different depth estimation strategies for indoor and outdoor scenes due to their distinct geometric challenges. For indoor scenes, we use EGformer [61], which excels in predicting the axis-aligned structure typical of indoor environments. For outdoor scenes, we instead use the 360MonoDepth framework [37], which aligns multiple perspective depth estimates obtained with Depth-Anything-V2 and possesses strong zero-shot generalization capabilities. To address unreliable depth predictions in low-detail regions, such as the sky, we explicitly detect and segment these areas using GroundedSAM [35] and assign the maximum depth value. By construction, we obtain two independent depth estimates for I_F and I_B and perform an alignment step from I_B to I_F to ensure that objects are placed properly in their corresponding locations. Similarly to concurrent work, we find that a global affine transform is not sufficient and instead employ a per-pixel shift and global scale that minimizes the $L2$ depth error in overlapping regions between the images, subject to as-rigid-as-possible constraints [42].

3.2. High Quality Object Generation

To enable stylistically aligned and geometrically consistent 3D object reconstructions, we condition image-to-3D models using RGBD data derived from the segmented panorama. However, segmented object images are often incomplete, low-resolution, seen from suboptimal viewpoints, or contain generative artifacts (such as distortions or incorrect limbs). To overcome these problems, we introduce a resynthesis step that produces a high-quality reference image of the object, utilizing both semantic and geometric cues. This image is then used to guide the 3D reconstruction of each object. Our pipeline consists of three stages: reference image generation, 3D object generation, and pose estimation for accurate placement in the scene. We illustrate this process in Figure 4.

3.2.1. Reference Image Generation

We synthesize a high-resolution reference image for each object by conditioning a diffusion model on the available

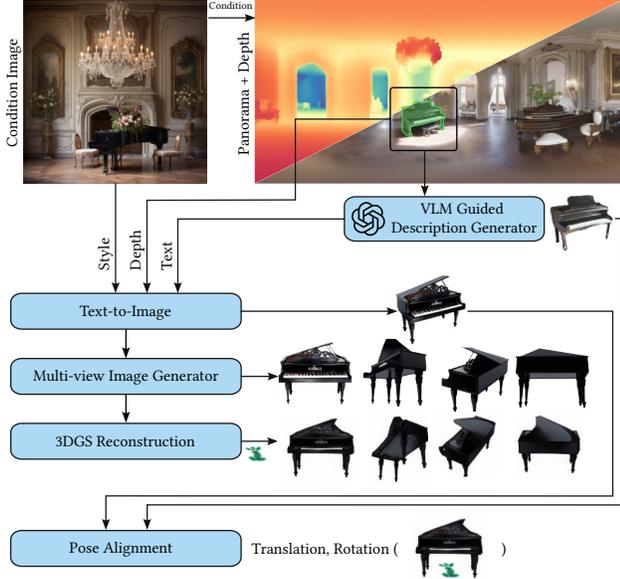


Figure 4. Our object reconstruction pipeline leverages the panorama and style information to generate a high-resolution reference image to be used for multi-view generation. The generated multi-view images are then transformed into 3D Gaussian splats through a reconstruction pipeline. Finally, we align the generated object with the original and place it in the scene.

multimodal cues. First, we extract a textual description of the object by prompting a vision-language model with a perspective view of a crop containing the object. This includes a description of the object’s visual attributes such as pose and color, serving as the primary conditioning signal. Second, we compute a partial depth map by combining the object’s segmentation mask with the estimated depth, offering geometric guidance for shape. Finally, we include a style image, reused from the panorama generation process, to ensure visual consistency and thematic coherence across the scene. As shown in Figure 5, our approach results in realistic and style-consistent object views. This approach is better suited for 3D reconstruction compared to other diffusion-based amodal completion methods (see Figure 11).

3.2.2. Image-based Object Generation

Our object reconstruction pipeline first generates six multi-view images at fixed camera poses, and then transforms them into a NeRF representation that captures both the geometry and color of the object. We found NeRF to be more robust than 3DGS as a differentiable representation when learning the reconstruction. Therefore, we opt to transform it into 3DGS in a separate post-processing step. We leverage Zero123++ [39] for the multi-view generation task and use the triplane decoder of InstantMesh [52] to convert the multi-view images into NeRF. In practice, generative models often introduce artifacts, including transparency, incomplete

