
VolFill: Single-View Amodal 3D Scene Reconstruction with Volumetric Flow Matching

Tuan Duc Ngo¹ Chuang Gan¹ Evangelos Kalogerakis^{1,2}
¹UMass Amherst ²TU Crete
github.com/VolFill

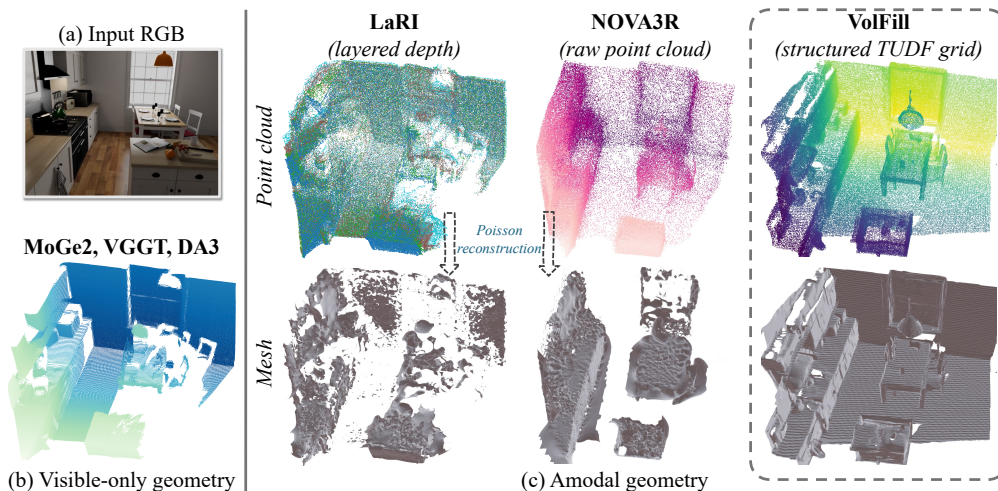


Figure 1: **VolFill** synthesizes *structured amodal 3D geometry* from (a) a single-view image, recovering holistic scene layouts from partial visibility. (b) Pixel-aligned methods are restricted to visible surfaces. (c) Amodal baselines produce sparse, noisy or artifact-heavy geometry, yielding fragmented meshes. Our approach delivers clean, sharp point clouds and smooth, consistent meshes.

Abstract

Reconstructing the complete geometry of a scene from a single RGB image remains challenging—especially when inferring hidden structures where visual evidence is incomplete. We introduce VolFill, a generative framework that predicts the 3D structure of the complete scene rather than relying on traditional pixel-aligned regression. Our method utilizes a hybrid 3D VAE to compress sparse *truncated unsigned distance function* grids into a compact latent space, paired with a latent Diffusion Transformer that denoises this representation to recover the complete scene. We condition the generation on geometry foundation models, leveraging rich spatial priors for robust reasoning. Unlike existing methods limited by per-ray constraints or unstructured point-cloud queries, VolFill provides a structured representation that supports direct surface extraction and occupancy queries at scale. Extensive experiments on the SCRREAM and NRGB-D datasets demonstrate that our approach significantly outperforms current baselines, providing a robust foundation for holistic spatial understanding.

1 Introduction

This paper addresses amodal 3D scene reconstruction, seeking to recover the complete geometry, including observed and occluded structures, from a single RGB image. Beyond the inherent ill-posed

problem, amodal reconstruction must infer hidden structures where geometric evidence is partially absent. This requires synthesizing under-determined regions that remain physically plausible and consistent with the observed scene. Such spatial awareness is critical for practical applications, as navigating or interacting with an environment requires a comprehensive understanding of the scene that extends well beyond the immediate line of sight.

Recent feed-forward methods [5, 37, 77, 78, 79, 80, 81, 87, 89] can recover accurate 3D structure from images in seconds, but all share a hard constraint: *pixel alignment*. Every predicted point lies on a source camera ray, bounding reconstruction strictly to visible surfaces and producing duplicated geometry in overlapping views. Optimization-based scene representations [38, 49, 52, 56, 69] reconstruct complete geometry but require dense captures and per-scene optimization which can take hours to compute. Large reconstruction models [30, 97] and generative 3D models [42, 73, 84, 85] demonstrate impressive 3D reconstruction quality from sparse inputs, but are predominantly optimized for object-centric scenarios, limiting their applicability to complex scenes with unbounded extents and challenging amodal structures.

At the scene level, LaRI [41] proposes a *layered ray intersection* representation, regressing multiple ordered surface intersections per camera ray. While this enables amodal recovery within a pixel-aligned framework, it remains a partial solution; the fixed layer count under-represents densely occluded structures while wasting capacity in open regions. Furthermore, it frequently yields layered artifacts (the “multiple wall” effect) rather than accurately recovering the actual unobserved surfaces. Alternatively, NOVA3R [11] employs flow matching to directly generate *amodal 3D point clouds*, but often yields noisy and disintegrated outputs (Fig. 1).

To address these limitations, we propose VolFill, a latent diffusion framework that predicts a complete Truncated Unsigned Distance Function (TUDF) voxel grid from a single image. A TUDF encodes the distance to the nearest scene surface, including occluded regions, as a continuous scalar field, enabling direct surface extraction via isosurfacing [23] and without requiring post-processing reconstruction of point clouds. Unlike point clouds, it scales well with scene complexity; unlike layered ray representations, it places no constraint on the number or layers of recoverable occluded surfaces. VolFill consists of a hybrid 3D VAE that encodes the sparse TUDF grid into a compact dense latent space, and a Diffusion Transformer (DiT) trained with flow matching. To ensure robust in-the-wild generalization, we introduce a dual conditioning strategy combining high-level image tokens with explicit visible geometry from a frozen MoGe2 [79] model, grounding amodal reasoning in the observed scene while leveraging strong geometric priors to compensate for limited 3D training data.

Trained on 3D-FRONT [17] and ScanNet++ [91], VolFill achieves state-of-the-art performance on the SCRREAM [33] and NRGB-D [2] benchmarks, synthesizing high-fidelity amodal geometry with significantly greater sharpness and structural accuracy than existing methods (Fig. 1). In summary, our main contributions are:

- **VolFill**, a generative framework that utilizes volumetric flow matching to recover complete scene-level geometry from single-view images.
- **A hybrid 3D VAE** enabling efficient spatial compression of high-resolution TUDF grids to a compact latent space, facilitating high-fidelity reconstruction of complex amodal structures.
- **A dual-conditioning strategy** leveraging geometry foundation models to integrate high-level image tokens with explicit visible geometry.

2 Related Work

Volumetric 3D Representations provide a flexible foundation for 3D vision, spanning applications from scene understanding to generative modeling. To mitigate the cubic complexity of dense grids, sparse convolutional engines [13, 22, 71] have become the standard for efficiently processing geometry along active surfaces. In scene-level reconstruction, “lifting” paradigms map 2D features into these volumes through depth-guided projection [70], ray-sampling [6], or transformer-based voxel queries [43, 82]. Within the generative domain, early models applied diffusion directly to raw voxels [53], whereas contemporary frameworks leverage hierarchical latent diffusion [12, 62] or rectified flow on sparse latents [85] to achieve high-fidelity structural synthesis. Building on these advances, our approach adopts a TUDF representation within a structured 3D grid, using sparse convolutions to capture complex scene geometry while maintaining high computational efficiency.

Pixel-Aligned Single-View 3D Reconstruction recovers *visible* surface geometry, evolving from early handcrafted features [29, 34, 66, 67] to deep architectures [3, 15, 18, 39, 75]. Progress in scaling has yielded models with remarkable zero-shot generalization [4, 5, 24, 31, 59, 61, 89, 90, 92]. Recently, methods have further refined geometric accuracy by jointly predicting intrinsics [5, 59], regressing dense 3D pointmaps [78, 79], or utilizing diffusion priors [60, 65] for high-fidelity, arbitrary-resolution depth synthesis [19, 21, 26, 36, 58, 87, 93]. Nonetheless, these approaches remain fundamentally confined to 2.5D visible surfaces. Their inability to infer occluded regions or amodal structure directly motivates our shift from surface-level estimation to full-scene reconstruction.

