

WorldSculpt: Generating Compositional Worlds from Grounded Videos

Muyao Niu[△]♣, Jixuan He[△], Ruihan Yu[△], Lian Fu[△], Yonghao Yu[△], Zheng-Hui Huang[△], Yifan Zhan[△], Fengbo Lan[△], Yongtao Ge[△], Yinqiang Zheng[♣], Kaipeng Zhang[△], Zhixiang Wang[△]

[△]Alaya Lab ♣The University of Tokyo

We study the problem of generating a compositional 3D representation of a cluttered scene containing hundreds of objects. The goal is to represent the scene as a collection of individual object meshes placed in a shared world frame, as required by downstream applications such as gaming, AR/VR, simulation, and robotics. This task is challenging in densely cluttered scenes, where objects heavily occlude one another and each view reveals only a fraction of their geometry. Geometry-based approaches typically reconstruct the scene as a single representation and leave incomplete geometry in occluded regions, while existing compositional methods with generative priors are largely limited to relatively simple scenes. We show that complex scenes with hundreds of objects can instead be generated compositionally by adapting a strong single-object 3D generative prior to multi-view observations. We instantiate this paradigm with Pixa3D, extending it with a multi-view conditioning pathway that grounds object generation in multiple posed observations. Although the model is finetuned entirely on single objects in canonical space, it generalizes to large scenes with severe occlusion without any scene-level training, demonstrating the feasibility and scalability of this paradigm. We further introduce UE-MeshyScene, a photorealistic benchmark of densely cluttered scenes with hundreds of objects, per-object annotations, and ground-truth meshes. Across single-object, controlled multi-object, and UE-MeshyScene evaluations, our method consistently outperforms prior approaches, with larger gains as scene complexity and occlusion increase. Finally, we demonstrate broader applicability by converting generated 3DGS worlds, such as Marble and HY-World 2.0, into compositional mesh scenes.

Project page: <https://alaya-lab.github.io/WorldSculpt>

Code: <https://github.com/AlayaLab/WorldSculpt>

Correspondence: Zhixiang Wang (Project Lead), Kaipeng Zhang

Date: September 4, 2026



1 Introduction

Recent generative world models have demonstrated impressive capabilities in synthesizing persistent and explorable 3D environments from only a single image or a short text prompt. Systems such as Marble [104] and HY-World 2.0 [90] can generate plausible geometry and appearance well beyond the directly observed content. However, the resulting world is typically represented as a unified scene representation, such as a fused mesh or a set of Gaussians, without explicitly separating individual objects. As a result, objects such as a chair or a vase cannot be independently selected, moved, or transferred into a physics simulator, since they are not represented as separate entities. This limitation creates a fundamental mismatch with the requirements of many downstream applications, including gaming and content creation, AR/VR, simulation, and robotics. These applications do not consume an undifferentiated soup of surface geometry. Instead, they require a scene to be decomposed into its constituent objects, with each represented by a mesh that can be independently selected, moved, and re-simulated.

Generating compositional 3D content from multiple images of a real cluttered scene remains a long-standing and largely unsolved problem in computer vision. In this setting, the structure of the output is just as important as its geometric accuracy. The task is particularly challenging in complex scenes, where many objects are densely packed and heavily occlude one another, leaving only partial observations of each object




Figure 1 WorldSculpt. Given a set of RGB images with instance-level 2D masks and 3D bounding boxes, WorldSculpt produces a compositional mesh representation for very complex scenes consisting of hundreds of individual objects.

in any individual view. A faithful result must therefore align the visible evidence across images, recover plausible geometry for unobserved regions, and preserve each object as an independently usable asset.

Existing approaches address different aspects of this problem, but none fully satisfy all of these requirements. Geometry-centric methods, including per-scene optimization, feed-forward volumetric or Gaussian regressors, and pointmap foundation models [99, 95, 45], can recover accurate geometry in observed regions. However, they typically reconstruct the scene as a single fused representation and leave missing geometry in areas hidden by occlusion. Generative priors offer a complementary strength. By learning a distribution over plausible 3D content, they can infer coherent geometry beyond the visible evidence, but existing methods generally fall into two categories, each with its own limitation. Scene-level world generation models such as Marble process the entire scene jointly and can synthesize unobserved regions, yet their output is still a single monolithic representation. Native-3D generative priors [108, 123, 42], on the other hand, model complete individual objects, but are designed for a single pose-free image of one object. They do not jointly condition on a consistent set of views or generate objects directly in a known scene coordinate frame, making them difficult to apply to complex scenes without additional machinery. Recent compositional approaches [29, 57, 82, 46] instead extend object-level generative priors to multi-object scenes, either by jointly generating multiple instances or by combining object generation with scene-level layout estimation. In practice, however, these methods have primarily been demonstrated on relatively simple scenes that can be described by a single image, *e.g.*, a small collection of objects on a tabletop or a sparse indoor furniture arrangement.

In this paper, we show that a complex scene containing hundreds of objects can be compositionally generated from a strong single-view, single-object generative prior. We take Pixel3D [42] as a case study to demonstrate the feasibility and scalability of this paradigm. Our key idea is to retain its object-level generative prior while extending it with a multi-view conditioning pathway. For each object, we first construct an anchor-aligned canonical frame from its scene observations, then lift per-view DINOv3 [84] features into the corresponding canonical voxel volume. Features from different views are fused with an IBR-style aggregator [96, 79] that is permutation-invariant and supports a variable number of inputs. The aggregated 3D condition is injected into the frozen prior through zero-initialized projection layers, while low-rank adapters (LoRA) [26] adapt the pretrained network to exploit the additional multi-view evidence. As a result, unobserved regions can be plausibly generated rather than left incomplete, while

the recovered geometry remains grounded in the available multi-view evidence. Each object is generated as an individual mesh in its anchor-aligned canonical frame and then placed into the scene through its canonical-to-world transformation, without cross-object fusion or joint shape optimization. Importantly, this formulation requires no scene-level training. The prior is finetuned entirely on individual objects in canonical space, yet generalizes at test time to large scenes containing hundreds of densely occluded objects. We further introduce a conditioning-view augmentation curriculum to improve robustness to the partial and degraded observations commonly encountered in cluttered scenes.

To enable meaningful evaluation on genuinely cluttered scenes, we introduce *UE-MeshyScene*, a benchmark of large-scale indoor and outdoor environments rendered in Unreal Engine. Each scene contains up to several hundred assets arranged in natural configurations with substantial and complex mutual occlusion, together with annotated camera parameters, 3D bounding boxes, and ground-truth meshes for individual objects. UE-MeshyScene provides both the realism and scale of complex scenes and the exact geometric ground truth that is difficult to obtain from real-world captures. We evaluate our method on controlled single-object stress tests using Toys4k [86], controlled multi-object scenes including the real-captured HouseCat6D [34] and synthetic Toys4k [86]-Scene datasets, as well as our UE-MeshyScene benchmark. Across all settings, our method consistently outperforms prior approaches, with larger gains as scenes become increasingly crowded and occluded.

In summary, our contributions are:

- We propose *WorldSculpt*, a framework for generating compositional 3D scenes containing hundreds of objects, each represented by an individual mesh. WorldSculpt maps multi-view scene observations into an anchor-aligned canonical frame, conditions a strong object-level 3D generative prior on the aligned observations, and places the generated meshes into a shared world frame through canonical-to-world transformations.
- We show that a generative prior finetuned entirely on individual objects in canonical space can generalize to large-scale scenes containing hundreds of densely occluded objects, without any scene-level training, through multi-view conditioning and a targeted augmentation curriculum.
- We introduce *UE-MeshyScene*, a photorealistic benchmark of densely cluttered scenes containing hundreds of objects, with exact per-object annotations and ground-truth meshes for evaluating compositional 3D generation under complex occlusion.

2 Related Work

Image and multi-view to 3D generation. Feed-forward 3D reconstruction methods directly infer geometry from one or a small number of input images. LRM [25] and subsequent mesh- and Gaussian-based variants, including MeshLRM [102], InstantMesh [109], GRM [110], LGM [89], TripoSR [94], and SF3D [5], use large transformer-based architectures to predict meshes or Gaussian representations in a single forward pass. Another family of methods transfers 2D diffusion priors to 3D generation by first synthesizing novel views. Methods such as Zero-1-to-3 [50], SyncDreamer [52], Wonder3D [53], CRM [101], and One-2-3-45 [49] generate consistent multi-view observations that are subsequently fused into a 3D shape. Native-3D generative models instead learn distributions directly in 3D representations, including point clouds in Point-E [59], implicit functions in Shap-E [33], aligned shape-image-text latent spaces in Michelangelo [124], and more recent high-resolution structured representations used by TRELIS [108], Hunyuan3D [123], and Pixa3D [42]. Such models are commonly trained on large-scale 3D asset datasets such as Objaverse and Objaverse-XL [18, 17]. Complementary pose-free stereo approaches, including DUST3R [99] and MAST3R [41], jointly estimate scene geometry and camera poses and can provide the camera information required by our pipeline.

Amodal 3D reconstruction. Amodal 3D reconstruction seeks to recover the complete geometry of an object from partial observations affected by occlusion. Early work studies this problem in cluttered robotic environments, where incomplete 3D observations are completed using physical constraints such as object stability and connectivity [1]. Another direction performs completion first in image space, as

in pix2gestalt [69], and then reconstructs the completed image with an image-to-3D model. However, image-space completion alone does not explicitly enforce 3D consistency or consistency across multiple views. More recently, Amodal3R [106] adapts a native-3D generative prior to directly recover complete object geometry from partially occluded images. Concurrent work, AmodalGen3D [125], extends this setting to amodal object generation from sparse, unposed views.