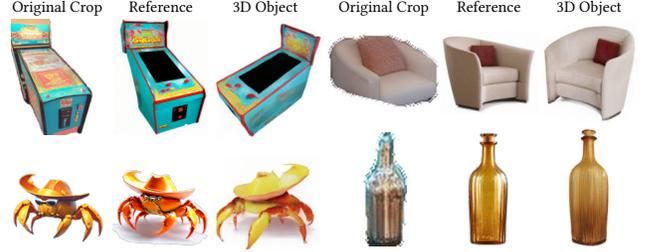


Figure 5. Qualitative evaluation of our object reconstruction module. Employing multimodal cues faithfully captures the stylistic and structural essence of low-fidelity or incomplete objects, reconstructing a suitable 3D counterpart to be later placed in the environment.

regions or structural inconsistencies. To enable a reliable and automated object reconstruction pipeline, we introduce a single-round refinement stage after generation. During this stage, we generate multiple object candidates and rank them by visual quality and alignment with the original input image using a VLM, similar to [71]. The VLM selects the best result from five candidates, one of which is generated using only the original cropped object image.

3.2.3. Pose Estimation

We decompose object pose estimation into two parts: the absolute pose of the original object O_{crop} with respect to the viewing origin and the relative transformation between O_{crop} and its high-quality counterpart O_{hq} . The absolute pose is estimated by projecting the crop’s midpoint along a ray from the camera origin to the estimated mean depth. Orientation and scale are heuristically derived based on the crop geometry and panorama alignment. This provides a coarse but coherent placement aligned with the image-plane and panorama geometry. To semantically align the point cloud of the high-quality object O_{hq} with the original O_{crop} we compute the relative pose between them. We employ MAST3R [23] on image pairs (perspective projections) of the original and high-quality object to obtain the relative pose. The final pose of O_{hq} preserves semantic alignment with the original object’s placement.

3.3. 3D Background Generation

Generating a 3D scene suitable for exploration from a single 360° image is an underconstrained problem, as the available information is limited to a single viewpoint. To address this, we employ a three-step process that reconstructs missing scene content through a combination of inpainting and score distillation sampling (SDS) [34]. First, we optimally initialize the 3DGS point cloud from the panoramic image and prepare disocclusions for inpainting through a pretuning stage; Second, we incrementally populate disocclusions with new Gaussians through a 3D inpainting process conditioned on the existing point cloud; Finally, we ensure multi-view consistency by applying SDS on a larger set of possible

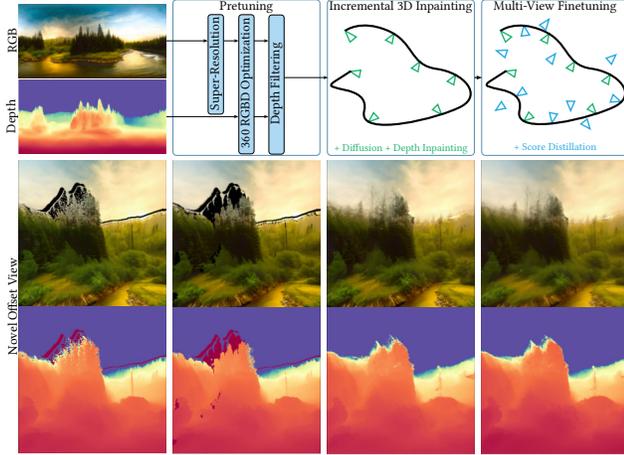


Figure 6. The background generation stage starts by identifying large occluded regions and preparing the 3DGS point cloud for inpainting, then proceeds to incrementally fill in new content and finally enforces multi-view consistency when rendering from novel views with score distillation.

views. During all optimization steps, we leverage MCMC-densification [22], a process demonstrated in Figure 6.

3.3.1. Background Initialization and Pretuning.

We initialize the 3DGS point cloud by unprojecting the pixels through a spherical camera placed at the origin, where the covariance matrices are initialized from normals derived from the depth map D_B , resulting in pixel-tight views when viewed from the origin. For pretuning, we optimize the initial Gaussians using mixed supervision from an upscaled version of the background panorama I_B (90%) and random views (10%). The I_B loss follows standard 3DGS supervision, while the random-view supervision focuses on identifying disoccluded regions.

3.3.2. Incremental Inpainting

To synthesize realistic new content in disoccluded regions, we leverage a pre-trained inpainting model conditioned on the rendered projections of the existing 3DGS splats. Specifically, we select a set of keyframes that target important regions with missing content. For each keyframe we render the point cloud and generate an inpainted RGB image and an inpainting mask, computed by applying max-pooling (kernel size 25) and thresholding ($T < 0.5$) on the transmittance map T , isolating low-opacity regions as targets for inpainting. For each pixel in the target region, we instantiate a new Gaussian by unprojecting it with a scene-aligned depth estimate. This depth is obtained using a diffusion model repurposed and fine-tuned for 3DGS depth inpainting [28], and provides a strong geometric prior. The inpainting process is done incrementally per keyframe, minimizing inconsistencies across overlapping views, similarly to [18].