Geometry Foundation Models as Priors. Recent geometry foundation models (GFM) [16, 44, 77, 78, 79, 80] have demonstrated strong performance in 3D reconstruction and serve as robust feature extractors for downstream tasks requiring explicit spatial priors [8, 54, 83, 87]. Current approaches typically adopt one of two strategies: (1) using explicit geometric outputs (e.g., depth or point clouds) as structural anchors to guide synthesis [7, 10, 94, 95], or (2) injecting latent foundation features into transformer blocks to maintain geometric consistency [32, 83, 86]. Our approach unifies these paradigms via a dual-conditioning strategy. By leveraging both explicit visible geometry for structural grounding and foundation latent tokens for high-level context, we enable robust and physically plausible amodal reconstruction that generalizes to diverse scenarios.

Amodal 3D Reconstruction aims to recover complete geometry from partial visual observations. Object-centric strategies often utilize two-stage pipelines—generating novel views via multi-view diffusion [46, 48, 50, 68] before reconstructing geometry [30, 40, 72, 88, 100]—or synthesize 3D representations directly [9, 42, 55, 73, 85], yet remain limited to isolated assets. For scenes, recent methods leverage camera trajectories for view synthesis [20, 25, 63, 76, 95, 98] but typically lack direct 3D outputs. Most related are LaRI [41], using layered ray intersections, and NOVA3R, which regresses point clouds from scene tokens. In contrast, our structured volumetric representation enables more complete and accurate amodal reconstruction, recovering geometry with smoother surfaces and sharper structural detail than unstructured point-based approaches.

3 Method

3.1 Problem Formulation

We address the task of amodal scene reconstruction from a single RGB image. Unlike traditional pixel-aligned methods [5, 78, 79, 80, 89] that recover only visible surfaces (e.g., *depthmaps* or *pointmaps*), our goal is to predict the complete 3D geometry within the observed view frustum. This includes recovering amodal geometry, surfaces occluded by foreground objects while maintaining the structural consistency of the scene-level environment.

Input and Output: Given a single input image $I \in \mathbb{R}^{H \times W \times 3}$, our model predicts a complete volumetric grid $V \in \mathbb{R}^{N \times N \times N}$ at a resolution of $N = 256$. We represent the scene geometry using a *Truncated Unsigned Distance Function*. Let $\mathcal{S} \subset \mathbb{R}^3$ denote the set of all physical surfaces in the scene contained within the view frustum. For each voxel center p within our grid, the value $V(p)$ is defined as the distance to the nearest surface in \mathcal{S} , truncated to a maximum value τ :

$$V(p) = \min(\text{dist}(p, \mathcal{S}), \tau) \quad (1)$$

We employ an unsigned formulation because indoor scenes are frequently non-watertight. A signed distance function requires a globally consistent inside/outside orientation — well-defined only for closed, watertight surfaces. For the common open geometries (single-sided walls, floors, furniture), no such orientation exists, making the sign geometrically meaningless beyond the immediate voxel neighborhood. The unsigned formulation avoids this ambiguity, as distance-to-nearest-surface is well-defined for any geometry regardless of topology.

Amodal Geometry Voxelization. The amodal surfaces \mathcal{S} can be obtained either directly from ground-truth dataset 3D meshes (if available) or by fusing depth from multiple surrounding views and transforming them into the target camera coordinate frame. To translate these amodal surfaces into our discrete target grid V , we employ a dynamic, frustum-aligned discretization strategy:

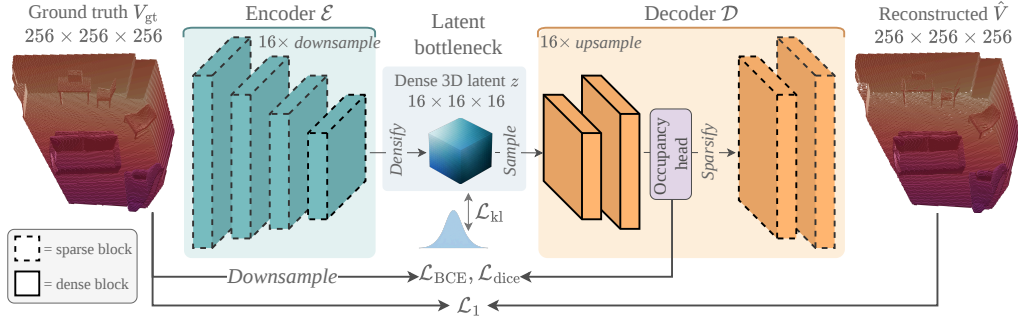


Figure 2: **3D VAE architecture.** The encoder compresses high-resolution sparse TUDF grids into a regularized dense latent via sparse convolutions. The decoder upsamples through dense layers, applies occupancy-guided sparsification, then restores the full-resolution TUDF via sparse convolutions.

- *Spatial scope:* To establish the precise boundaries of our volumetric grid, we compute an axis-aligned bounding box \mathcal{B} based on the *visible* geometry $P_{vis} \subset \mathcal{S}$. Any amodal surfaces in \mathcal{S} extending beyond \mathcal{B} are discarded, focusing the representation on the immediate captured scene.
- *Dynamic scaling and TUDF computation:* Rather than using fixed spatial extents [6, 43, 82], we calculate a dynamic, isotropic voxel size derived from the maximum dimension of \mathcal{B} and our target resolution N . We then discretize this bounded space and compute the truncated distance from each voxel center p to the cropped amodal surfaces, yielding the final TUDF grid V .

Downstream Geometric Extractions. The TUDF grid V serves as a flexible geometric proxy supporting diverse downstream formats. Surfaces can be efficiently extracted as point clouds by thresholding voxels with distance values below a small threshold $|V(p)| < \epsilon_{\text{tudf}}$, or converted into meshes via MeshUDF [23], producing clean, topologically consistent reconstructions.

3.2 Amodal 3D Reconstruction via Generative Framework

We formulate amodal scene reconstruction as estimating the conditional distribution $P(V|I)$ rather than a deterministic point estimate, overcoming a critical limitation of regression-based pipelines [5, 78]. Standard L_1 or L_2 losses predict the conditional expectation $\mathbb{E}[V|I]$, averaging over plausible geometries and producing over-smoothed reconstructions in occluded regions where evidence is absent. In contrast, our generative framework samples from $P(V|I)$, producing sharp and physically plausible completions by following the learned manifold of real-world scene structures.

3.3 Hybrid 3D VAE

To bridge the gap between high-resolution TUDF grids and the manageable latent space required by our generative model, we propose a hybrid sparse-dense 3D VAE (Fig. 2). This architecture exploits the inherent sparsity of scene geometry, where typically only 3–5% of voxels are surface-adjacent, while providing a fixed-size dense latent representation suitable for denoising via a latent transformer.

The encoder \mathcal{E} is designed to process natively sparse TUDF inputs with high computational efficiency. We convert the input grids into sparse tensors, retaining only active voxels within the truncation band τ , and apply a sequence of sparse 3D convolutions [13, 14, 71] with strided downsampling to reduce the spatial resolution by $16\times$. At the bottleneck, these features are densified into a regular 4D tensor $z \in \mathbb{R}^{16 \times 16 \times 16 \times C}$; by densifying only at this highly compressed scale, we maintain a manageable memory footprint while ensuring compatibility with dense generative architectures. Finally, a standard KL-divergence bottleneck regularizes the latent space to facilitate stable sampling [64].

The decoder \mathcal{D} faces the challenge of recovering a 256^3 grid from a dense 16^3 latent without prior knowledge of the target scene’s sparsity pattern. We address this through a hybrid dense-to-sparse decoding schedule:

- *Dense Upsampling ($16^3 \rightarrow 64^3$):* We first employ standard transposed 3D convolutions to upsample the latent to a 64^3 dense feature map, where the total voxel count ($\sim 262\text{K}$) remains tractable for dense computation.