Object-level 3D datasets. Evaluating image-to-3D methods requires datasets with ground-truth 3D geometry. Large synthetic or artist-created collections, including Objaverse and Objaverse-XL [18, 17] and Toys4k [86], provide clean object meshes for both training and held-out evaluation. Google Scanned Objects [20] and ABO [15] further provide scanned or catalog-based assets, although evaluation is typically performed on rendered images with clean backgrounds. Datasets such as CO3D [75] and MVImgNet [118] contain real captures with natural backgrounds and lighting, but provide SfM point clouds rather than ground-truth meshes. Datasets that pair real in-the-wild photographs with accurate scanned mesh geometry remain comparatively scarce.

Scene-level and compositional generation. Multi-object scene generation has been studied from both single images and real-world captures. Holistic image-based approaches, including InstPIFu [48], PartCrafter [46], MIDI [29], SAM3D [10], SceneGen [57], SceneMaker [82], and RecGen [119], jointly recover object geometry and scene layout, and are commonly evaluated on synthetic datasets such as 3D-FRONT/3D-FUTURE [22, 23] and Hypersim [77]. Real-world datasets, in contrast, typically provide either fused but incomplete scene-level meshes, as in ScanNet [16], ScanNet++ [116], Matterport3D [8], MultiScan [56], and Replica [87], or object-level annotations in the form of oriented bounding boxes, as in ARKitScenes [4], SUN RGB-D [85], and Objectron [2]. Other datasets provide aligned CAD models as object proxies, including Scan2CAD [3], ROCA [24], and CAD-Estate [55]. Accurate and complete per-object mesh ground truth for real, cluttered scenes, however, remains difficult to obtain.

Structure and motion from multiple images. Reconstructing objects from multiple images typically requires estimates of camera motion, coarse 3D structure, and object localization, for which a broad range of existing methods can be used. Camera poses are traditionally recovered with structure-from-motion [80, 81, 70] or deep visual SLAM [92, 93], while more recent systems target casual, low-parallax, and dynamic video specifically [44, 28]. Dense scene structure can be obtained from monocular geometry estimators, which have evolved from affine-invariant relative depth [73, 112, 113, 45] to metric depth and camera-aware point maps jointly estimated with intrinsics [117, 72, 98], as well as temporally consistent depth for long videos [27, 9]. More recently, feed-forward pointmap models have begun to unify camera estimation and geometry reconstruction by directly predicting both from unposed image sets or video streams [99, 41, 58, 122, 111, 97, 95, 35, 100, 45]. Object-level processing can be handled separately using promptable or open-world video segmentation [74, 7, 12, 11, 13], optionally initialized by open-vocabulary detectors [51, 76, 105, 14], to obtain temporally consistent object masks. These masks can then be combined with 3D detection or object-level SLAM methods [6, 78, 121, 39, 19, 114, 43, 103, 40], or with classical silhouette-based geometry when camera poses are known [38, 37], to recover object locations in 3D. Taken together, these components can convert raw monocular footage into camera trajectories, dense scene geometry, object mask tracks, and object-level 3D localizations.

3 Method

Problem setup. Given a set of N posed images $\{I_n\}_{n=1}^N$, we assume known camera intrinsics $\{K_n\}_{n=1}^N$ and camera-to-world extrinsics $\{T_n^{\text{cw}}\}_{n=1}^N$. For each object k , we assume access to a per-view instance mask S_{kn} whenever the object is visible [7, 74, 12, 11, 13, 105, 51, 76, 14], together with a coarse world-space localization box B_k^{loc} [6, 78, 121, 39, 19, 40, 114, 43, 103]. Recovering these masks, camera parameters, and coarse object localizations is well studied and lies outside the scope of this work.

Our goal is to produce a compositional scene representation

$$\mathcal{M} = \{(\mathcal{M}_k^c, T_k^{\text{ow}})\}_{k=1}^K, \quad (1)$$




Figure 2 Overview of WorldSculpt. Given posed scene images, instance masks, and coarse 3D object localizations, the anchor-aligned canonicalization stage constructs a virtual canonical cube for each object and derives crop-aware projections from its input views. Per-view DINOv3 features are lifted into the canonical voxel volume and fused with a permutation-invariant aggregator. The resulting 3D condition, together with global tokens from the anchor view, drives the two geometry stages of a Pixal3D prior through zero-initialized injection layers and LoRA adaptation. The generated canonical meshes are finally placed into the scene using their canonical-to-world transformations.

where $\mathcal{M}_k^c$ is an individual mesh generated in the canonical frame of object k , and T_k^{ow} maps that canonical frame into the world coordinate system. The corresponding world-space mesh is

$$\mathcal{M}_k^w = T_k^{\text{ow}}(\mathcal{M}_k^c). \quad (2)$$

The scene remains a collection of individually addressable object meshes rather than a single fused representation, making it directly suitable for downstream rendering, simulation, and content-authoring pipelines.

Our method consists of three main steps. We first construct an anchor-aligned virtual canonical frame for each object and map its scene observations into that frame. We then generate an individual object mesh with a multi-view conditioned 3D generative prior. Finally, the generated mesh is transformed back into the world frame using the same canonical-to-world transformation. Figure 2 provides an overview.

3.1 Anchor-aligned object canonicalization

Anchor view and canonical cube. The supplied coarse localization box B_k^{loc} is not directly used as the generation volume. Its side lengths are generally unequal, while the generative model operates in a normalized cubic domain. Its orientation is also determined by the localization procedure and need not agree with the camera-relative canonical orientation expected by the pretrained object prior.

For each object, we therefore construct an anchor-aligned virtual canonical cube. Let $\mathcal{I}_k$ denote the set of views in which object k is observed. We select one view $a_k \in \mathcal{I}_k$ as the *anchor*. At inference time, the anchor is chosen as the view in which the object is most fully observed. During training, it is sampled randomly from the available views. The anchor determines the orientation of the virtual canonical frame, with the anchor camera observing the object from the canonical front-view direction.

We use the normalized cube

$$\mathcal{V} = \left[-\frac{1}{2}, \frac{1}{2} \right]^3 \quad (3)$$

as the canonical spatial domain. Let c_k be the center of B_k^{loc} , and let s_k initially be its largest side length. The rotation $R_k \in \text{SO}(3)$ is induced by the anchor camera orientation so that the anchor view is mapped to the

canonical viewing convention. The canonical-to-world transformation is then

$$T_k^{\text{ow}} = \begin{bmatrix} s_k R_k & c_k \\ 0 & 1 \end{bmatrix}, \quad (4)$$

which maps a canonical point $x \in \mathcal{V}$ to

$$x_k^{\text{w}} = s_k R_k x + c_k. \quad (5)$$

Since s_k is isotropic, T_k^{ow} is a similarity transformation that preserves the proportions of the generated object.

Localization boxes may be inaccurate (*i.e.*, slightly loose or tight). We therefore allow s_k to increase while keeping c_k and R_k fixed until the projection of the resulting cube covers the object masks in all selected views. This adjustment prevents the object from being clipped by the per-view crops while preserving a single consistent canonical frame across observations.

Canonical observations and crop-aware projection. For each selected view $n \in \mathcal{I}_k$, we project the anchor-aligned cube into the image, crop the image to its projected extent, mask out the background using S_{kn} , and resize the crop to the input resolution of the generative model. We denote the resulting object-centric image by $\bar{I}_{kn}$.

Cropping and resizing change the image coordinate system. Let A_{kn} denote the corresponding transformation from the original image coordinates to the coordinates of the resized crop. The adjusted camera intrinsics are

$$\bar{K}_{kn} = A_{kn} K_n. \quad (6)$$

For a canonical voxel x , with homogeneous coordinate $\tilde{x} = [x^\top, 1]^\top$, its location in the cropped observation is

$$\pi_{kn}(x) = \Pi \left(\bar{K}_{kn} \begin{bmatrix} I_3 & 0 \end{bmatrix} (T_n^{\text{cw}})^{-1} T_k^{\text{ow}} \tilde{x} \right), \quad (7)$$

where $\Pi([u, v, d]^\top) = (u/d, v/d)^\top$ denotes perspective division. Equation (7) first maps the voxel from the object canonical frame into the world frame, then into the camera frame, and finally into the resized crop. It therefore accounts for the object’s position, orientation, scale, and off-center image location without assuming that the object is centered in any input view.

3.2 Multi-view conditioned object generation

Generative prior. We instantiate the per-object generator Φ from Pixa3D [42], a native-3D generative model built on the TRELLIS.2 structured-latent backbone [107]. Pixa3D generates an object through a cascade of flow-matching stages [47]. We use its first two stages, which determine object geometry. The *sparse structure* stage predicts coarse occupancy on a 64^3 grid, and the *shape* stage generates a high-resolution sparse structured latent over the occupied voxels. The resulting latents are decoded into a mesh. We focus on geometry and leave the subsequent texture and material stage to future work.

The original Pixa3D model consumes a single pose-free RGB image through cross-attention. Its output orientation, position, and scale are therefore implicit in the single-image training convention rather than explicitly tied to camera geometry. Our setting instead provides a coherent set of posed observations that must jointly constrain the generated object in a known frame. We retain the pretrained object prior and introduce a multi-view conditioning pathway that grounds its generation in the anchor-aligned canonical volume.

Multi-view feature lifting. Each canonical observation $\bar{I}_{kn}$ is independently encoded with DINOv3 [84], producing a dense feature map

$$F_{kn} = E_{\text{DINO}}(\bar{I}_{kn}). \quad (8)$$

For each voxel x in the canonical volume, the feature from view n is sampled at its projected image location:

$$g_n(x) = F_{kn}(\pi_{kn}(x)). \quad (9)$$

This operation lifts each 2D feature map into a 3D feature grid whose spatial coordinates are shared across all views. Observations of the same canonical location are therefore aligned before cross-view aggregation.

Permutation-invariant view aggregation. Let $\mathcal{J}_k \subseteq \mathcal{I}_k$ be the selected conditioning views for object k , and let $L_k = |\mathcal{J}_k|$. We aggregate the lifted features at each voxel using an IBRNet-style module [79, 96]. The module accepts a variable number of views and does not depend on their ordering.