3.3.3. Multi-view Fine-tuning

Finally, to consolidate the inpainted views with the existing 3D content, we employ mixed supervision from the panorama image (25%), the inpainted images (25%) and SDS [34] from random views around the inpainted regions (50%). This final stage of SDS fine-tuning helps to remove view-inconsistencies and ensures that the Gaussians introduced during inpainting are consistent and aligned with the existing scene as shown in Figure 6.

3.4. Composition

The last stage of the pipeline composes the final scene, by placing the 3D objects onto the background point cloud’s world frame, resulting in a navigable 3D scene. Since the scene consists of a set of splat parameters and their explicit transformations, post-processing effects and re-positioning of objects can be applied trivially, a key strength of the proposed system.

To further improve visual coherence, we address two key sources of error that can disrupt immersion. The first is the elimination of contact shadows—an unavoidable consequence of the mask dilation step used during object removal. Since achieving full control over lighting in diffusion models [10, 15] and relighting remain challenging problems, we reintroduce shadows through a splat-aware shadow mapping pass (details in the supplemental). The second issue involves inaccurate pose estimation, which can lead to objects violating physical plausibility. To mitigate this, we automatically snap object positions to dominant support planes (e.g. floor or walls), ensuring a more realistic and physically grounded scene integration.

4. Results

In this section we provide a quantitative and qualitative evaluation on the scene generation task, as well as an ablation of the core components of our pipeline. To assess the effectiveness of our approach in generating 3D scenes from text prompts, we focus on its ability to produce coherent and immersive scenes with high-quality and structurally consistent novel view synthesis. We compare our method against three state-of-the-art approaches for immersive scene generation: LayerPano3D [53], Text2Room [18], and DreamScene360 [71]. Additionally, in Figure 7 we provide a qualitative comparison with other methods that share a similar goal but would require substantial integration effort to directly compare against.

Our method is evaluated on 17 text prompts commonly found in recent literature [40, 53]. For each scene we recorded a 6-second walk-through video (at 1024×1024 px resolution) along a predefined camera trajectory *different* from the one used for the optimization described in Section 3.3, covering large parts of the scene’s geometry as well as occluded surfaces. Examples of video frames are shown

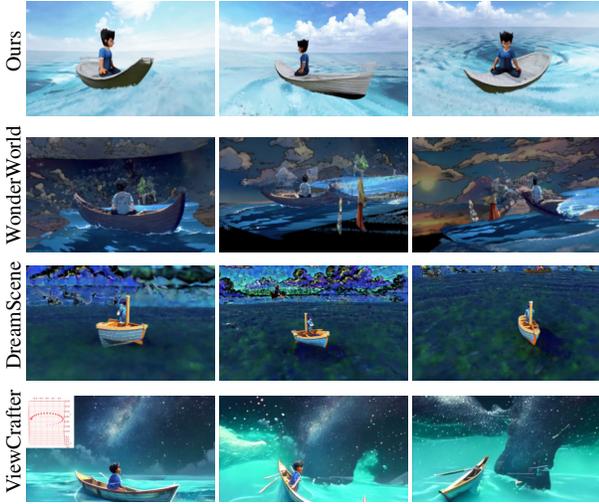


Figure 7. Qualitative evaluation against WonderWorld [59], DreamScene [25] and ViewCrafter [60]. WonderWorld degenerates into flat, view-dependent renderings when shown from out-of-context viewpoints, whereas DreamScene exhibits reduced fidelity in complex regions and object poses while expecting the positioning of objects to be known a priori. ViewCrafter combines video diffusion and reconstruction models but it is not able to maintain 3D coherence for long camera paths.



Figure 8. Qualitative ablation of out-of-domain sampling for panorama generator models. Our method captures the stylistic guidance faithfully, but without IP-Adapter it reverts back to its base panoramic distribution similarly to other recent methods. Simply using a perspective IP-Adapter does not work because of the mismatching distributions at training time.

in Figure 10. We generated all scenes on a single H100, with each scene taking around 15 mins to generate, given an average of 7.5 objects per scene. We include more results in the supplementary material.