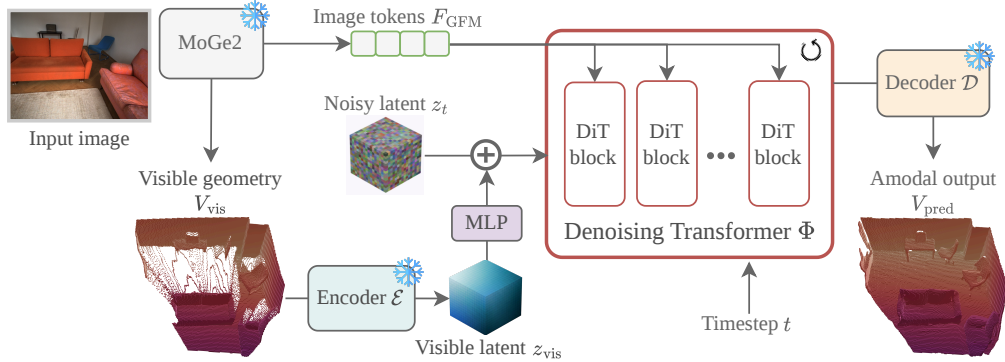


Figure 3: **Latent DiT architecture.** It operates in the compressed VAE latent space using a denoising transformer backbone with a flow-matching objective. We leverage a **dual conditioning** strategy, integrating high-level image tokens and explicit visible geometry, to guide the generative process and synthesize sharp, scene-consistent amodal structures.

- *Structure Prediction:* At the 64^3 stage, a binary occupancy head predicts a mask \hat{O} , indicating surface-adjacent voxels. This mask is used to sparsify the dense feature map, discarding empty regions and focusing subsequent computation solely on relevant geometric structures. We supervise this with a ground-truth occupancy mask O_{gt} derived from the TUDF grid.
- *Sparse Upsampling* ($64^3 \rightarrow 256^3$): The sparse features are processed by sparse 3D convolutions to reach the final 256^3 resolution, where memory cost scales linearly with surface area rather than cubically with volume, enabling high-resolution TUDF reconstruction at surface-adjacent voxels. The resulting sparse predictions are finally scattered back into a dense grid to obtain \hat{V} .

The VAE is trained end-to-end on our collected TUDF datasets using a composite loss function:

$$\mathcal{L}_{VAE} = \mathcal{L}_1(\hat{V}, V_{gt}) + \lambda_{bce} \text{BCE}(\hat{O}, O_{gt}) + \lambda_{dice} \text{Dice}(\hat{O}, O_{gt}) + \lambda_{kl} \mathcal{L}_{KL} \quad (2)$$

where \mathcal{L}_1 ensures precise distance reconstruction in active voxels, BCE and Dice losses supervise binary occupancy in the dense-to-sparse transition, and \mathcal{L}_{KL} regularizes the latent distribution. This hybrid design facilitates aggressive $16\times$ spatial compression while preserving the fine-grained details essential for high-fidelity amodal reconstruction.

3.4 Geometry-Conditioned Flow Matching

We implement a latent Diffusion Transformer Φ that learns the conditional velocity field to transport noise toward the distribution of scene TUDF latents. Starting from $z_0 \sim \mathcal{N}(0, \mathbf{I})$, we iteratively integrate the learned flow to obtain a clean latent z_1 , which is decoded by \mathcal{D} to produce the final TUDF prediction V_{pred} . To guide this process, we propose a dual-conditioning strategy that supplements global semantic-geometric context with local structural anchors. This architecture is shown in Fig. 3.

Global Geometric Priors. To inherit a robust, zero-shot understanding of spatial layouts, we condition the DiT on frozen features (F_{GFM}) extracted from a geometry foundation model. We utilize MoGe2 [79] as our backbone for its strong geometric estimation performance. Following standard practice, these features are processed via layer normalization and a linear projection before being injected into the cross-attention layers of each DiT block as keys and values (details in Appendix A.2).

Visible-Latent Inpainting. While global features provide context, we explicitly anchor the generative process to the observed scene by integrating known visible geometry directly into the latent space. Conceptually motivated by unprojection-based synthesis [10, 95], this strategy provides a partial guide for the model to complete occluded regions. Concretely, we obtain the visible pointmap P_{vis} from the GFM and convert it into a TUDF grid. This is then encoded into a visible latent z_{vis} via the frozen VAE encoder \mathcal{E} , and fused with the noisy latent as:

$$\tilde{z}_t = z_t + \text{MLP}_{zero}(z_{vis}) \quad (3)$$

where MLP_{zero} is a zero-initialized projection layer. This provides the model with an explicit structural prior over observed regions to guide completion of occluded geometry.

Table 1: **3D Reconstruction results** on SCRREAM dataset. We mark **best** and **second-best**.

Method	Visible			Occluded			Complete		
	CD↓	APD _{0.1} ↑	APD _{0.02} ↑	CD↓	APD _{0.1} ↑	APD _{0.02} ↑	CD↓	FS _{0.1} ↑	FS _{0.02} ↑
TripoSG [†] [42]	13.75	44.74	10.08	10.73	55.34	13.37	11.10	53.56	13.30
TRELLIS [†] [85]	11.63	54.05	15.00	9.65	60.70	17.52	10.72	56.69	15.59
DUST3R [80]	4.46	92.13	30.81	23.21	38.43	8.40	10.85	69.90	24.17
VGGT [77]	3.04	96.98	45.62	17.40	44.60	10.04	6.77	80.45	36.77
DA3 [44]	2.37	98.98	54.34	16.16	48.31	10.76	6.16	82.86	43.10
DepthPro [5]	3.85	95.10	33.30	15.31	48.66	9.71	6.90	80.36	28.44
MoGe2 [79]	2.28	98.38	37.42	15.29	50.09	10.30	5.74	83.96	45.84
LaRI [41]	3.43	95.48	41.61	5.23	85.25	29.31	5.19	85.43	30.85
NOVA3R [11]	3.20	96.77	43.75	3.56	94.86	41.21	3.43	95.26	45.12
VolFill (ours)	2.84	96.10	56.73	3.46	92.70	55.53	3.03	95.03	54.83

Note: All reported numbers are scaled by a factor of 10^2 . [†]: methods trained on object-centric datasets.

Flow Matching Objective. We train our model using a rectified flow objective [1, 45, 47]. Given a ground-truth TUDF latent z_1 and Gaussian noise $\epsilon \sim \mathcal{N}(0, I)$, we define a linear interpolation path $z_t = (1 - t)\epsilon + tz_1$. Our training objective is to minimize the expected mean squared error between the predicted velocity u_Φ and the constant velocity target $u^* = z_1 - \epsilon$:

$$\mathcal{L}_{\text{gen}} = \mathbb{E}_{t, \epsilon, z_1} [\|u_\Phi(\tilde{z}_t, t, F_{\text{GFM}}) - u^*\|^2] \quad (4)$$

During inference, we integrate the learned ODE from $t = 0$ to $t = 1$ via Euler steps, then decode the resulting latent through the VAE decoder \mathcal{D} to produce the final TUDF grid V_{pred} .

3.5 Model Architecture

The 3D VAE employs a symmetric encoder-decoder architecture, where each stage consists of two ResNet-style blocks and doubles the channel dimension at each resolution step. The primary distinction lies in the convolution implementation: while the encoder uses sparse 3D convolutions across all stages for efficiency, the decoder utilizes standard dense convolutions at lower resolutions ($16^3 \rightarrow 64^3$) before transitioning to sparse convolutions for the remaining higher stages ($64^3 \rightarrow 256^3$). Finally, the binary occupancy and TUDF regression heads are both implemented as 2-layer MLPs.

The latent DiT features 12 transformer blocks with a 768-dimensional hidden state and 12-head self-attention [74]. Each block integrates image tokens F_{GFM} through cross-attention, utilizes AdaLN [57] for timestep conditioning, and employs a GELU feed-forward network [27] with $4\times$ expansion. To stabilize flow-matching training, we apply QK-normalization [28] via RMSNorm [96] before the attention layers. The detailed architecture of the DiT block is provided in Appendix A.2.

4 Experiments

4.1 Implementation details

Training datasets. We train VolFill on a combination of 3D-FRONT [17] and ScanNet++ [91], following the data splits established by LaRI [41] and NOVA3R [11]. For the synthetic 3D-FRONT dataset, we utilize 18k room-level scenes to construct 96k image-TUDF pairs. For real-world ScanNet++ data, we aggregate multi-view depth maps and back-project them into a unified point cloud, followed by a filtering and voxelization pipeline to obtain amodal TUDF grids for 46k samples.