For each canonical voxel x , we first compute the cross-view mean and variance:

$$\mu(x) = \frac{1}{L_k} \sum_{n \in \mathcal{J}_k} g_n(x), \quad (10)$$

$$\sigma^2(x) = \frac{1}{L_k} \sum_{n \in \mathcal{J}_k} g_n(x) \odot g_n(x) - \mu(x) \odot \mu(x), \quad (11)$$

where $\odot$ denotes element-wise multiplication. Two lightweight MLPs then refine each feature and predict its aggregation logit:

$$g'_n(x) = \text{MLP}_{\text{feat}} \left(\left[g_n(x), \mu(x), \sigma^2(x) \right] \right), \quad (12)$$

$$w_n(x) = \text{MLP}_{\text{weight}} \left(\left[g_n(x), \mu(x), \sigma^2(x) \right] \right). \quad (13)$$

The final voxel feature is the cross-view mean plus a softmax-weighted residual:

$$g_{\text{out}}(x) = \mu(x) + \sum_{n \in \mathcal{J}_k} \alpha_n(x) g'_n(x), \quad \alpha_n(x) = \frac{\exp(w_n(x))}{\sum_{m \in \mathcal{J}_k} \exp(w_m(x))}. \quad (14)$$

The final layer of MLP_{feat} is initialized to zero. The aggregator therefore begins as an exact cross-view mean and gradually learns view-dependent residual corrections during training. Applying this operation to every voxel produces the aggregated 3D conditioning grid G_k .

All selected views contribute to G_k , while the global image tokens used by the original Pixel3D cross-attention interface are extracted only from the anchor observation $\bar{I}_{ka_k}$. The anchor tokens preserve the single-image conditioning convention of the pretrained model, and the voxel-aligned grid introduces additional evidence from the remaining views.

Conditioning injection. We inject G_k into both Pixel3D geometry stages. At each transformer block, the aggregated condition is aligned with the current 3D token locations, projected to the block feature dimension, and added to the corresponding block features. We keep the original Pixel3D parameters frozen and adapt its attention and per-block projection layers using low-rank adapters (LoRA) [26]. The multi-view aggregator and the zero-initialized conditioning projections are trained in full. This design retains the completion ability of the pretrained object prior while allowing the denoiser to use spatially aligned evidence from multiple views.

3.3 Training

Random-view training. The sparse-structure and shape stages are trained independently using their respective ground-truth latent targets. For each training object, we sample a variable number of conditioning views and randomly choose one of them as the anchor. The sampled anchor defines the camera-relative canonical orientation for that training example, and the conditioning views and target latents are expressed in the same anchor-aligned frame. Varying the anchor and the number of observations exposes the model to different camera configurations while the aggregation architecture remains permutation-invariant.

Only the multi-view aggregator, the conditioning injection layers, and the LoRA parameters are optimized. All remaining parameters of the pretrained Pixel3D generator stay frozen.

Conditioning-view augmentation. The clean object renders used for training differ substantially from the partial and degraded observations found in cluttered scenes. We reduce this gap with a conditioning-view augmentation curriculum. The augmentations are applied only to the input observations and never to the supervision target.

We independently sample four types of degradation. First, we simulate occlusion using either random 2D masks or 3D-consistent occluders placed between the camera and the target object. The latter produce occlusion patterns that move coherently across views. Second, we perturb the camera poses of the non-anchor observations while keeping the anchor pose unchanged. Third, we introduce mask errors by perturbing and degrading the segmentation boundaries. Fourth, we downsample the observations to simulate objects that occupy only a small image region. The strength of each degradation is gradually increased from zero to its maximum value during training, allowing the model to first learn from clean observations before encountering more difficult inputs.

3.4 Inference and scene composition

At inference time, we process each object independently. We first select the anchor as the view in which the object is most fully observed and construct the anchor-aligned canonical cube described in Section 3.1. If the initial cube does not cover the object masks in all candidate views, its side length is increased while its center and orientation remain fixed. We then select a bounded number of conditioning views, prioritizing observations with larger visible object regions, and produce the corresponding masked crops and crop-aware projection functions.

Each object is generated by running the two Pixel3D geometry stages with the multi-view conditioning described in Section 3.2. The sparse-structure stage predicts the coarse occupied voxels, which define the sparse support used by the shape stage. The second-stage latent is then decoded into an individual canonical-space mesh $\mathcal{M}_k^c$.

The mesh is placed into the scene using the canonical-to-world transformation from Eq. (4):

$$\mathcal{M}_k^w = T_k^{\text{ow}}(\mathcal{M}_k^c). \quad (15)$$

The final compositional scene is the collection

$$\mathcal{M}^w = \{\mathcal{M}_k^w\}_{k=1}^K. \quad (16)$$

3.5 Implementation details

We initialize the two Pixel3D geometry stages from their released checkpoints and finetune them independently on the TexVerse dataset. Conditioning images are resized to 512px for the sparse-structure stage and 1024px for the shape stage. The LoRA adapters use rank $r=32$ and scaling factor $\alpha=32$ and are applied to the DiT attention and per-block projection layers. The multi-view aggregator and the zero-initialized conditioning projections are optimized without low-rank parameterization.

Each stage is trained for 15k iterations in bf16 using AdamW with a learning rate of 10^{-4} , betas (0.9, 0.95), no weight decay, and an exponential moving average of 0.9999. We use a batch size of 48 for the sparse-structure stage and 12 for the shape stage. At each training iteration, we randomly sample 1–20 views to form the multi-view conditioning set. To limit memory usage, the shape stage applies gradient checkpointing to half of its transformer blocks and caps each conditioning view at 32,768 tokens. The augmentation curriculum reaches its maximum strength over the first 3k training iterations.

4 Experiments

4.1 Datasets and Metrics

We evaluate our method under three settings of increasing difficulty. (i) We first evaluate canonical-space generation of individual objects on Toys4k [86] (Sec. 4.2), isolating the ability of the model to recover object geometry from partial multi-view observations. (ii) We then evaluate compositional scene generation on synthetic Toys4k-Scene and the real-world HouseCat6D [34] dataset (Sec. 4.3). (iii) Finally, we evaluate large-scale compositional generation on UE-MeshyScene, where scenes contain hundreds of objects under severe mutual occlusion (Sec. 4.4).

Table 1 Per-scene statistics of UE-MeshyScene. Every scene is rendered in Unreal Engine 5.8 at 2560×1440 ; “#Views” denotes the number of rendered frames.

	<i>Hangar</i>	<i>Abandoned City</i>	<i>Cathedral</i>	<i>Office</i>	<i>Japanese School</i>	<i>Desert Town</i>	Total
#Objects	171	93	145	678	511	701	2,299
#Views	463	265	956	1268	1758	1254	5,964




Figure 3 Visualization of the UE-MeshyScene dataset. Our benchmark contains six photorealistic environments spanning relatively orderly to extremely cluttered layouts, with up to 701 objects per scene. Each environment is captured along an orbiting camera trajectory and provides exact per-object ground truth, enabling world-frame evaluation under complex mutual occlusion.

Datasets. *Toys4k* [86] consists of clean, isolated object models and is used for controlled single-object evaluation. We vary both the number of input views (1–16) and the per-view occlusion fraction (0–75%), allowing us to measure how generation quality changes as observations become fewer and increasingly incomplete, without introducing scene-level factors.

Toys4k-Scene is a synthetic benchmark constructed by placing multiple *Toys4k* objects into cluttered layouts and rendering short orbiting sequences. It preserves the clean ground-truth geometry of *Toys4k* while introducing substantial inter-object occlusion and challenging multi-object configurations. Many objects contain thin or intricate structures and are densely arranged, leaving only partial observations from typical viewpoints. This setting is therefore particularly challenging for methods that rely on a single image.

HouseCat6D [34] provides a complementary real-world evaluation setting. It contains real captures of tabletop scenes together with scanned ground-truth meshes for individual objects. Compared with *Toys4k-Scene*, its layouts are relatively sparse, the objects are mostly geometrically simple household items, and inter-object occlusion is generally mild. It therefore allows us to evaluate how well a model trained on synthetic renders transfers to real imagery.

UE-MeshyScene is our large-scale benchmark for evaluating compositional generation in densely cluttered scenes. We describe its construction and statistics in detail below.

UE-MeshyScene. Existing scene benchmarks remain limited in either scale, occlusion, or the availability of individual per-object ground truth. *HouseCat6D* contains relatively sparse layouts with mild inter-

Table 2 Canonical-space generation on Toys4k [86] under varying numbers of input views (#V) and per-view occlusion levels. We report $CD-\ell_2$ ($\times 10^{-3}$, ↓), $CD-\ell_1$ ($\times 10^{-2}$, ↓), EMD ($\times 10^{-2}$, ↓), and F-Score@{.02, .05} (↑). The best result in each column is shown in **bold**.