4.1. Quantitative Evaluation

We quantitatively evaluate prompt alignment and image quality. To measure image quality and aesthetics, we use CLIP-IQA+ [48], Q-Align and A-Align [49] metrics. To quantify prompt alignment we compute the CLIPScore [17] between the prompt and each rendered frame. The results in Table 1 demonstrate that our method outperforms competing work in terms of both image quality and aesthetics metrics. For completeness, we also evaluate CLIP scores on single views

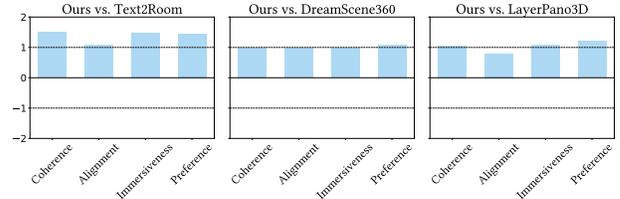


Figure 9. User study results for pairwise comparisons between our method, DreamScene360, Text2Room and LayerPano3D. Ratings are given on a 5-point Likert scale capturing strong, weak and no preference between the approaches. The participants rated our approach higher than the comparison methods on all axes.

Table 1. Quantitative evaluation of our method against Text2Room [18], DreamScene360 [71] and Layerpano3D [53]. Metrics are averaged across all scenes and computed on 136 frames per scene. Our method has state-of-the-art performance in terms of perceptual image quality and visual aesthetics. The differences in CLIP scores (evaluated against the original text prompt) are small and should be interpreted with care (see main text).

Method	CLIP \uparrow	Q-ALIGN \uparrow	A-ALIGN \uparrow	CLIP-IQA \uparrow
Ours	20.565	3.427	3.121	0.461
Text2Room	20.779	2.998	<u>2.872</u>	0.432
DreamScene360	20.615	2.510	2.517	0.347
Layerpano3D	<u>20.721</u>	<u>3.383</u>	2.833	<u>0.421</u>

Table 2. Ablation of the perspective conditioned IP-adapter for panorama generation. Out-of-domain prompt following improves without compromising generation quality. FID, KID and IS are evaluated on the test split of Matterport3d [4]. The relative difference in distribution similarity is significantly less relevant than CLIP improvement.

Method	CLIP \uparrow	CLIP-persp \uparrow	FID \downarrow	KID \downarrow	IS \uparrow
Full model	26.49	25.18	73.09	2.04	7.51
w/o IP-adapter	24.82	24.10	74.34	2.40	7.14
w/o joint fine-tuning	17.54	14.48	437.67	33.91	2.98

against the scene’s text prompt. However, individual views typically capture only a small portion of the scene, making CLIP scores a less reliable metric for our novel view evaluation.

Figure 10 provides a side-by-side comparison of the scenes generated by the different methods. In general, the results from Text2Room exhibit repetitive content and resemble stitched two-dimensional paintings—particularly in outdoor scenes—which limits the creation of immersive environments. However, the resulting images naturally excel in sharpness and detail, due to the method’s use of meshes for 3D representation. LayerPano3D and DreamScene360 demonstrate some capability for novel view synthesis, but struggle when scenes are viewed with significant offsets, such as sideway perspectives of objects.



Figure 10. Qualitative comparison with Text2Room [18], DreamScene360 [71] and LayerPano3D [53]. Whereas scenes from these competing method suffer from disturbing artifacts for significant viewpoint changes, our method achieves coherent novel view synthesis for virtually all viewpoints.

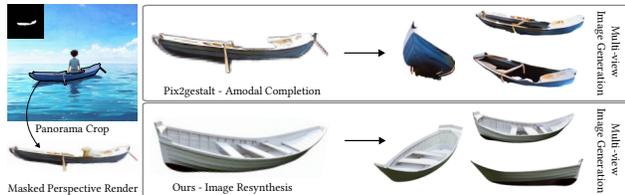


Figure 11. Diffusion-based amodal completion methods (such as pix2gestalt [30]) face challenges in producing high-quality inputs for multi-view image generation. Our method prioritizes structural fidelity by resynthesizing the reference image to yield higher quality 3D reconstructions.

4.2. User Study

To assess the effectiveness of our model in generating immersive and navigable 3D scenes, we conducted a user study with 28 participants, including graduate students and professionals in graphics and vision. Participants evaluated short videos of scenes on: *Coherence*: Do the spatial layout, object placement, and interactions in the scene form a logical and believable whole?; *Compatibility*: Do the objects, materials, lighting, and textures align well with the given text prompt?; *Immersiveness*: Does the scene convincingly represent a 3D environment, evoking a sense of depth, realism and presence?; *Overall Preference*. Each comparison was rated from -2 to $+2$. As shown in Figure 9, our method was consistently preferred across all metrics and in particular was praised for its geometric structure in novel view synthesis, as shown by the *coherence* and *immersiveness* scores. The results were statistically significant ($p < .001$) according to the Wilcoxon signed-rank tests.