Training protocol. Training is conducted on two NVIDIA A6000/L40S GPUs in two stages: in Stage 1, the 3D VAE is trained for 20 epochs with a total batch size of 24, requiring 3 days to converge; in Stage 2, the latent transformer is trained for 100 epochs with a total batch size of 48, completing in 2.5 days. Both stages utilize the AdamW optimizer with an initial learning rate of 10^{-4} , a cosine scheduler, and mixed-precision training. More training details are provided in Appendix A.3.

Table 2: **3D Reconstruction results** on NRGB-D dataset.

Method	Visible			Occluded			Complete		
	CD↓	APD _{0.1} ↑	APD _{0.05} ↑	CD↓	APD _{0.1} ↑	APD _{0.05} ↑	CD↓	FS _{0.1} ↑	FS _{0.05} ↑
DUST3R [80]	4.46	92.13	69.08	23.21	38.43	22.18	10.85	69.90	51.61
VGGT [77]	3.25	96.51	81.63	21.38	42.15	26.35	9.44	73.85	58.91
DA3 [44]	2.69	98.68	89.24	21.65	43.34	28.15	8.88	74.87	63.52
DepthPro [5]	4.38	92.56	67.41	22.19	41.69	24.64	10.03	71.30	51.60
MoGe2 [79]	3.67	98.46	89.55	21.12	43.31	28.43	9.31	75.35	61.93
LaRI [41]	4.09	93.77	73.51	13.04	64.53	45.82	9.79	68.37	43.47
NOVA3R [11]	5.05	89.63	64.44	8.31	80.50	58.55	8.19	73.99	49.92
VolFill (ours)	4.52	90.06	69.22	6.97	81.41	62.48	6.44	84.35	63.92

Latent Sampling and Inference. During training, we sample the flow-matching timestep $t \in [0, 1]$ from a logit-normal distribution to prioritize training on higher-noise regimes following [85]. We apply classifier-free guidance (CFG) by dropping conditions with a probability of 0.1. At inference, we utilize an ODE solver with 50 steps and set the CFG guidance scale to 3.0.

Our source code and evaluation procedure will be published on our project page upon acceptance.

4.2 Comparison with prior work

Datasets. We evaluate our model on two benchmarks. (1) SCREAM [33] consists of 460 samples with complete, high-quality scanned meshes, providing reliable ground truth for both visible and occluded surfaces. (2) Neural RGB-D [2] is adapted for amodal evaluation by aggregating its depth maps into a global coordinate frame via camera trajectories and multi-view fusion to reconstruct occluded surfaces. We manually curate the dataset to retain only samples with high “amodal richness”, excluding scenes where occluded coverage is sparse or hidden geometry negligible. This filtering ensures the benchmark effectively tests scene completion in complex environments.

Metrics. Following [11, 41], we evaluate reconstruction quality using Chamfer Distance (CD↓) and F-score (FS_γ↑) at thresholds $\gamma \in \{0.02, 0.05, 0.10\}$. Since our non-pixel-aligned formulation lacks point-wise correspondences, we recover the optimal similarity transformation (scale, rotation, translation) by minimizing Chamfer Distance via gradient descent prior to evaluation. We assess performance across *visible*, *occluded*, and *complete* subsets. We employ a one-way Chamfer Distance (measuring from ground truth to prediction) in the *visible* and *occluded* subsets to assess regional coverage without penalizing valid predicted geometry that falls outside the target subset, and replace the F-score with a threshold coverage score APD_γ↑, measuring the percentage of ground truth points successfully reconstructed within threshold γ . We further evaluate generative quality via Fréchet Point Cloud Distance (FPD↓) [51], using a pretrained Uni3D [99] to measure distributional similarity in a semantically aligned 3D feature space.

Baselines. We compare against four groups. *Object-level generative models* (TRELIS [85], TripoSG [42]) show the effect of naively applying strong 3D object priors to full unmasked scenes. *Multi-view geometry* models (DUST3R [80], VGGT [77], DepthAnything3 [44]) and *single-view pixel-aligned* estimators (MoGe2 [79], DepthPro [5]) establish visible-surface reference performance but fundamentally cannot recover occlusions. Finally, *scene-level amodal reconstruction* methods (LaRI [41], NOVA3R [11]) serve as our primary competitors.

Results. Tables 1 and 2 report performance on SCREAM and NRGB-D datasets respectively. While pixel-aligned estimators [44, 79] naturally dominate visible metrics, they inherently fail to reconstruct occluded regions. Among methods designed for amodal completion [11, 41], our approach establishes a new state-of-the-art for complete scene reconstruction, achieving a significant performance leap under the stringent FS_{0.02} metric. On NRGB-D visible metrics, we maintain performance near that of LaRI [41], while achieving substantially stronger reconstruction across occluded and complete geometry — confirming that our volumetric formulation provides a far more robust foundation for full-scene understanding than existing point-cloud or ray-based baselines.




Figure 4: **Qualitative comparison.** VolFill synthesizes sharp, high-fidelity geometry, whereas LaRI produces layered artifacts (red circle) and holes, and NOVA3R yields noisy, unstructured point scatters (green circle).




Figure 5: **Mesh reconstruction comparison.** LaRI and NOVA3R produce fragmented and noisy meshes due to their unstructured outputs, whereas VolFill directly extracts clean, topologically consistent surfaces from the structured TUDF grid.

Table 3 presents our metric-scale evaluation, where we assess the reconstruction of amodal point cloud by applying a rigid transformation to align the predicted and ground-truth coordinates. While our model leverages internal MoGe2 predictions to infer absolute scale, we provide baselines [11, 41] with a standalone MoGe2 model to lift their normalized outputs into metric space. Our approach significantly outperforms these methods, confirming its ability to reconstruct faithful metric-scale amodal geometry. Our superior FPD score demonstrates that the synthesized geometry is distributionally well-aligned with ground-truth scenes within a semantically meaningful feature space.

Fig. 4 presents qualitative comparisons of predicted point clouds. LaRI produces severe layered artifacts inherited from its ray-aligned formulation and leaves many occluded regions empty; NOVA3R recovers broader scene coverage but produces a noisy, unstructured point cloud. VolFill recovers a more complete geometry with higher structural fidelity in both visible and occluded regions.

Furthermore, the structured nature of our TUDF representation enables direct surface extraction [23], yielding smooth and topologically consistent meshes, as qualitatively evidenced in Fig. 5. In contrast, [11, 41] require Poisson reconstruction [35] to produce meshes, a slower post-processing step that nevertheless results in significant noise and disconnected regions wherever predicted density is sparse — further demonstrating that representing amodal scenes as structured distance functions is essential for recovering physically plausible surface geometry. See Appendix A.4 for more results.

4.3 Ablation study

We evaluate our core design choices through ablation studies on the SCRREAM dataset to justify our architectural components and conditioning strategies for effective amodal scene reconstruction.

VAE Decoder Design is analyzed in Table 4. We evaluate reconstruction fidelity using the TUDF L1 distance within active cells and the occupancy Intersection over Union (IoU) from the thresholded TUDF. While a dense-only decoder is highly resource-intensive, a purely sparse variant reduces

Table 3: **Metric geometry evaluation.**

Method	CD↓	FS _{0.02} ↑	FPD↓
LaRI	11.43	11.47	–
NOVA3R	6.56	22.23	16.50
VolFill (ours)	3.88	44.90	4.86

Table 5: **Conditioning ablation.**

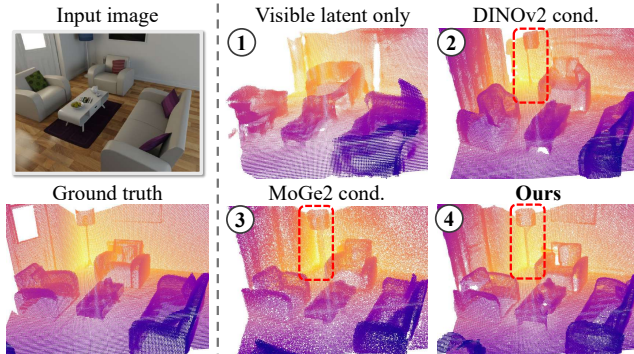
Config.	F_{GFM}	z_{vis}	CD↓	FS _{0.02} ↑
Single	\times	<i>Add</i>	6.80	31.72
	<i>MoGe2</i>	\times	3.81	46.80
Dual	<i>MoGe2</i>	<i>Concat</i>	3.66	48.94
	<i>MoGe2</i>	<i>Add</i>	3.03	54.83

Table 6: **Foundation model ablation.**

Conditioning	CD↓	FS _{0.02} ↑
DINOv2	5.46	35.98
VGGT	4.32	42.30
MoGe2	3.81	46.80

Table 4: **VAE decoder design ablation.**

Decoder	L1 Dist↓	IoU↑	Time (s)↓	Mem. (GB)↓
Dense-only	0.49	93.2	0.63	9.2
Sparse-only	0.48	93.0	0.09	2.5
Hybrid (Ours)	0.43	93.5	0.07	1.7

Figure 6: **Qualitative conditioning ablation.** ① Visible-only geometry fails to complete the scene; ②, ③ image-only tokens result in distorted results; ④ our dual-conditioning synthesizes sharp, high-fidelity amodal geometry.

overhead but compromises accuracy. By integrating both, our hybrid design achieves the best reconstruction quality alongside the lowest latency and memory footprint, validating its efficiency for compressing sparse 3D structures into a compact latent space.