Method	#V	Occl. 0%					Occl. 25%					Occl. 50%					Occl. 75%				
		$CD-\ell_2$	$CD-\ell_1$	EMD	F@.02	F@.05	$CD-\ell_2$	$CD-\ell_1$	EMD	F@.02	F@.05	$CD-\ell_2$	$CD-\ell_1$	EMD	F@.02	F@.05	$CD-\ell_2$	$CD-\ell_1$	EMD	F@.02	F@.05
TRELLIS.2 [107]	1	7.55	6.35	10.7	0.587	0.839	14.19	8.86	15.1	0.480	0.739	27.09	13.17	20.4	0.344	0.587	56.71	20.63	28.5	0.189	0.367
Pixal3D [42]	1	2.17	3.74	7.8	0.749	0.934	7.68	6.80	15.2	0.556	0.813	22.81	11.69	21.6	0.405	0.651	63.95	21.51	31.8	0.213	0.401
TRELLIS [108]	1	9.39	6.88	11.1	0.556	0.824	14.18	8.69	13.3	0.466	0.753	25.02	11.89	16.9	0.375	0.645	44.46	17.16	22.5	0.256	0.487
	2	8.29	6.50	10.7	0.574	0.837	14.87	8.62	13.0	0.474	0.758	21.31	11.34	16.3	0.368	0.649	39.15	16.58	22.1	0.244	0.481
	4	7.74	6.29	10.5	0.583	0.845	13.07	8.16	12.6	0.492	0.773	19.58	10.65	15.5	0.386	0.672	38.49	16.51	21.8	0.235	0.471
	8	8.91	6.51	10.7	0.579	0.839	11.53	7.73	12.2	0.508	0.789	18.05	10.25	15.1	0.399	0.683	33.81	15.68	21.4	0.241	0.480
	16	9.46	6.65	10.8	0.577	0.835	11.59	7.73	12.0	0.510	0.787	18.51	10.29	15.1	0.401	0.687	34.39	15.71	21.2	0.249	0.484
Ours	1	2.59	3.85	7.8	0.739	0.932	3.70	4.73	8.7	0.669	0.896	8.90	6.84	11.2	0.554	0.818	25.45	12.30	16.8	0.370	0.631
	2	1.54	3.23	7.2	0.790	0.954	2.57	3.87	7.9	0.734	0.931	4.01	4.85	8.8	0.651	0.893	10.03	7.96	12.1	0.494	0.762
	4	1.35	3.02	7.0	0.813	0.964	1.62	3.35	7.3	0.778	0.950	2.56	4.01	7.8	0.719	0.925	5.67	5.86	10.1	0.590	0.849
	8	1.23	2.94	6.9	0.822	0.965	1.50	3.20	7.2	0.796	0.956	1.91	3.57	7.5	0.755	0.943	3.52	4.65	8.6	0.663	0.902
	16	1.19	2.89	6.9	0.827	0.967	1.29	3.01	7.0	0.813	0.964	1.44	3.23	7.2	0.788	0.955	2.25	3.97	8.0	0.713	0.929

Table 3 Compositional scene generation on the real-world HouseCat6D [34] dataset and synthetic Toys4k [86]-Scene benchmark. We report $CD-\ell_2$ ($\times 10^{-3}$, ↓), $CD-\ell_1$ ($\times 10^{-2}$, ↓), EMD ($\times 10^{-2}$, ↓), and F-Score (↑). The best result in each column is shown in **bold**.

Method	HouseCat6D [34]				Toys4k-Scene [86]			
	$CD-\ell_2$ ↓	$CD-\ell_1$ ↓	EMD↓	F-Score↑	$CD-\ell_2$ ↓	$CD-\ell_1$ ↓	EMD↓	F-Score↑
SAM3D [10]	17.7	11.45	12.1	0.630	52.7	17.45	17.5	0.457
SceneGen [57]	31.4	13.08	15.9	0.626	25.3	11.79	15.6	0.664
SceneMaker [82]	16.3	12.98	12.3	0.519	26.8	16.85	17.3	0.422
RecGen [119]	7.22	6.98	8.6	0.805	13.2	9.74	11.0	0.676
ShapeR [83]	1.26	1.88	4.8	0.973	8.38	7.81	9.4	0.746
Ours	0.28	1.53	3.4	0.995	0.61	1.78	3.7	0.981

Table 4 Compositional generation on six scenes of the UE-MeshyScene dataset. Protocol and metrics as in Table 3 (F-Score at $\tau=0.05$), additionally reporting per-instance medians (CD^{med} , same scale as the corresponding mean). Best per column in **bold**.

Method	$CD-\ell_2$ ↓	$CD-\ell_2^{\text{med}}$ ↓	$CD-\ell_1$ ↓	$CD-\ell_1^{\text{med}}$ ↓	EMD↓	F-Score↑
ShapeR [83]	7.42	2.47	6.37	5.07	8.8	0.813
Ours	2.48	0.25	2.66	1.69	5.0	0.951

object occlusion, while Toys4k-Scene introduces substantially stronger occlusion but still contains only a modest number of objects. Realistic environments are considerably more challenging, as they may contain dense clutter and large numbers of mutually occluding objects. To evaluate this setting, we introduce UE-MeshyScene, a collection of six large-scale scenes authored and rendered in Unreal Engine. The scenes contain between 93 and 701 objects arranged naturally across shelves, desks, floors, and other cluttered regions, producing complex patterns of mutual occlusion. Since every asset is a known 3D model placed with a known transformation, UE-MeshyScene provides complete per-object mesh ground truth, exact camera poses, and 3D bounding boxes. This enables direct evaluation in a common world coordinate frame while retaining the complexity of densely cluttered environments.

UE-MeshyScene comprises six environments authored and rendered in Unreal Engine 5.8: an aircraft hangar, an abandoned-city interior, an old cathedral, an architectural-visualization office, a Japanese school, and a desert town. Together, they span a range of layouts from relatively orderly to extremely cluttered. Each scene is rendered along a smooth camera trajectory at a resolution of 2560×1440 , producing 265–1758 views per scene and 5,964 views in total. The scenes contain between 93 and 701 objects each, with 2,299 objects overall. Table 1 summarizes the statistics of each scene, and Figure 3 shows representative examples. For every frame, the dataset provides an RGB image, camera pose, per-object instance masks, per-object 3D bounding boxes, and a metric depth map. Although depth is not used by our method, we release the full set of annotations to support related tasks.

Metrics. For each object, we decode the prediction into a mesh and compare it with the corresponding ground-truth geometry. We report Chamfer distance using both ℓ_2 and ℓ_1 distances ($CD-\ell_2$ and $CD-\ell_1$), Earth Mover’s Distance (EMD), and F-Score.

For single-object evaluation in canonical space, both the prediction and ground truth are normalized with respect to the ground-truth unit sphere. Since different baselines adopt different canonical coordinate conventions, we perform per-instance ICP alignment before computing the metrics, making the comparison invariant to differences in canonical frame.




Figure 4 Qualitative comparison on Toys4k-Scene. Given the input observations, each method generates per-object meshes, whose surface normals are rendered in the shared world frame. Our method produces more detailed and accurately registered object geometry, closely matching the ground truth in the rightmost column. For each example, we enlarge the boxed region and visualize its pixelwise discrepancy with the ground truth in the row below. On overlapping surfaces, color denotes the angular error between surface normals, ranging from blue at 0° to red at 90° . Non-overlapping regions are shown separately: *pink* indicates spurious predicted surfaces, while *gray* indicates missing ground-truth surfaces. These regions also contribute to the mask-IoU error.

For compositional scene evaluation, each generated object is transformed from its canonical frame into the world frame using the canonical-to-world transformation. We then compare it directly with the corresponding ground-truth mesh in the shared world frame, without per-object ICP alignment. Distance-based metrics are normalized by the diagonal length of the ground-truth bounding box for each object. This protocol evaluates both object geometry and its placement within the composed scene.




Figure 5 Qualitative comparison on HouseCat6D. Our method recovers detailed object geometry that is accurately placed in the world frame and closely matches the scanned ground truth shown in the rightmost column, while the baselines often exhibit incomplete or distorted shapes and inaccurate placement. As in Fig. 4, each example includes an enlarged region of interest and a pixelwise discrepancy map with respect to the ground truth in the row below. On overlapping surfaces, color indicates the angular error between surface normals, ranging from blue at 0° to red at 90° . Non-overlapping regions are shown separately: *pink* denotes spurious predicted surfaces, while *gray* denotes missing ground-truth surfaces. Both contribute to the mask IoU error.

4.2 Canonical-Space Single-Object Generation

Setup and baselines. We evaluate on a fixed set of 500 randomly selected objects from Toys4k [86]. We compare against the official Pixal3D [42] pipeline, whose released model supports only single-view input, TRELIS [108] in both its single-image and multi-image modes, and TRELIS.2 [107], whose released model also supports only a single input view. Our model is evaluated with 1–16 views. All predictions are decoded and aligned using the same evaluation protocol. Table 2 reports generation quality across different numbers of input views (1, 2, 4, 8, 16) and per-view occlusion fractions (0, 25, 50, 75%).

Results. On clean inputs, our model is competitive with the official single-view Pixal3D pipeline when using one view and outperforms the baselines as additional views become available. Performance improves consistently with the number of input views. Since our model is trained with a variable number of views rather than optimized specifically for single-image generation, matching a dedicated single-view prior in the one-view setting indicates that the multi-view adaptation largely preserves the original generative capability while enabling substantial gains from additional observations.

The advantage becomes more pronounced under occlusion. As the number of views increases, our model remains robust even when individual observations are heavily incomplete. With 16 views, performance changes only slightly up to 50% per-view occlusion, whereas both the single-view baselines and the single-view variant of our model degrade substantially. Across all metrics, additional views considerably reduce the effect of occlusion, with distance errors remaining stable and F-Scores decreasing only mildly in the high-view regime. These results show that multi-view conditioning, together with the conditioning-view augmentation, provides strong robustness to partial observations.

4.3 Multi-object Scene Generation

Setup and baselines. We evaluate compositional scene generation on both synthetic and real-world multi-object scenes. For Toys4k-Scene, we construct cluttered scenes from Toys4k [86] objects and render each scene along a short orbiting camera trajectory. We additionally evaluate on HouseCat6D [34], which provides real-world captures together with per-object ground-truth meshes. For both datasets, each predicted object is placed into the shared world frame and evaluated directly against its corresponding ground-truth geometry.




Figure 6 Qualitative comparison on UE-MeshyScene. We compare per-object meshes generated in densely cluttered and heavily occluded scenes against the multi-view baseline ShapeR [83]. All predicted meshes are placed in the shared world frame and visualized using surface normals. Our method recovers more detailed and accurately placed geometry, including for small and severely occluded objects, while the baseline often produces over-smooth or misplaced shapes in highly cluttered regions. These qualitative differences are consistent with the quantitative results reported in Table 4.