4.2.1. Ablation

We present a quantitative ablation of the panorama generator in Table 2, highlighting the need for joint fine-tuning as well as the substantial prompt adherence obtained with the introduction of the IP-Adapter mechanism. A qualitative

comparison with [53] and [63] is shown in Figure 8, noting that each model is built on top of different foundation models. Despite relying on only a subset of the public data used by [53], our approach allows for more style-adhering images.

Our object generation module effectively addresses several common challenges in object extraction from images, such as panoramic distortion, low-resolution crops, and imperfect segmentation masks. Examples of these issues are illustrated in Figure 11. One alternative is to apply super-resolution and amodal completion but this approach typically requires multiple specialized modules, each with its own issues (Figure 5). Generating the same level of detail with super-resolution techniques on the panoramic image is still a challenge.

5. Conclusion

This work presented DreamAnywhere, a novel framework for generating immersive 3D scenes from text prompts by leveraging 360° images as foundational priors. Its modular, object-centric pipeline separates background and object generation, enabling high-fidelity object reconstruction and coherent scene completion through hybrid inpainting techniques. The modular nature of our pipeline has the disadvantage of having multiple possible points of failure. Conversely, this very modularity, serves as a significant advantage for targeted user intervention, which is *often desired in generative applications*. It also facilitates replacement of individual modules (e.g. the image generator method with [7, 21, 53], the 3D inpainting step with [60] or the object reconstruction pipeline with [55]). Compared to existing methods, our framework significantly enhances novel view synthesis coherence and offers robust performance for navigable 3D environments, as validated by quantitative results and strong user preference. This approach facilitates downstream applications like scene editing and simulation, paving the way for more scalable and interactive 3D content creation.