Ablation on visual and geometric priors is shown in Table 5 and Fig. 6. Relying exclusively on the visible latent (①) fails to recover unobserved structures. While image-only conditioning (F_{GFM}) enables amodal reasoning, using DINOv2 (②) leads to distorted geometry, and replacing it with MoGe2 (③)—though more realistic—remains insufficient for recovering sharp, fine-grained details. Our full dual-conditioning strategy (④) achieves the highest fidelity, where additive fusion via zero-initialized MLPs (*Add*) outperforms concatenation (*Concat*), demonstrating that the synergy of global visual context and explicit geometric grounding is essential for accurate amodal reconstruction.

Ablation on different geometry foundation models is analyzed in Table 6. Despite its semantic richness, DINOv2 performs worst due to a lack of 3D spatial anchors. Geometry models like VGGT significantly improve results, while MoGe2 achieves the best performance. This suggests that structural consistency is a stronger driver for amodal reconstruction than general semantic reasoning.

5 Limitations and Conclusion

Limitations. While VolFill produces sharper geometry with fewer artifacts than NOVA3R [11] and LaRI [41], small holes can occasionally persist in unobserved regions, reflecting the challenge of synthesizing geometry without visual evidence. Scaling model capacity and data diversity will be essential to bridge these remaining topological gaps. Additionally, as an iterative generative framework, inference with 50 steps and CFG requires 1.4s on an RTX 4090, remaining slower than regression-based baselines.

Conclusion. We introduced VolFill, a generative framework for amodal 3D scene reconstruction. A hybrid 3D VAE compresses sparse TUDF grids into a compact latent space, while a latent Diffusion Transformer denoises this representation to recover the complete scene geometry; frozen geometric priors from MoGe2 provide the necessary spatial context. This shifts reconstruction from deterministic pixel-aligned regression to a structured generative manifold, enabling physically plausible completions in occluded regions. VolFill produces sharper geometry and outperforms existing point-cloud and ray-based methods.

Acknowledgements Evangelos Kalogerakis has received funding from the European Research Council (ERC) under the Horizon research and innovation programme (Grant agreement No. 101124742).

References

- [1] Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*, 2023.
- [2] Dejan Azinović, Ricardo Martin-Brualla, Dan B Goldman, Matthias Nießner, and Justus Thies. Neural rgb-d surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6290–6301, 2022.
- [3] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4009–4018, 2021.
- [4] Shariq Farooq Bhat, Reiner Birkl, Diana Wofk, Peter Wonka, and Matthias Muller. Zoedepth: Zero-shot transfer by combining relative and metric depth. *ArXiv*, abs/2302.12288, 2023.
- [5] Aleksei Bochkovskii, Amaël C. Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan R Richter, and Vladlen Koltun. Depth pro: Sharp monocular metric depth in less than a second. *arXiv preprint arXiv:2410.02073*, 2024.
- [6] Anh-Quan Cao and Raoul de Charette. Monoscene: Monocular 3d semantic scene completion. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3981–3991, 2021.
- [7] Anh-Quan Cao and Tuan-Hung Vu. Occany: Generalized unconstrained urban 3d occupancy. In *CVPR*, 2026.
- [8] Jiahao Chang, Chongjie Ye, Yushuang Wu, Yuantao Chen, Yidan Zhang, Zhongjin Luo, Chenghong Li, Yihao Zhi, and Xiaoguang Han. Reconviagen: Towards accurate multi-view 3d object reconstruction via generation. *arXiv preprint arXiv:2510.23306*, 2025.
- [9] Hansheng Chen, Jiatao Gu, Anpei Chen, Wei Tian, Zhuowen Tu, Lingjie Liu, and Haoran Su. Single-stage diffusion nerf: A unified approach to 3d generation and reconstruction. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2416–2425, 2023.
- [10] Kaihua Chen, Tarasha Khurana, and Deva Ramanan. Reconstruct, inpaint, test-time finetune: Dynamic novel-view synthesis from monocular videos. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.
- [11] Weirong Chen, Chuanxia Zheng, Ganlin Zhang, Andrea Vedaldi, and Daniel Cremers. Nova3r: Non-pixel-aligned visual transformer for amodal 3d reconstruction. In *The Fourteenth International Conference on Learning Representations*, 2026.
- [12] Yen-Chi Cheng, Hsin-Ying Lee, S. Tulyakov, Alexander G. Schwing, and Liangyan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4456–4465, 2022.
- [13] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3075–3084, 2019.
- [14] Spconv Contributors. Spconv: Spatially sparse convolution library. <https://github.com/traveller59/spconv>, 2022.
- [15] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27, 2014.
- [16] Xianze Fang, Jingnan Gao, Zhe Wang, Zhuo Chen, Xingyu Ren, Jiangjing Lyu, Qiaomu Ren, Zhonglei Yang, Xiaokang Yang, Yichao Yan, and Chengfei Lyu. Dens3r: A foundation model for 3d geometry prediction. *ArXiv*, abs/2507.16290, 2025.

- [17] Huan Fu, Bowen Cai, Lin Gao, Ling-Xiao Zhang, Cao Li, Zengqi Xun, Chengyue Sun, Yiyun Fei, Yu qiong Zheng, Ying Li, Yi Liu, Peng Liu, Lin Ma, Le Weng, Xiaohang Hu, Xin Ma, Qian Qian, Rongfei Jia, Binqiang Zhao, and Hao Helen Zhang. 3d-front: 3d furnished rooms with layouts and semantics. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10913–10922, 2020.
- [18] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2002–2011, 2018.
- [19] Xiao Fu, Wei Yin, Mu Hu, Kaixuan Wang, Yuexin Ma, Ping Tan, Shaojie Shen, Dahua Lin, and Xiaoxiao Long. Geowizard: Unleashing the diffusion priors for 3d geometry estimation from a single image. In *European Conference on Computer Vision*, pages 241–258. Springer, 2024.
- [20] Ruiqi Gao, Aleksander Holynski, Philipp Henzler, Arthur Brussee, Ricardo Martin-Brualla, Pratul P. Srinivasan, Jonathan T. Barron, and Ben Poole. Cat3d: Create anything in 3d with multi-view diffusion models. *ArXiv*, abs/2405.10314, 2024.
- [21] Gonzalo Martin Garcia, Karim Abou Zeid, Christian Schmidt, Daan De Geus, Alexander Hermans, and Bastian Leibe. Fine-tuning image-conditional diffusion models is easier than you think. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 753–762. IEEE, 2025.
- [22] Benjamin Graham and Laurens van der Maaten. Submanifold sparse convolutional networks. *ArXiv*, abs/1706.01307, 2017.
- [23] Benoît Guillard, Federico Stella, and P. Fua. Meshudf: Fast and differentiable meshing of unsigned distance field networks. In *European Conference on Computer Vision*, 2021.
- [24] Vitor Campanholo Guizilini, Igor Vasiljevic, Di Chen, Rares Ambrus, and Adrien Gaidon. Towards zero-shot scale-aware monocular depth estimation. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9199–9209, 2023.
- [25] Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, and Ceyuan Yang. Cameractrl: Enabling camera control for text-to-video generation. *ArXiv*, abs/2404.02101, 2024.
- [26] Jing He, Haodong Li, Wei Yin, Yixun Liang, Leheng Li, Kaiqiang Zhou, Hongbo Zhang, Bingbing Liu, and Ying-Cong Chen. Lotus: Diffusion-based visual foundation model for high-quality dense prediction. *arXiv preprint arXiv:2409.18124*, 2024.
- [27] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv: Learning*, 2016.
- [28] Alex Henry, Prudhvi Raj Dachapally, Shubham Vivek Pawar, and Yuxuan Chen. Query-key normalization for transformers. In *Findings*, 2020.
- [29] Derek Hoiem, Alexei A Efros, and Martial Hebert. Recovering surface layout from an image. *International Journal of Computer Vision*, 75(1):151–172, 2007.
- [30] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. Lrm: Large reconstruction model for single image to 3d. *ArXiv*, abs/2311.04400, 2023.
- [31] Mu Hu, Wei Yin, Chi Zhang, Zhipeng Cai, Xiaoxiao Long, Hao Chen, Kaixuan Wang, Gang Yu, Chunhua Shen, and Shaojie Shen. Metric3d v2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):10579–10596, 2024.
- [32] Wooseok Jang, Seonghu Jeon, Jisang Han, Jinhyeok Choi, Minkyung Kwon, Seungryong Kim, Saining Xie, and Sainan Liu. Repurposing geometric foundation models for multi-view diffusion. 2026.