We compare against a representative set of compositional and generative scene reconstruction methods, including SAM3D [10], SceneGen [57], SceneMaker [82], RecGen [119], and ShapeR [83]. Methods that require depth additionally receive ground-truth depth maps. We use the same evaluation protocol for UE-MeshyScene.

Qualitative comparison. Figures 4 and 5 compare the methods qualitatively on Toys4k-Scene and HouseCat6D, respectively. Alongside renderings of the composed scenes, we visualize per-pixel geometric discrepancies within selected zoomed-in regions. For pixels where the predicted and ground-truth surfaces overlap, color indicates the angular error between their surface normals, ranging from blue at 0° to red at 90° . Non-overlapping regions are visualized separately: pink denotes spurious predicted surfaces, while gray denotes ground-truth surfaces that are missing from the prediction. These regions also contribute to the mask-IoU error.

4.4 UE-MeshyScene Generation

Setup. We evaluate our method on the six UE-MeshyScene environments introduced above. This benchmark provides a substantially more challenging setting than previous dataset, with hundreds of objects, severe mutual occlusion, and a broader range of viewpoints. In such densely cluttered scenes, a single image often reveals only a small fraction of many objects. We therefore focus our comparison on ShapeR [83], a multi-view baseline that can incorporate observations across multiple viewpoints. Both methods are evaluated using the same world-frame protocol described above, without per-object alignment.

Results. Table 4 reports the quantitative results on UE-MeshyScene. Our method outperforms ShapeR

Table 5 Fusion ablation on the Toys4k [86] canonical-space setting of Table 2. We compare the learned IBR aggregator with arithmetic mean fusion (Avg) of the per-view feature grids under varying numbers of input views (#V) and per-view occlusion levels. We report $CD-\ell_2$ ($\times 10^{-3}$, $\downarrow$), $CD-\ell_1$ ($\times 10^{-2}$, $\downarrow$), EMD ($\times 10^{-2}$, $\downarrow$), and F-Score@{.02, .05} ($\uparrow$). The better fusion strategy for each setting and metric is shown in **bold**.

#V	Fusion	Occl. 0%					Occl. 25%					Occl. 50%					Occl. 75%				
		$CD-\ell_2$	$CD-\ell_1$	EMD	F@.02	F@.05	$CD-\ell_2$	$CD-\ell_1$	EMD	F@.02	F@.05	$CD-\ell_2$	$CD-\ell_1$	EMD	F@.02	F@.05	$CD-\ell_2$	$CD-\ell_1$	EMD	F@.02	F@.05
1	Avg	2.83	3.96	7.7	0.731	0.927	4.07	4.83	8.7	0.666	0.893	9.53	7.03	11.2	0.557	0.812	25.43	12.35	16.9	0.372	0.633
	IBR	2.59	3.85	7.8	0.739	0.932	3.70	4.73	8.7	0.669	0.896	8.90	6.84	11.2	0.554	0.818	25.45	12.30	16.8	0.370	0.631
2	Avg	1.58	3.33	7.2	0.778	0.951	2.65	3.95	7.9	0.723	0.930	4.30	5.05	8.9	0.642	0.884	12.91	8.72	13.2	0.471	0.742
	IBR	1.54	3.23	7.2	0.790	0.954	2.57	3.87	7.9	0.734	0.931	4.01	4.85	8.8	0.651	0.893	10.03	7.96	12.1	0.494	0.762
4	Avg	1.43	3.17	7.1	0.796	0.958	2.00	3.52	7.3	0.760	0.947	2.71	4.13	7.9	0.703	0.923	6.10	6.04	10.2	0.579	0.845
	IBR	1.35	3.02	7.0	0.813	0.964	1.62	3.35	7.3	0.778	0.950	2.56	4.01	7.8	0.719	0.925	5.67	5.86	10.1	0.590	0.849
8	Avg	1.55	3.13	7.1	0.806	0.960	1.51	3.26	7.2	0.783	0.957	1.89	3.63	7.5	0.744	0.942	3.04	4.65	8.6	0.650	0.898
	IBR	1.23	2.94	6.9	0.822	0.965	1.50	3.20	7.2	0.796	0.956	1.91	3.57	7.5	0.755	0.943	3.52	4.65	8.6	0.663	0.902
16	Avg	1.30	3.02	6.9	0.812	0.964	1.37	3.11	7.1	0.800	0.962	1.57	3.33	7.2	0.775	0.954	2.43	4.13	8.1	0.695	0.922
	IBR	1.19	2.89	6.9	0.827	0.967	1.29	3.01	7.0	0.813	0.964	1.44	3.23	7.2	0.788	0.955	2.25	3.97	8.0	0.713	0.929

Table 6 Scene-level fusion ablation. We compare the learned IBR aggregator with arithmetic mean fusion (Avg) on the compositional benchmarks in Table 3 and UE-MeshyScene in Table 4. All methods are evaluated using the same world-frame, no-ICP protocol, with F-Score computed at $\tau=0.05$. We report $CD-\ell_2$ ($\times 10^{-3}$, $\downarrow$), $CD-\ell_1$ ($\times 10^{-2}$, $\downarrow$), and EMD ($\times 10^{-2}$, $\downarrow$). The best result in each column is shown in **bold**.

Fusion	HouseCat6D [34]				Toys4k-Scene [86]				UE-MeshyScene			
	$CD-\ell_2\downarrow$	$CD-\ell_1\downarrow$	EMD $\downarrow$	F-Score $\uparrow$	$CD-\ell_2\downarrow$	$CD-\ell_1\downarrow$	EMD $\downarrow$	F-Score $\uparrow$	$CD-\ell_2\downarrow$	$CD-\ell_1\downarrow$	EMD $\downarrow$	F-Score $\uparrow$
Avg	0.27	1.60	3.3	0.996	0.65	1.88	3.8	0.981	2.83	2.92	5.1	0.944
IBR	0.28	1.53	3.4	0.995	0.61	1.78	3.7	0.981	2.48	2.66	5.0	0.951

across all reported metrics. The improvement is also more pronounced in the median than in the mean, indicating that the performance gap is not driven by a small number of difficult objects but is consistently observed across typical objects in the scenes. Figure 6 shows qualitative comparisons. Our method recovers more detailed object geometry while maintaining more accurate placement in the shared world frame, particularly for small and heavily occluded objects.

4.5 Ablation: Multi-View Fusion

Setup. We ablate the multi-view fusion module by replacing the learned aggregator in Eq. (14) with a simple arithmetic mean over the per-view 3D feature grids (“Avg”). Both variants start from the same fused representation: since the final layer of the IBR feature MLP is zero-initialized, its residual branch initially contributes nothing, and the learned aggregator reduces exactly to the cross-view mean. During training, the learned variant can move beyond this initialization by assigning voxel-wise, view-dependent residual contributions. We follow the same canonical-space evaluation protocol on Toys4k as in Sec. 4.2 and Table 2, using the same 500 objects, input-view counts, and per-view occlusion levels. The fusion module is the only component changed between the two variants. Results are reported in Table 5.

Results. The learned IBR aggregator performs comparably to or better than mean averaging in most settings, with its advantage becoming more evident as the number of input views and the occlusion level increase. With a single input view, the two variants perform similarly, since there are no additional observations to reweight. Once multiple views are available, IBR shows a consistent benefit. For example, with 2 views and 75% occlusion, IBR reduces $CD-\ell_2$ from 12.91 to 10.03 and improves F@.05 from 0.742 to 0.762. With 16 views, it outperforms mean fusion on nearly every metric across all occlusion levels. This trend suggests that learned view weighting is most useful when multiple observations provide features of varying reliability. The few settings in which mean averaging performs slightly better, primarily at 8 views, show only small differences. We therefore use the learned IBR aggregator as the default multi-view fusion module.

Scene-level ablation. We further compare the two fusion modules end-to-end in the compositional setting using the same world-frame, no-ICP protocol as Table 3. The three benchmarks form a natural progression




Figure 7 Application to a Marble-generated 3DGS world. Given a generated 3DGS scene, we render a sequence of posed observations and apply WorldSculpt to convert the scene into a compositional representation of individual per-object meshes in a shared world frame.

in scene complexity, from the relatively sparse and mildly occluded HouseCat6D scenes, to the denser Toys4k-Scene layouts, and finally to the heavily cluttered UE-MeshyScene benchmark with hundreds of mutually occluding objects. As shown in Table 6, the benefit of learned IBR aggregation becomes increasingly pronounced along this progression. On HouseCat6D, the two fusion strategies perform similarly, suggesting that learned view weighting offers limited benefit when objects are well observed. IBR consistently improves the distance metrics on the more cluttered Toys4k-Scene and achieves its largest gains on UE-MeshyScene, where it improves all reported metrics over mean fusion. For example, on UE-MeshyScene, IBR reduces $CD-\ell_2$ by 12% and increases F-Score from 0.944 to 0.951. This trend is consistent with the canonical-space ablation in Table 5: learned aggregation becomes most useful when multiple views provide observations of varying reliability, as commonly occurs in densely cluttered and heavily occluded scenes.

4.6 Application: Converting Generated 3D Worlds into Compositional Meshes

Generative world models such as Marble [104] and HY-World 2.0 [90] can synthesize photorealistic 3D environments represented as 3D Gaussian Splatting (3DGS). While such representations support high-quality free-viewpoint rendering, they do not explicitly separate the scene into editable object-level meshes required by downstream authoring, simulation, and asset reuse. As a proof of concept, we apply WorldSculpt to a world generated by Marble and convert its 3DGS representation into a compositional mesh scene without any additional training.

(i) Render posed observations from the 3DGS. We obtain the reconstructed 3DGS of a Marble scene and render a sequence of RGB frames along a virtual camera trajectory. For each frame, we additionally render a pixel-aligned metric depth map from the same Gaussian representation using the expected-depth renderer in gsplat. Since both RGB and depth are rendered with a controlled virtual camera in the coordinate system of the 3DGS, the camera intrinsics and poses are known exactly, and the depth maps and gravity direction are expressed in the same world frame.