References

- [1] Hassan Abu Alhaija, Jose Alvarez, Maciej Bala, Tiffany Cai, Tianshi Cao, Liz Cha, Joshua Chen, Mike Chen, Francesco Ferroni, Sanja Fidler, et al. Cosmos-transfer1: Conditional world generation with adaptive multimodal control, 2025.
- [2] Andreea Ardelean, Mert Özer, and Bernhard Egger. Generalizable 3d scene reconstruction via divide and conquer from a single view, 2025.
- [3] Iro Armeni, Zhi-Yang He, JunYoung Gwak, Amir R Zamir, Martin Fischer, Jitendra Malik, and Silvio Savarese. 3d scene graph: A structure for unified semantics, 3d space, and camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5663–5672, 2019.
- [4] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Nießner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. pages 667–676, 2017.
- [5] Luxi Chen, Zihan Zhou, Min Zhao, Yikai Wang, Ge Zhang, Wenhao Huang, Hao Sun, Ji-Rong Wen, and Chongxuan Li. Flexworld: Progressively expanding 3d scenes for flexible-view synthesis, 2025.
- [6] Yongwei Chen, Tengfei Wang, Tong Wu, Xingang Pan, Kui Jia, and Ziwei Liu. Comboverse: Compositional 3d assets creation using spatially-aware diffusion guidance. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 128–146, 2025.
- [7] Zhaoxi Chen, Guangcong Wang, and Ziwei Liu. Text2light: Zero-shot text-driven hdr panorama generation. *ACM Trans. Graph.*, 41(6), 2022.
- [8] Zhaoxi Chen, Guangcong Wang, and Ziwei Liu. Scenedreamer: Unbounded 3d scene generation from 2d image collections, 2023.
- [9] Dana Cohen-Bar, Elad Richardson, Gal Metzger, Raja Giryes, and Daniel Cohen-Or. Set-the-scene: Global-local training for generating controllable nerf scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 2920–2929, 2023.
- [10] Riccardo Corvi, Davide Cozzolino, Giada Zingarini, Giovanni Poggi, Koki Nagano, and Luisa Verdoliva. On the detection of synthetic images generated by diffusion models. In *ICASSP 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, 2023.
- [11] Dave Epstein, Ben Poole, Ben Mildenhall, Alexei A. Efros, and Aleksander Holynski. Disentangled 3d scene generation with layout learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2024.
- [12] Chuan Fang, Yuan Dong, Kunming Luo, Xiaotao Hu, Rakesh Shrestha, and Ping Tan. Ctrl-room: Controllable text-to-3d room meshes generation with layout constraints, 2025.
- [13] Mengyang Feng, Jinlin Liu, Miaomiao Cui, and Xuansong Xie. Diffusion360: Seamless 360 degree panoramic image generation based on diffusion models, 2023.
- [14] Rafail Fridman, Amit Abecasis, Yoni Kasten, and Tali Dekel. Scenescape: Text-driven consistent scene generation. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2024.
- [15] Lanqing Guo, Chong Wang, Wenhan Yang, Siyu Huang, Yufei Wang, Hanspeter Pfister, and Bihan Wen. Shadowdiffusion: When degradation prior meets diffusion model for shadow removal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14049–14058, 2023.
- [16] Haonan Han, Rui Yang, Huan Liao, Jiankai Xing, Zunnan Xu, Xiaoming Yu, Junwei Zha, Xiu Li, and Wanhua Li. Reparo: Compositional 3d assets generation with differentiable 3d layout alignment, 2024.
- [17] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. CLIPScore: A reference-free evaluation metric for image captioning, 2021.
- [18] Lukas Höllein, Ang Cao, Andrew Owens, Justin Johnson, and Matthias Nießner. Text2room: Extracting textured 3d meshes from 2d text-to-image models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7909–7920, 2023.
- [19] Tianyu Huang, Wangguandong Zheng, Tengfei Wang, Yuhao Liu, Zhenwei Wang, Junta Wu, Jie Jiang, Hui Li, Rynson WH Lau, Wangmeng Zuo, et al. Voyager: Long-range and world-consistent video diffusion for explorable 3d scene generation. *arXiv preprint arXiv:2506.04225*, 2025.
- [20] Zehuan Huang, Yuan-Chen Guo, Xingqiao An, Yunhan Yang, Yangguang Li, Zi-Xin Zou, Ding Liang, Xihui Liu, Yan-Pei Cao, and Lu Sheng. Midi: Multi-instance diffusion for single image to 3d scene generation, 2025.
- [21] Nikolai Kalischek, Michael Oechsle, Fabian Manhardt, Philipp Henzler, Konrad Schindler, and Federico Tombari. Cubediff: Repurposing diffusion-based image models for panorama generation, 2025.
- [22] Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Yang-Che Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 3D Gaussian Splatting as Markov Chain Monte Carlo. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2024.
- [23] Vincent Leroy, Yohann Cabon, and Jerome Revaud. Grounding image matching in 3d with mast3r. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024.
- [24] Chenghao Li, Chaoning Zhang, Joseph Cho, Atish Waghvase, Lik-Hang Lee, Francois Rameau, Yang Yang, Sung-Ho Bae, and Choong Seon Hong. Generative ai meets 3d: A survey on text-to-3d in aigc era, 2023.
- [25] Haoran Li, Haolin Shi, Wenli Zhang, Wenjun Wu, Yong Liao, Lin Wang, Lik-Hang Lee, and Peng Yuan Zhou. Dreamscene: 3d gaussian-based text-to-3d scene generation via formation pattern sampling. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 214–230, 2025.
- [26] Yixun Liang, Xin Yang, Jiantao Lin, Haodong Li, Xiaogang Xu, and Yingcong Chen. Luciddreamer: Towards high-fidelity text-to-3d generation via interval score matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6517–6526, 2024.
- [27] Andrew Liu, Richard Tucker, Varun Jampani, Ameesh Makadia, Noah Snavely, and Angjoo Kanazawa. Infinite nature: Perpetual view generation of natural scenes from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.