- [33] Hyunjun Jung, Weihang Li, Shun cheng Wu, William Bittner, Nikolas Brasch, Jifei Song, Eduardo P'erez-Pellitero, Zhensong Zhang, Arthur Moreau, Nassir Navab, and Benjamin Busam. Screeam : Scan, register, render and map:a framework for annotating accurate and dense 3d indoor scenes with a benchmark. *ArXiv*, abs/2410.22715, 2024.
- [34] Kevin Karsch, Ce Liu, and Sing Bing Kang. Depth Transfer: Depth Extraction from Video Using Non-Parametric Sampling . *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 36(11):2144–2158, November 2014.
- [35] Michael M. Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Eurographics Symposium on Geometry Processing*, 2006.
- [36] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9492–9502, 2024.
- [37] Nikhil Keetha, Norman Müller, Johannes Schönberger, Lorenzo Porzi, Yuchen Zhang, Tobias Fischer, Arno Knapitsch, Duncan Zauss, Ethan Weber, Nelson Antunes, Jonathon Luiten, Manuel López-Antequera, Samuel Rota Bulò, Christian Richardt, Deva Ramanan, Sebastian Scherer, and Peter Kotschieder. Mapanything: Universal feed-forward metric 3d reconstruction. *ArXiv*, abs/2509.13414, 2025.
- [38] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), July 2023.
- [39] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019.
- [40] Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model. *ArXiv*, abs/2311.06214, 2023.
- [41] Rui Li, Biao Zhang, Zhenyu Li, Federico Tombari, and Peter Wonka. Lari: Layered ray intersections for single-view 3d geometric reasoning. *ArXiv*, abs/2504.18424, 2025.
- [42] Yangguang Li, Zi-Xin Zou, Zexiang Liu, Dehui Wang, Yuanzhi Liang, Zhipeng Yu, Xingchao Liu, Yuanchen Guo, Ding Liang, Wanli Ouyang, and Yan-Pei Cao. Triposg: High-fidelity 3d shape synthesis using large-scale rectified flow models. *IEEE transactions on pattern analysis and machine intelligence*, PP, 2025.
- [43] Yiming Li, Zhiding Yu, Christopher Bongsoo Choy, Chaowei Xiao, José Manuel Álvarez, Sanja Fidler, Chen Feng, and Anima Anandkumar. Voxformer: Sparse voxel transformer for camera-based 3d semantic scene completion. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9087–9098, 2023.
- [44] Haotong Lin, Sili Chen, Jun Hao Liew, Donny Y. Chen, Zhenyu Li, Guang Shi, Jiashi Feng, and Bingyi Kang. Depth anything 3: Recovering the visual space from any views. *ArXiv*, abs/2511.10647, 2025.
- [45] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *ArXiv*, abs/2210.02747, 2022.
- [46] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9264–9275, 2023.
- [47] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *ArXiv*, abs/2209.03003, 2022.

- [48] Yuan Liu, Chu-Hsing Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. *ArXiv*, abs/2309.03453, 2023.
- [49] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Trans. Graph.*, 38(4):65:1–65:14, July 2019.
- [50] Xiaoxiao Long, Yuanchen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, and Wenping Wang. Wonder3d: Single image to 3d using cross-domain diffusion. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9970–9980, 2023.
- [51] Quan Meng, Yujin Chen, Lei Li, Matthias Nießner, and Angela Dai. Seen2scene: Completing realistic 3d scenes with visibility-guided flow, 2026.
- [52] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [53] Norman Muller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulò, Peter Kotschieder, and Matthias Nießner. Diffrrf: Rendering-guided 3d radiance field diffusion. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4328–4338, 2022.
- [54] Tuan Duc Ngo, Jiahui Huang, Seoung Wug Oh, Kevin Blackburn-Matzen, Evangelos Kalogerakis, Chuang Gan, and Joon-Young Lee. Dage: Dual-stream architecture for efficient and fine-grained geometry estimation. In *CVPR*, 2026.
- [55] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *ArXiv*, abs/2212.08751, 2022.
- [56] Michael Niemeyer, Lars M. Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3501–3512, 2019.
- [57] William S. Peebles and Saining Xie. Scalable diffusion models with transformers. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4172–4182, 2022.
- [58] Duc-Hai Pham, Tung Do, Phong Nguyen, Binh-Son Hua, Khoi Nguyen, and Rang Nguyen. Sharpdepth: Sharpening metric depth predictions using diffusion distillation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 17060–17069, 2025.
- [59] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc van Gool, and Fisher Yu. Unidepth: Universal monocular metric depth estimation. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10106–10116, 2024.
- [60] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023.
- [61] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen 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.
- [62] Xuanchi Ren, Jiahui Huang, Xiaohui Zeng, Ken Museth, Sanja Fidler, and Francis Williams. Xcube: Large-scale 3d generative modeling using sparse voxel hierarchies. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4209–4219, 2023.
- [63] Xuanchi Ren, Tianchang Shen, Jiahui Huang, Huan Ling, Yifan Lu, Merlin Nimier-David, Thomas Muller, Alexander Keller, Sanja Fidler, and Jun Gao. Gen3c: 3d-informed world-consistent video generation with precise camera control. *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6121–6132, 2025.

- [64] Robin Rombach, A. Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, 2021.
- [65] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021.
- [66] Ashutosh Saxena, Sung Chung, and Andrew Ng. Learning depth from single monocular images. *Advances in neural information processing systems*, 18, 2005.
- [67] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Make3d: Learning 3d scene structure from a single still image. *IEEE transactions on pattern analysis and machine intelligence*, 31(5):824–840, 2008.
- [68] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and X. Yang. Mvdream: Multi-view diffusion for 3d generation. *ArXiv*, abs/2308.16512, 2023.
- [69] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.
- [70] Shuran Song, Fisher Yu, Andy Zeng, Angel X. Chang, Manolis Savva, and Thomas A. Funkhouser. Semantic scene completion from a single depth image. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 190–198, 2016.
- [71] Haotian Tang, Zhijian Liu, Xiuyu Li, Yujun Lin, and Song Han. TorchSparse: Efficient Point Cloud Inference Engine. In *Conference on Machine Learning and Systems (MLSys)*, 2022.
- [72] Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. In *European Conference on Computer Vision*, 2024.
- [73] Tencent Hunyuan3D Team. Hunyuan3d 1.0: A unified framework for text-to-3d and image-to-3d generation, 2024.
- [74] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Neural Information Processing Systems*, 2017.
- [75] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2022–2030, 2018.
- [76] Chaoyang Wang, Peiye Zhuang, Tuan Duc Ngo, Willi Menapace, Aliaksandr Siarohin, Michael Vasilkovsky, Ivan Skorokhodov, Sergey Tulyakov, Peter Wonka, and Hsin-Ying Lee. 4real-video: Learning generalizable photo-realistic 4d video diffusion. *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17723–17732, 2024.
- [77] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vgg: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5294–5306, 2025.
- [78] Ruicheng Wang, Sicheng Xu, Cassie Dai, Jianfeng Xiang, Yu Deng, Xin Tong, and Jiaolong Yang. Moge: Unlocking accurate monocular geometry estimation for open-domain images with optimal training supervision. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5261–5271, 2025.
- [79] Ruicheng Wang, Sicheng Xu, Yue Dong, Yu Deng, Jianfeng Xiang, Zelong Lv, Guangzhong Sun, Xin Tong, and Jiaolong Yang. Moge-2: Accurate monocular geometry with metric scale and sharp details. *arXiv preprint arXiv:2507.02546*, 2025.
- [80] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024.