(ii) Recover object masks and coarse 3D localizations. We apply SAM3 [7] to the rendered sequence to obtain per-object mask tracks. A physical object may occasionally be split into multiple tracklets because of duplicate detections or temporary tracking failures. We merge such fragments using both temporal mask overlap and 3D consistency. Tracklets that coexist with high mask IoU are first merged in 2D, after which same-category tracklets with strongly overlapping world-space boxes are merged when they occur in disjoint temporal intervals. Object boxes are re-estimated after each merge until the tracks stabilize.

For each consolidated track, we estimate a coarse 3D localization from its masks and rendered depth. Valid masked pixels are back-projected through the known camera intrinsics and poses to form partial world-space point clouds. We discard frames with outlying point-cloud centroids, merge the remaining observations, and retain the dominant cluster using DBSCAN [21]. An axis-aligned bounding box of the resulting point cloud is then used as the coarse localization required by WorldSculpt.

(iii) Apply WorldSculpt. The resulting posed RGB frames, per-object mask tracks, and coarse 3D localization boxes match the input assumptions of Sec. 3.4. We therefore apply the standard WorldSculpt inference pipeline without modification to generate a collection of individual per-object meshes in the shared world frame, as shown in Figure 7. This example demonstrates that a generated 3DGS world can be converted into a compositional mesh representation using the same object-generation model, without scene- or domain-specific retraining.

4.7 Limitations

WorldSculpt has several limitations. First, our method relies on upstream estimates of camera poses, per-object instance masks, and coarse 3D localization boxes. The conditioning-view augmentation curriculum in Sec. 3.3 improves robustness to moderate errors in these inputs, but large inaccuracies, such as severely misplaced boxes or erroneous masks, can still affect the generated geometry. A more integrated system that jointly performs object detection, localization, and generation could reduce this dependence on upstream predictions. Performance may also deteriorate under challenging capture conditions that are not explicitly modeled during training, including low-light degradation [64, 65, 68, 63] and motion-induced artifacts [32, 31, 36, 66, 61].

Second, the current model generates geometry only. We build on the first two geometry stages of Pixa3D and do not model object appearance. Both Pixa3D [42] and TRELIS.2 [107] support physically based rendering materials, making it possible to extend our multi-view conditioning pathway to their texture and material generation stages. This would allow compositional scenes to be generated with both detailed geometry and consistent object appearance.

Third, our formulation assumes a static scene. Multi-view observations of each object are treated as measurements of the same fixed geometry and pose, so moving or deforming objects are not currently supported. Extending WorldSculpt to dynamic scenes would require estimating object motion and accounting for non-rigid changes across observations. Possible directions include incorporating optical-flow-based motion estimation [88, 91, 62, 30] and introducing articulated representations such as skeletal models [115] or SMPL-X-based models [71, 54, 67, 120, 60] for objects undergoing articulated deformation.

5 Conclusion

We presented *WorldSculpt*, a framework for generating compositional 3D scenes from grounded multi-view observations. WorldSculpt adapts a strong object-level generative prior with a multi-view conditioning pathway that integrates observations of each object in an anchor-aligned canonical frame. The generative prior provides plausible completion of occluded and unobserved geometry, while multi-view conditioning grounds the generated shape in the available image evidence. Although the model is trained entirely on individual objects in canonical space, it generalizes without scene-level retraining to large, densely cluttered scenes at inference time.

To support evaluation in this setting, we introduced UE-MeshyScene, a photorealistic synthetic benchmark containing six scenes with 93 to 701 objects each. Across controlled single-object evaluations, multi-object scene benchmarks, and large-scale UE-MeshyScene experiments, WorldSculpt consistently outperforms the evaluated baselines, with the largest gains appearing under dense clutter and severe occlusion. Our fusion ablations show a similar trend, with learned multi-view aggregation becoming increasingly beneficial as observations grow more incomplete and heterogeneous.

Future work could extend the same conditioning mechanism to texture and material generation for joint geometry and appearance synthesis. Another important direction is to relax the static-scene assumption and support moving or deforming objects. Overall, our results show that adapting object-level generative

priors to multi-view compositional generation provides a practical path toward editable 3D representations of complex, cluttered environments.

References

- [1] W. Agnew, C. Xie, A. Walsman, O. Murad, Y. Wang, P. Domingos, and S. Srinivasa. Amodal 3d reconstruction for robotic manipulation via stability and connectivity. In *Conference on robot learning*, pages 1498–1508. PMLR, 2021.
- [2] A. Ahmadyan, L. Zhang, A. Ablavatski, J. Wei, and M. Grundmann. Objectron: A large scale dataset of object-centric videos in the wild with pose annotations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7822–7831, 2021.
- [3] A. Avetisyan, M. Dahnert, A. Dai, M. Savva, A. X. Chang, and M. Nießner. Scan2cad: Learning cad model alignment in rgb-d scans. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 2614–2623, 2019.
- [4] G. Baruch, Z. Chen, A. Dehghan, T. Dimry, Y. Feigin, P. Fu, T. Gebauer, B. Joffe, D. Kurz, A. Schwartz, et al. Arkitscenes: A diverse real-world dataset for 3d indoor scene understanding using mobile rgb-d data. *arXiv preprint arXiv:2111.08897*, 2021.
- [5] M. Boss, Z. Huang, A. Vasishta, and V. Jampani. Sf3d: Stable fast 3d mesh reconstruction with uv-unwrapping and illumination disentanglement. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 16240–16250, 2025.
- [6] G. Brazil, A. Kumar, J. Straub, N. Ravi, J. Johnson, and G. Gkioxari. Omni3d: A large benchmark and model for 3d object detection in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13154–13164, 2023.
- [7] N. Carion, L. Gustafson, Y.-T. Hu, S. Debnath, R. Hu, D. Suris, C. Ryali, K. V. Alwala, H. Khedr, A. Huang, et al. Sam 3: Segment anything with concepts. *arXiv preprint arXiv:2511.16719*, 2025.
- [8] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017.
- [9] S. Chen, H. Guo, S. Zhu, F. Zhang, Z. Huang, J. Feng, and B. Kang. Video depth anything: Consistent depth estimation for super-long videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22831–22840, 2025.
- [10] X. Chen, F.-J. Chu, P. Gleize, K. J. Liang, A. Sax, H. Tang, W. Wang, M. Guo, T. Hardin, X. Li, et al. Sam 3d: 3dfy anything in images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7220–7232, 2026.
- [11] H. K. Cheng, S. W. Oh, B. Price, J.-Y. Lee, and A. Schwing. Putting the object back into video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3151–3161, 2024.
- [12] H. K. Cheng, S. W. Oh, B. Price, A. Schwing, and J.-Y. Lee. Tracking anything with decoupled video segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1316–1326, 2023.
- [13] H. K. Cheng and A. G. Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European conference on computer vision*, pages 640–658. Springer, 2022.
- [14] T. Cheng, L. Song, Y. Ge, W. Liu, X. Wang, and Y. Shan. Yolo-world: Real-time open-vocabulary object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16901–16911, 2024.

- [15] J. Collins, S. Goel, K. Deng, A. Luthra, L. Xu, E. Gundogdu, X. Zhang, T. F. Y. Vicente, T. Dideriksen, H. Arora, et al. Abo: Dataset and benchmarks for real-world 3d object understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 21126–21136, 2022.
- [16] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017.
- [17] M. Deitke, R. Liu, M. Wallingford, H. Ngo, O. Michel, A. Kusupati, A. Fan, C. Laforte, V. Voleti, S. Y. Gadre, et al. Objaverse-xl: A universe of 10m+ 3d objects. *Advances in Neural Information Processing Systems*, 36:35799–35813, 2023.
- [18] M. Deitke, D. Schwenk, J. Salvador, L. Weihs, O. Michel, E. VanderBilt, L. Schmidt, K. Ehsani, A. Kembhavi, and A. Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13142–13153, 2023.
- [19] D. DeTone, T. Shen, F. Zhang, L. Ma, J. Straub, R. Newcombe, and J. Engel. Boxer: Robust lifting of open-world 2d bounding boxes to 3d. *arXiv preprint arXiv:2604.05212*, 2026.
- [20] L. Downs, A. Francis, N. Koenig, B. Kinman, R. Hickman, K. Reymann, T. B. McHugh, and V. Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2553–2560. Ieee, 2022.
- [21] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, pages 226–231, 1996.
- [22] H. Fu, B. Cai, L. Gao, L.-X. Zhang, J. Wang, C. Li, Q. Zeng, C. Sun, R. Jia, B. Zhao, et al. 3d-front: 3d furnished rooms with layouts and semantics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10933–10942, 2021.
- [23] H. Fu, R. Jia, L. Gao, M. Gong, B. Zhao, S. Maybank, and D. Tao. 3d-future: 3d furniture shape with texture. *International Journal of Computer Vision*, 129(12):3313–3337, 2021.
- [24] C. Gümeli, A. Dai, and M. Nießner. Roca: Robust cad model retrieval and alignment from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4022–4031, 2022.
- [25] Y. Hong, K. Zhang, J. Gu, S. Bi, Y. Zhou, D. Liu, F. Liu, K. Sunkavalli, T. Bui, and H. Tan. Lrm: Large reconstruction model for single image to 3d. In *International Conference on Learning Representations*, volume 2024, pages 50678–50702, 2024.
- [26] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al. Lora: Low-rank adaptation of large language models. *Iclr*, 1(2):3, 2022.
- [27] W. Hu, X. Gao, X. Li, S. Zhao, X. Cun, Y. Zhang, L. Quan, and Y. Shan. Depthcrafter: Generating consistent long depth sequences for open-world videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2005–2015, 2025.
- [28] J. Huang, Q. Zhou, H. Rabeti, A. Korovko, H. Ling, X. Ren, T. Shen, J. Gao, D. Slepichev, C.-H. Lin, et al. Vipe: Video pose engine for 3d geometric perception. *arXiv preprint arXiv:2508.10934*, 2025.
- [29] Z. Huang, Y.-C. Guo, X. An, Y. Yang, Y. Li, Z.-X. Zou, D. Liang, X. Liu, Y.-P. Cao, and L. Sheng. Midi: Multi-instance diffusion for single image to 3d scene generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23646–23657, 2025.
- [30] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017.
- [31] X. Ji, G. Lin, Z. Yin, J. Zhao, and Y. Zheng. Moment-reenacting: Inverse motion degradation with cross-shutter guidance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2026.