- [28] Zhiheng Liu, Hao Ouyang, Qiuyu Wang, Ka Leong Cheng, Jie Xiao, Kai Zhu, Nan Xue, Yu Liu, Yujun Shen, and Yang Cao. Infusion: Inpainting 3d gaussians via learning depth completion from diffusion prior, 2024.
- [29] Lykon. Dreamshaper v8 inpainting. <https://huggingface.co/Lykon/dreamshaper-8-inpainting>, 2023.
- [30] Ege Ozguroglu, Ruoshi Liu, Dídac Surís, Dian Chen, Achal Dave, Pavel Tokmakov, and Carl Vondrick. pix2gestalt: Amodal segmentation by synthesizing wholes. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3931–3940, 2024.
- [31] Avinash Paliwal, Xilong Zhou, Andrii Tsarov, and Nima Khademi Kalantari. Panodreamer: 3d panorama synthesis from a single image, 2024.
- [32] Ryan Po, Wang Yifan, Vladislav Golyanik, Kfir Aberman, Jonathan T Barron, Amit Bermano, Eric Chan, Tali Dekel, Aleksander Holynski, Angjoo Kanazawa, et al. State of the art on diffusion models for visual computing. *Computer Graphics Forum*, 43(2):e15063, 2024.
- [33] PolyHaven. Polyhaven, 2025.
- [34] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D Diffusion. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- [35] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, Zhaoyang Zeng, Hao Zhang, Feng Li, Jie Yang, Hongyang Li, Qing Jiang, and Lei Zhang. Grounded sam: Assembling open-world models for diverse visual tasks, 2024.
- [36] Xuanchi Ren, Tianchang Shen, Jiahui Huang, Huan Ling, Yifan Lu, Merlin Nimier-David, Thomas Müller, Alexander Keller, Sanja Fidler, and Jun Gao. Gen3c: 3d-informed world-consistent video generation with precise camera control, 2025.
- [37] Manuel Rey-Area, Mingze Yuan, and Christian Richardt. 360MonoDepth: High-resolution 360° monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3752–3762, 2022.
- [38] Jonas Schult, Sam Tsai, Lukas Höllein, Bichen Wu, Jialiang Wang, Chih-Yao Ma, Kunpeng Li, Xiaofang Wang, Felix Wimbauer, Zijian He, Peizhao Zhang, Bastian Leibe, Peter Vajda, and Ji Hou. Controlroom3d: Room generation using semantic proxy rooms. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6201–6210, 2024.
- [39] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base model, 2023.
- [40] Jaidev Shriram, Alex Trevithick, Lingjie Liu, and Ravi Ramamoorthi. Realdreamer: Text-driven 3d scene generation with inpainting and depth diffusion, 2024.
- [41] Mira Slavcheva, Dave Gausebeck, Kevin Chen, David Buchhofer, Azwad Sabik, Chen Ma, Sachal Dhillon, Olaf Brandt, and Alan Dolhasz. An empty room is all we want: Automatic defurnishing of indoor panoramas. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 7384–7394, 2024.
- [42] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Proceedings of EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing*, pages 109–116, 2007.
- [43] Gabriela Ben Melech Stan, Diana Wofk, Estelle Aflalo, Shao-Yen Tseng, Zhipeng Cai, Michael Paulitsch, and Vasudev Lal. Ldm3d-vr: Latent diffusion model for 3d vr, 2023.
- [44] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3172–3182, 2022.
- [45] Shitao Tang, Fuyang Zhang, Jiacheng Chen, Peng Wang, and Yasutaka Furukawa. Mvdifusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2024.
- [46] Xiang Tang, Ruotong Li, and Xiaopeng Fan. Recent advance in 3d object and scene generation: A survey, 2025.
- [47] Shubham Tulsiani, Saurabh Gupta, David Fouhey, Alexei A. Efros, and Jitendra Malik. Factoring shape, pose, and layout from the 2d image of a 3d scene. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 302–310, 2018.
- [48] Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *AAAI*, 2023.
- [49] 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. Q-align: Teaching Imms for visual scoring via discrete text-defined levels, 2023. Equal Contribution by Wu, Haoning and Zhang, Zicheng. Project Lead by Wu, Haoning. Corresponding Authors: Zhai, Guangtai and Lin, Weisi.
- [50] Tianhao Wu, Chuanxia Zheng, and Tat-Jen Cham. Panodifusion: 360-degree panorama outpainting via diffusion. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- [51] Zhennan Wu, Yang Li, Han Yan, Taizhang Shang, Weixuan Sun, Senbo Wang, Ruikai Cui, Weizhe Liu, Hiroyuki Sato, Hongdong Li, and Pan Ji. Blockfusion: Expandable 3d scene generation using latent tri-plane extrapolation. *ACM Transactions on Graphics*, 43(4), 2024.
- [52] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models, 2024.
- [53] Shuai Yang, Jing Tan, Mengchen Zhang, Tong Wu, Gordon Wetzstein, Ziwei Liu, and Dahua Lin. Layerpano3d: Layered 3d panorama for hyper-immersive scene generation. pages 1–10, 2025.
- [54] Xiuyu Yang, Yunze Man, Jun-Kun Chen, and Yu-Xiong Wang. Scenecraft: Layout-guided 3d scene generation. In *Proceed-*