- [81] Yifan Wang, Jianjun Zhou, Haoyi Zhu, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen, Jiangmiao Pang, Chunhua Shen, and Tong He. Scalable permutation-equivariant visual geometry learning. *arXiv preprint arXiv:2507.13347*, 2025.
- [82] Yi Wei, Linqing Zhao, Wenzhao Zheng, Zhengbiao Zhu, Jie Zhou, and Jiwen Lu. Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 21672–21683, 2023.
- [83] Haoyu Wu, Diankun Wu, Tianyu He, Junliang Guo, Yang Ye, Yueqi Duan, and Jiang Bian. Geometry forcing: Marrying video diffusion and 3d representation for consistent world modeling. *ArXiv*, abs/2507.07982, 2025.
- [84] Jianfeng Xiang, Xiaoxue Chen, Sicheng Xu, Ruicheng Wang, Zelong Lv, Yu Deng, Hong Zhu, Yue Dong, Hao Zhao, Nicholas Jing Yuan, and Jiaolong Yang. Native and compact structured latents for 3d generation. *ArXiv*, abs/2512.14692, 2025.
- [85] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21469–21480, 2024.
- [86] Mingyang Xie, Numair Khan, Tianfu Wang, Naina Dhir, Seonghyeon Nam, Haitao Yang, Zhuo Hui, Christopher Metzler, Andrea Vedaldi, Hamed Pirsiavash, and Lei Luo. Lavr: Scene latent conditioned generative video trajectory re-rendering using large 4d reconstruction models. *ArXiv*, abs/2601.14674, 2026.
- [87] Gangwei Xu, Haotong Lin, Hongcheng Luo, Xianqi Wang, Jingfeng Yao, Lianghai Zhu, Yuechuan Pu, Cheng Chi, Haiyang Sun, Bing Wang, et al. Pixel-perfect depth with semantics-prompted diffusion transformers. *arXiv preprint arXiv:2510.07316*, 2025.
- [88] Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and Gordon Wetzstein. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. *ArXiv*, abs/2403.14621, 2024.
- [89] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang 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.
- [90] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *Advances in Neural Information Processing Systems*, 37:21875–21911, 2024.
- [91] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela 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.
- [92] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9009–9019, 2023.
- [93] Hao Yu, Haotong Lin, Jiawei Wang, Jiaxin Li, Yida Wang, Xueyang Zhang, Yue Wang, Xiaowei Zhou, Ruizhen Hu, and Sida Peng. Infinidepth: Arbitrary-resolution and fine-grained depth estimation with neural implicit fields. *ArXiv*, abs/2601.03252, 2026.
- [94] Mark Yu, Wenbo Hu, Jinbo Xing, and Ying Shan. Trajectorycrafter: Redirecting camera trajectory for monocular videos via diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 100–111, October 2025.
- [95] 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. *IEEE transactions on pattern analysis and machine intelligence*, PP, 2024.

- [96] Biao Zhang and Rico Sennrich. Root Mean Square Layer Normalization. In *Advances in Neural Information Processing Systems 32*, Vancouver, Canada, 2019.
- [97] Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. Gs-irm: Large reconstruction model for 3d gaussian splatting. *ArXiv*, abs/2404.19702, 2024.
- [98] Jensen Zhou, Hang Gao, Vikram S. Voleti, Aaryaman Vasishtha, Chun-Han Yao, Mark Boss, Philip Torr, Christian Rupprecht, and Varun Jampani. Stable virtual camera: Generative view synthesis with diffusion models. *ArXiv*, abs/2503.14489, 2025.
- [99] Junsheng Zhou, Jinsheng Wang, Baorui Ma, Yu-Shen Liu, Tiejun Huang, and Xinlong Wang. Uni3d: Exploring unified 3d representation at scale. In *International Conference on Learning Representations (ICLR)*, 2024.
- [100] Peiye Zhuang, Songfang Han, Chaoyang Wang, Aliaksandr Siarohin, Jiaxu Zou, Michael Vasilkovsky, Vladislav Shakhrai, Sergey Korolev, S. Tulyakov, and Hsin-Ying Lee. Gtr: Improving large 3d reconstruction models through geometry and texture refinement. *ArXiv*, abs/2406.05649, 2024.

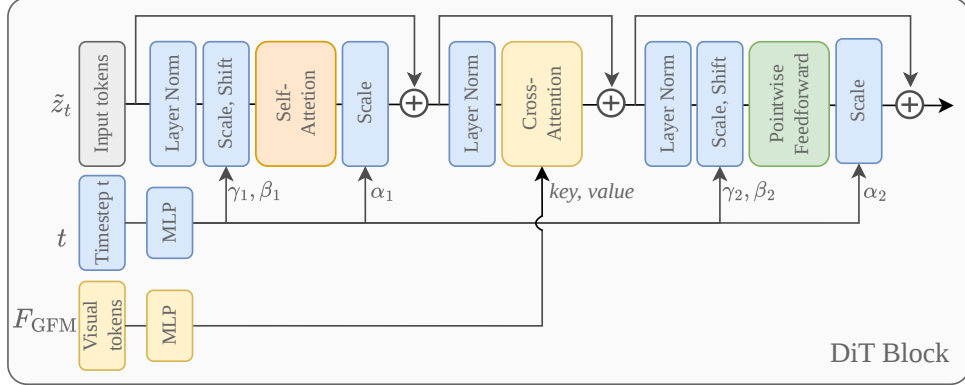


Figure 7: Illustration of the DiT block.

A Technical appendices and supplementary material

A.1 Volumetric Representation Design Choice

We represent amodal 3D geometry as a Truncated Unsigned Distance Field (TUDF), where each voxel stores its distance to the nearest surface, clipped to a maximum of τ voxels. This choice is motivated by the limitations of the two below alternatives.

Binary occupancy marks each voxel as occupied or empty, providing only a single bit of information per voxel, while TUDF encodes the continuous distance to the nearest surface. This continuous formulation provides smooth gradients that guide the decoder toward surface boundaries, whereas binary losses offer no directional signal within the truncation band. Additionally, TUDF enables precise sub-voxel surface localization through isosurface extraction, while binary occupancy is restricted to discrete voxel centers.

Truncated Signed Distance Functions (TSDF) require a globally consistent inside/outside orientation, which is well-defined only for closed, watertight surfaces. Indoor scenes are frequently non-watertight (single-sided walls, floors, and furniture surfaces do not enclose a volume) so no such orientation exists. Consequently, TSDF provides no practical advantage over TUDF for open-geometry scenes. The unsigned formulation avoids this ambiguity entirely, as distance-to-nearest-surface remains well-defined regardless of scene topology.

A.2 Architecture Details

Fig. 7 describes the internal components of our DiT block, which processes latent tokens \tilde{z}_t through sequential self-attention, cross-attention, and pointwise feedforward stages. To inject temporal information, the timestep t is mapped via an MLP to generate adaptive Layer Norm parameters $\{\gamma_i, \beta_i, \alpha_i\}$ that modulate and scale the features. Visual priors are integrated by passing foundation model tokens F_{GFM} through an MLP to serve as key and value pairs for cross-attention.