- [32] X. Ji, Z. Wang, Z. Zhong, and Y. Zheng. Rethinking video frame interpolation from shutter mode induced degradation. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12225–12234. IEEE, 2023.
- [33] H. Jun and A. Nichol. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463*, 2023.
- [34] H. Jung, S.-C. Wu, P. Ruhkamp, G. Zhai, H. Schieber, G. Rizzoli, P. Wang, H. Zhao, L. Garattoni, S. Meier, et al. Housecat6d-a large-scale multi-modal category level 6d object perception dataset with household objects in realistic scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22498–22508, 2024.
- [35] N. Keetha, N. Müller, J. Schönberger, L. Porzi, Y. Zhang, T. Fischer, A. Knapitsch, D. Zauss, E. Weber, N. Antunes, et al. Mapanything: Universal feed-forward metric 3d reconstruction; map-anything. github. io. In *2026 International Conference on 3D Vision (3DV)*, pages 499–509. IEEE, 2026.
- [36] O. Kupyn, T. Martyniuk, J. Wu, and Z. Wang. Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8877–8886. IEEE, 2019.
- [37] K. N. Kutulakos and S. M. Seitz. A theory of shape by space carving. *International journal of computer vision*, 38(3):199–218, 2000.
- [38] A. Laurentini. The visual hull concept for silhouette-based image understanding. *IEEE Transactions on pattern analysis and machine intelligence*, 16(2):150–162, 1994.
- [39] J. Lazarow, D. Griffiths, G. Kohavi, F. Crespo, and A. Dehghan. Cubify anything: Scaling indoor 3d object detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22225–22233, 2025.
- [40] A. Lemesko, B. Gabdullin, N. Drozdov, A. Konushin, D. Rukhovich, and M. Kolodiazhnyi. Zoo3d: Zero-shot 3d object detection at scene level. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25820–25829, 2026.
- [41] V. Leroy, Y. Cabon, and J. Revaud. Grounding image matching in 3d with mast3r. In *European conference on computer vision*, pages 71–91. Springer, 2024.
- [42] D.-Y. Li, W. Zhao, Y. Chen, W. Hu, M.-H. Guo, F.-L. Zhang, Y. Shan, and S.-M. Hu. Pixel3d: Pixel-aligned 3d generation from images. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pages 1–12, 2026.
- [43] K. Li, D. DeTone, Y. F. S. Chen, M. Vo, I. Reid, H. Rezatofighi, C. Sweeney, J. Straub, and R. Newcombe. Odam: Object detection, association, and mapping using posed rgb video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5998–6008, 2021.
- [44] Z. Li, R. Tucker, F. Cole, Q. Wang, L. Jin, V. Ye, A. Kanazawa, A. Holynski, and N. Snavely. Megasam: Accurate, fast and robust structure and motion from casual dynamic videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10486–10496, 2025.
- [45] H. Lin, S. Chen, J. Liew, D. Y. Chen, Z. Li, G. Shi, J. Feng, and B. Kang. Depth anything 3: Recovering the visual space from any views. *arXiv preprint arXiv:2511.10647*, 2025.
- [46] Y. Lin, C. Lin, P. Pan, H. Yan, F. Yiqiang, Y. Mu, and K. Fragkiadaki. Partcrafter: Structured 3d mesh generation via compositional latent diffusion transformers. *Advances in neural information processing systems*, 38:35387–35415, 2026.
- [47] Y. Lipman, R. T. Chen, H. Ben-Hamu, M. Nickel, and M. Le. Flow matching for generative modeling. In *The eleventh international conference on learning representations*, 2022.
- [48] H. Liu, Y. Zheng, G. Chen, S. Cui, and X. Han. Towards high-fidelity single-view holistic reconstruction of indoor scenes. In *European Conference on Computer Vision*, pages 429–446. Springer, 2022.

- [49] M. Liu, C. Xu, H. Jin, L. Chen, M. Varma T, Z. Xu, and H. Su. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *Advances in Neural Information Processing Systems*, 36:22226–22246, 2023.
- [50] R. Liu, R. Wu, B. Van Hoorick, P. Tokmakov, S. Zakharov, and C. Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9298–9309, 2023.
- [51] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, Q. Jiang, C. Li, J. Yang, H. Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European conference on computer vision*, pages 38–55. Springer, 2024.
- [52] Y. Liu, C. Lin, Z. Zeng, X. Long, L. Liu, T. Komura, and W. Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. In *International conference on learning representations*, volume 2024, pages 27676–27697, 2024.
- [53] X. Long, Y.-C. Guo, C. Lin, Y. Liu, Z. Dou, L. Liu, Y. Ma, S.-H. Zhang, M. Habermann, C. Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9970–9980, 2024.
- [54] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. Smpl: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 851–866. 2023.
- [55] K.-K. Maninis, S. Popov, M. Nießner, and V. Ferrari. Cad-estate: Large-scale cad model annotation in rgb videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20189–20199, 2023.
- [56] Y. Mao, Y. Zhang, H. Jiang, A. Chang, and M. Savva. Multiscan: Scalable rgb-d scanning for 3d environments with articulated objects. *Advances in neural information processing systems*, 35:9058–9071, 2022.
- [57] Y. Meng, H. Wu, Y. Zhang, and W. Xie. Scenegen: Single-image 3d scene generation in one feedforward pass. In *2026 International Conference on 3D Vision (3DV)*, pages 543–553. IEEE, 2026.
- [58] R. Murai, E. Dexheimer, and A. J. Davison. Mast3r-slam: Real-time dense slam with 3d reconstruction priors. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 16695–16705, 2025.
- [59] A. Nichol, H. Jun, P. Dhariwal, P. Mishkin, and M. Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022.
- [60] M. Niu, M. Cao, Y. Zhan, Q. Zhu, M. Ma, J. Zhao, Y. Zeng, Z. Zhong, X. Sun, and Y. Zheng. Anicrafter: Customizing realistic human-centric animation via avatar-background conditioning in video diffusion models. *arXiv preprint arXiv:2505.20255*, 2025.
- [61] M. Niu, T. Chen, Y. Zhan, Z. Li, X. Ji, and Y. Zheng. Rs-nerf: Neural radiance fields from rolling shutter images. In *European Conference on Computer Vision*, pages 163–180. Springer, 2024.
- [62] M. Niu, X. Cun, X. Wang, Y. Zhang, Y. Shan, and Y. Zheng. Mofa-video: Controllable image animation via generative motion field adaptations in frozen image-to-video diffusion model. In *European conference on computer vision*, pages 111–128. Springer, 2024.
- [63] M. Niu, Z. Li, Y. Zhan, H. H. Nguyen, I. Echizen, and Y. Zheng. Physics-based adversarial attack on near-infrared human detector for nighttime surveillance camera systems. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 8799–8807, 2023.
- [64] M. Niu, Z. Li, Z. Zhong, and Y. Zheng. Visibility constrained wide-band illumination spectrum design for seeing-in-the-dark. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13976–13985. IEEE, 2023.
- [65] M. Niu, M. Ma, Y. Zhan, Q. Zhu, Z. Zhong, W. Guo, C. W. Chen, and Y. Zheng. Toward robust and 3d-aware rgb-nir imaging in the dark, 2026.

- [66] M. Niu, Y. Zhan, Q. Zhu, Z. Li, W. Wang, Z. Zhong, X. Sun, and Y. Zheng. Bundle adjusted gaussian avatars deblurring. *arXiv e-prints*, pages arXiv–2411, 2024.
- [67] M. Niu, Y. Zhan, Q. Zhu, Z. Li, W. Wang, Z. Zhong, X. Sun, and Y. Zheng. Motion-aware animatable gaussian avatars deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 40140–40151, 2026.
- [68] M. Niu, Z. Zhong, and Y. Zheng. Nir-assisted video enhancement via unpaired 24-hour data. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10744–10754. IEEE, 2023.
- [69] E. Ozguroglu, R. Liu, D. Surís, D. Chen, A. Dave, P. Tokmakov, and C. Vondrick. pix2gestalt: Amodal segmentation by synthesizing wholes. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3931–3940. IEEE Computer Society, 2024.
- [70] L. Pan, D. Barath, M. Pollefeys, and J. L. Schönberger. Global Structure-from-Motion Revisited. In *European Conference on Computer Vision (ECCV)*, 2024.
- [71] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. Osman, D. Tzionas, and M. J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10975–10985, 2019.
- [72] L. Piccinelli, Y.-H. Yang, C. Sakaridis, M. Segu, S. Li, L. Van Gool, and F. Yu. Unidepth: Universal monocular metric depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10106–10116, 2024.
- [73] R. Ranftl, K. Lasinger, D. Hafner, K. Schindler, and V. Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 44(3):1623–1637, 2020.
- [74] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryali, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson, et al. Sam 2: Segment anything in images and videos. In *International Conference on Learning Representations*, volume 2025, pages 28085–28128, 2025.
- [75] J. Reizenstein, R. Shapovalov, P. Henzler, L. Sbordone, P. Labatut, and D. Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10901–10911, 2021.
- [76] T. Ren, S. Liu, A. Zeng, J. Lin, K. Li, H. Cao, J. Chen, X. Huang, Y. Chen, F. Yan, et al. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*, 2024.
- [77] M. Roberts, J. Ramapuram, A. Ranjan, A. Kumar, M. A. Bautista, N. Paczan, R. Webb, and J. M. Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10912–10922, 2021.
- [78] D. Rukhovich, A. Vorontsova, and A. Konushin. Imvoxelnet: Image to voxels projection for monocular and multi-view general-purpose 3d object detection. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2397–2406, 2022.
- [79] K. Schmid, N. von Lützw, J. Hladký, A. Dai, and M. Nießner. Genrecon: Bridging generative priors for multi-view 3d scene reconstruction. *arXiv preprint arXiv:2605.23888*, 2026.
- [80] J. L. Schönberger and J.-M. Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [81] J. L. Schönberger, E. Zheng, M. Pollefeys, and J.-M. Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016.
- [82] Y. Shi, W. Li, Z. Wang, H. Li, X. Chen, P. Tan, and L. Zhang. Scenemaker: Open-set 3D scene generation with decoupled de-occlusion and pose estimation model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27146–27156, 2026.