- ings of the Conference on Neural Information Processing Systems (NeurIPS), 2024.
- [55] Kaixin Yao, Longwen Zhang, Xinhao Yan, Yan Zeng, Qixuan Zhang, Lan Xu, Wei Yang, Jiayuan Gu, and Jingyi Yu. Cast: Component-aligned 3d scene reconstruction from an rgb image. *ACM Transactions on Graphics (TOG)*, 44(4): 1–19, 2025.
- [56] Hu Ye, Jun Zhang, Sibao Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models, 2023.
- [57] Weicai Ye, Chenhao Ji, Zheng Chen, Junyao Gao, Xiaoshui Huang, Song-Hai Zhang, Wanli Ouyang, Tong He, Cairong Zhao, and Guofeng Zhang. Diffpano: Scalable and consistent text to panorama generation with spherical epipolar-aware diffusion. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2024.
- [58] Hong-Xing Yu, Haoyi Duan, Junhwa Hur, Kyle Sargent, Michael Rubinstein, William T. Freeman, Forrester Cole, Deqing Sun, Noah Snavely, Jiajun Wu, and Charles Herrmann. Wonderjourney: Going from anywhere to everywhere. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [59] Hong-Xing Yu, Haoyi Duan, Charles Herrmann, William T Freeman, and Jiajun Wu. Wonderworld: Interactive 3d scene generation from a single image, 2025.
- [60] Wangbo Yu, Jinbo Xing, Li Yuan, Wenbo Hu, Xiaoyu Li, Zhipeng Huang, Xiangjun Gao, Tien-Tsin Wong, Ying Shan, and Yonghong Tian. Viewcrafter: Taming video diffusion models for high-fidelity novel view synthesis, 2024.
- [61] Ilwi Yun, Chanyong Shin, Hyunku Lee, Hyuk-Jae Lee, and Chae Eun Rhee. Egformer: Equirectangular geometry-biased transformer for 360 depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6078–6089, 2023.
- [62] Shangjin Zhai, Zhichao Ye, Jialin Liu, Weijian Xie, Jiaqi Hu, Zhen Peng, Hua Xue, Danpeng Chen, Xiaomeng Wang, Lei Yang, et al. Stargen: A spatiotemporal autoregression framework with video diffusion model for scalable and controllable scene generation, 2025.
- [63] Cheng Zhang, Qianyi Wu, Camilo Cruz Gambardella, Xiaoshui Huang, Dinh Phung, Wanli Ouyang, and Jianfei Cai. Taming stable diffusion for text to 360°panorama image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [64] Jingbo Zhang, Xiaoyu Li, Ziyu Wan, Can Wang, and Jing Liao. Text2nerf: Text-driven 3d scene generation with neural radiance fields. *IEEE Transactions on Visualization and Computer Graphics*, 30:7749–7762, 2023.
- [65] Qihang Zhang, Chaoyang Wang, Aliaksandr Siarohin, Peiye Zhuang, Yinghao Xu, Ceyuan Yang, Dahua Lin, Bo Dai, Bolei Zhou, Sergey Tulyakov, and Hsin-Ying Lee. SceneWiz3D: Towards text-guided 3D scene composition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [66] Shengjun Zhang, Jinzhao Li, Xin Fei, Hao Liu, and Yueqi Duan. Scene splatter: Momentum 3d scene generation from single image with video diffusion model. *IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR)*, 2025.
- [67] Youcai Zhang, Xinyu Huang, Jinyu Ma, Zhaoyang Li, Zhaochuan Luo, Yanchun Xie, Yuzhuo Qin, Tong Luo, Yaqian Li, Shilong Liu, et al. Recognize anything: A strong image tagging model, 2023.
- [68] Jia Zheng, Junfei Zhang, Jing Li, Rui Tang, Shenghua Gao, and Zihan Zhou. Structured3d: A large photo-realistic dataset for structured 3d modeling. In *Proceedings of The European Conference on Computer Vision (ECCV)*, 2020.
- [69] Haiyang Zhou, Xinhua Cheng, Wangbo Yu, Yonghong Tian, and Li Yuan. Holodreamer: Holistic 3d panoramic world generation from text descriptions, 2024.
- [70] Junsheng Zhou, Yu-Shen Liu, and Zhizhong Han. Zero-shot scene reconstruction from single images with deep prior assembly. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2024.
- [71] Shijie Zhou, Zhiwen Fan, Dejia Xu, Haoran Chang, Pradyumna Chari, Tejas Bharadwaj, Suyu You, Zhangyang Wang, and Achuta Kadambi. Dreamscene360: Unconstrained text-to-3d scene generation with panoramic gaussian splatting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 324–342, 2025.