A.3 Implementation Details

For the construction of the Truncated Unsigned Distance Function (TUDF) grids (Section 3.1), the truncation distance τ is set to 3.0. In the VAE architecture (Section 3.3), the bottleneck latent space ($16 \times$ downsample) utilizes a channel dimension of $C = 16$. To derive the low-resolution ground-truth binary occupancy mask from the high-resolution TUDF grid, we apply 3D max-pooling to accurately preserve the structural bounds of the geometry. Finally, the weighting coefficients for the composite VAE objective (Equation 2) are empirically balanced and set to $\lambda_{bce} = 0.2$, $\lambda_{dice} = 0.2$, and $\lambda_{kl} = 1e - 6$.

A.4 Additional Qualitative Results

We refer the reader to the supplementary video for animated visualizations of the reconstructed 3D scenes.

Comparison with pixel-aligned geometry approaches is presented in Fig. 8. While these baselines [44, 79] achieve high accuracy on visible surfaces, they are inherently restricted to the camera’s line-of-sight, leaving significant holes and “shadows” in occluded regions. In contrast, VolFill synthesizes physically plausible hidden structures, producing a continuous and structurally coherent scene representation resolving the visibility constraints of prior work.

Comparison of mesh reconstruction. As shown in Fig. 9, LaRI and NOVA3R can produce reasonable point clouds in some cases, yet their unstructured outputs consistently yield fragmented, artifact-heavy meshes. Our approach produces cleaner point clouds and topologically consistent meshes that better reflect the physical scene geometry.

Qualitative results of the hybrid 3D VAE are shown in Fig. 10. Our VAE effectively compresses sparse high-resolution TUDF grids into a compact dense latent while preserving fine geometric details upon reconstruction. For larger scenes, a slight reduction in fidelity is observed given the fixed latent capacity, though overall structure and surface topology remain well-preserved — confirming the latent space provides a faithful scene representation for the downstream diffusion process.

Generative inference trajectory. Fig. 11 illustrates the iterative refinement of the amodal geometry throughout the denoising trajectory. Even at early stages ($t = 3$), the framework successfully predicts the coarse overall structure of the scene. As the latent representation becomes progressively cleaner, the synthesized geometry sharpens significantly around $t = 15$, followed by continuous, fine-grained structural refinement until the final inference step.

B Computing Resources

Our research utilized multiple computing servers, consisting of a high-performance cluster for model training and a dedicated local workstation for data preparation, evaluation, and visualization.

Hardware Specifications.

- **Cluster Training:** Primary training was conducted on a cluster node equipped with $2 \times$ NVIDIA RTX A6000 GPUs (48GB VRAM each). We also have access to a SLURM-managed server providing $2 \times$ NVIDIA L40S GPUs (48GB VRAM each).
- **Local Workstation:** Data preprocessing, quantitative evaluation, and visualization were performed on a local workstation featuring an NVIDIA RTX 4090 GPU (24GB VRAM) and an Intel Core i9 processor with 64GB of RAM.

Data Preparation and Training Duration. The computational timeline for the final models reported in the main text is as follows:

- **Offline Preprocessing:** The preparation of the training data, including TUDF preparation and the pre-extraction of VAE latents to accelerate training, was performed on the local workstation. This process required approximately 3–4 days for the complete dataset.
- **VAE Training:** The hybrid 3D VAE required approximately 3 days of training using 2 GPUs.
- **DiT Training:** The latent Diffusion Transformer (DiT) required approximately 2.5 days using 2 GPUs.
- **Ablation Studies:** The dense VAE ablation required approximately 6 days of compute, while DiT architecture ablations mirrored the 2.5-day schedule.

Total Compute Expenditure. Accounting for preliminary investigations and hyperparameter tuning, we estimate a total compute expenditure of approximately 1,680 GPU hours (2 GPUs over 35 days). This excludes the additional CPU-intensive preprocessing and evaluation time on the local workstation.

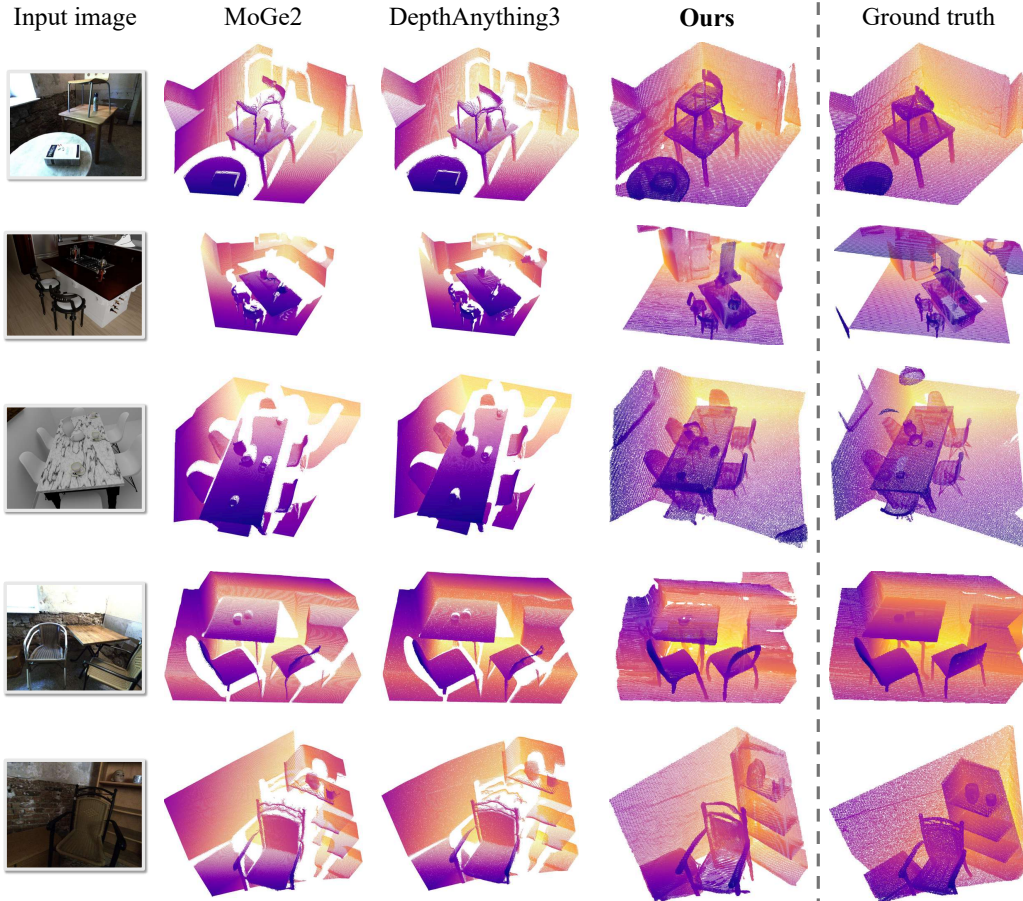


Figure 8: **Qualitative comparison with pixel-aligned approaches.** Unlike MoGe2 [79] and DepthAnything3, which are restricted to visible surfaces and leave significant holes, VolFill reconstructs complete, physically plausible amodal geometry.

C Broader societal impacts.

Our work advances single-image 3D scene reconstruction, which may benefit applications in robotics, assistive navigation, AR/VR, digital twins, architectural design, and embodied AI by enabling richer spatial understanding from limited visual input. In particular, complete-scene reconstruction could help autonomous systems reason about occluded geometry and improve safety in indoor navigation or human-robot interaction. However, the same capability may also raise concerns if deployed in privacy-sensitive settings, since inferred 3D layouts could reveal information about private homes, workplaces, or personal environments beyond what is directly visible in an image. The method could also be misused for surveillance, unauthorized mapping, or synthetic reconstruction of restricted spaces. In addition, models trained on datasets such as 3D-FRONT and ScanNet++ may inherit dataset biases toward particular indoor layouts, object categories, geographic regions, or socioeconomic settings, potentially limiting performance in underrepresented environments. Thus, we view this work as a research contribution rather than a deployment-ready system, and recommend that practical use include consent-aware data collection and privacy safeguards.




Figure 9: **Point cloud and mesh comparison.** Our method produces cleaner point clouds and significantly more coherent meshes than LaRI and NOVA3R.




Figure 10: **Qualitative results** of the hybrid 3D VAE.




Figure 11: **Evolution of amodal geometry.** Step-wise visualization of the denoising process. Our model rapidly converges on the coarse scene layout by $t = 3$ and recovers sharp, detailed structures around $t = 8$, followed by continuous refinement.