- [83] Y. Siddiqui, D. Frost, S. Aroudj, A. Avetisyan, H. Howard-Jenkins, D. DeTone, P. Moulon, Q. Wu, Z. Li, J. Straub, R. Newcombe, and J. Engel. Shaper: Robust conditional 3d shape generation from casual captures, 2026.
- [84] O. Siméoni, H. V. Vo, M. Seitzer, F. Baldassarre, M. Oquab, C. Jose, V. Khalidov, M. Szafraniec, S. Yi, M. Ramamonjisoa, F. Massa, D. Haziza, L. Wehrstedt, J. Wang, T. Darcet, T. Moutakanni, L. Sentana, C. Roberts, A. Vedaldi, J. Tolan, J. Brandt, C. Couprie, J. Mairal, H. Jégou, P. Labatut, and P. Bojanowski. DINOv3, 2025.
- [85] S. Song, S. P. Lichtenberg, and J. Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 567–576, 2015.
- [86] S. Stojanov, A. Thai, and J. M. Rehg. Using shape to categorize: Low-shot learning with an explicit shape bias. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1798–1808, 2021.
- [87] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijmans, S. Green, J. J. Engel, R. Mur-Artal, C. Ren, S. Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019.
- [88] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018.
- [89] J. Tang, Z. Chen, X. Chen, T. Wang, G. Zeng, and Z. Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. In *European Conference on Computer Vision*, pages 1–18. Springer, 2024.
- [90] H. Team, C. Cao, X. Zuo, Z. Wang, Y. Zhang, J. Wu, Z. Liu, Y. Gong, Y. Liu, B. Yuan, et al. Hy-world 2.0: A multi-modal world model for reconstructing, generating, and simulating 3d worlds. *arXiv e-prints*, pages arXiv–2604, 2026.
- [91] Z. Teed and J. Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European conference on computer vision*, pages 402–419. Springer, 2020.
- [92] Z. Teed and J. Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in neural information processing systems*, 34:16558–16569, 2021.
- [93] Z. Teed, L. Lipson, and J. Deng. Deep patch visual odometry. *Advances in Neural Information Processing Systems*, 36:39033–39051, 2023.
- [94] D. Tochilkin, D. Pankratz, Z. Liu, Z. Huang, A. Letts, Y. Li, D. Liang, C. Laforte, V. Jampani, and Y.-P. Cao. Triposr: Fast 3d object reconstruction from a single image. *arXiv preprint arXiv:2403.02151*, 2024.
- [95] J. Wang, M. Chen, N. Karaev, A. Vedaldi, C. Rupprecht, and D. Novotny. Vgggt: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5294–5306, 2025.
- [96] Q. Wang, Z. Wang, K. Genova, P. P. Srinivasan, H. Zhou, J. T. Barron, R. Martin-Brualla, N. Snavely, and T. Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2021.
- [97] Q. Wang, Y. Zhang, A. Holynski, A. A. Efros, and A. Kanazawa. Continuous 3d perception model with persistent state. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 10510–10522, 2025.
- [98] R. Wang, S. Xu, C. Dai, J. Xiang, Y. Deng, X. Tong, and J. Yang. Moge: Unlocking accurate monocular geometry estimation for open-domain images with optimal training supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5261–5271, 2025.
- [99] S. Wang, V. Leroy, Y. Cabon, B. Chidlovskii, and J. Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20697–20709, 2024.

- [100] Y. Wang, J. Zhou, H. Zhu, W. Chang, Y. Zhou, Z. Li, J. Chen, J. Pang, C. Shen, and T. He. π^3 : Permutation-equivariant visual geometry learning. *arXiv preprint arXiv:2507.13347*, 2025.
- [101] Z. Wang, Y. Wang, Y. Chen, C. Xiang, S. Chen, D. Yu, C. Li, H. Su, and J. Zhu. Crm: Single image to 3d textured mesh with convolutional reconstruction model. In *European conference on computer vision*, pages 57–74. Springer, 2024.
- [102] X. Wei, K. Zhang, S. Bi, H. Tan, F. Luan, V. Deschaintre, K. Sunkavalli, H. Su, and Z. Xu. Meshlrm: Large reconstruction model for high-quality meshes. *arXiv preprint arXiv:2404.12385*, 2024.
- [103] B. Wen, J. Tremblay, V. Blukis, S. Tyree, T. Müller, A. Evans, D. Fox, J. Kautz, and S. Birchfield. Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 606–617, 2023.
- [104] World Labs. Marble: Create and share 3D worlds, 2025.
- [105] J. Wu, Y. Jiang, Q. Liu, Z. Yuan, X. Bai, and S. Bai. General object foundation model for images and videos at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3783–3795, 2024.
- [106] T. Wu, C. Zheng, F. Guan, A. Vedaldi, and T.-J. Cham. Amodal3r: Amodal 3d reconstruction from occluded 2d images. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9181–9193. IEEE, 2025.
- [107] J. Xiang, X. Chen, S. Xu, R. Wang, Z. Lv, Y. Deng, H. Zhu, Y. Dong, H. Zhao, N. J. Yuan, and J. Yang. Native and compact structured latents for 3d generation. *Tech report*, 2025.
- [108] J. Xiang, Z. Lv, S. Xu, Y. Deng, R. Wang, B. Zhang, D. Chen, X. Tong, and J. Yang. Structured 3d latents for scalable and versatile 3d generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 21469–21480, 2025.
- [109] J. Xu, W. Cheng, Y. Gao, X. Wang, S. Gao, and Y. Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*, 2024.
- [110] Y. Xu, Z. Shi, W. Yifan, H. Chen, C. Yang, S. Peng, Y. Shen, and G. Wetzstein. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. In *European Conference on Computer Vision*, pages 1–20. Springer, 2024.
- [111] J. Yang, A. Sax, K. J. Liang, M. Henaff, H. Tang, A. Cao, J. Chai, F. Meier, and M. Feiszli. Fast3r: Towards 3d reconstruction of 1000+ images in one forward pass. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 21924–21935, 2025.
- [112] L. Yang, B. Kang, Z. Huang, X. Xu, J. Feng, and H. Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10371–10381, 2024.
- [113] L. Yang, B. Kang, Z. Huang, Z. Zhao, X. Xu, J. Feng, and H. Zhao. Depth anything v2. *Advances in Neural Information Processing Systems*, 37:21875–21911, 2024.
- [114] S. Yang and S. Scherer. Cubeslam: Monocular 3-d object slam. *IEEE Transactions on Robotics*, 35(4):925–938, 2019.
- [115] Z. Yang, A. Zeng, C. Yuan, and Y. Li. Effective whole-body pose estimation with two-stages distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4210–4220, 2023.
- [116] C. Yeshwanth, Y.-C. Liu, M. Nießner, and A. Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–22, 2023.
- [117] W. Yin, C. Zhang, H. Chen, Z. Cai, G. Yu, K. Wang, X. Chen, and C. Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9043–9053, 2023.

- [118] X. Yu, M. Xu, Y. Zhang, H. Liu, C. Ye, Y. Wu, Z. Yan, C. Zhu, Z. Xiong, T. Liang, et al. Mvimgnet: A large-scale dataset of multi-view images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9150–9161, 2023.
- [119] A. Zadaianchuk, L. Barcellona, L. Schuenemann, C. Gumbsch, Z. Wang, M. Z. Irshad, F. Despinoy, R. Aljundi, S. Gavves, and S. Zakharov. Recgen: Reconstructive generation of 3d scenes from rgb-d observations, 2026.
- [120] Y. Zhan, Q. Zhu, M. Niu, M. Ma, J. Zhao, Z. Zhong, X. Sun, Y. Qiao, and Y. Zheng. Towards explicit exoskeleton for the reconstruction of complicated 3d human avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14259–14269, 2025.
- [121] H. Zhang, H. Jiang, Q. Yao, Y. Sun, R. Zhang, H. Zhao, H. Li, H. Zhu, and Z. Yang. Detect anything 3d in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5048–5059, 2025.
- [122] J. Zhang, C. Herrmann, J. Hur, V. Jampani, F. Cole, D. Sun, M.-H. Yang, et al. Monst3r: A simple approach for estimating geometry in the presence of motion. In *International Conference on Learning Representations*, volume 2025, pages 82863–82886, 2025.
- [123] Z. Zhao, Z. Lai, Q. Lin, Y. Zhao, H. Liu, S. Yang, Y. Feng, M. Yang, S. Zhang, X. Yang, et al. Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation. *arXiv preprint arXiv:2501.12202*, 2025.
- [124] Z. Zhao, W. Liu, X. Chen, X. Zeng, R. Wang, P. Cheng, B. Fu, T. Chen, G. Yu, and S. Gao. Michelangelo: Conditional 3d shape generation based on shape-image-text aligned latent representation. *Advances in neural information processing systems*, 36:73969–73982, 2023.
- [125] J. Zhou and Y.-W. Tai. Amodalgen3d: Generative amodal 3d object reconstruction from sparse unposed views. *arXiv e-prints*, pages arXiv–2511, 2025.