

GENA3D: Generative Amodal 3D Modeling by Bridging 2D Priors and 3D Coherence

Junwei Zhou[✉] and Yu-Wing Tai[✉]

Dartmouth College
Hanover, NH, USA

<https://colezwhy.github.io/gena3d/>

Abstract. Generating complete 3D objects under partial occlusions (*i.e.*, amodal scenarios) is a practically important yet challenging problem, as large portions of object geometry are unobserved in real-world scenarios. Existing approaches either operate directly in 3D, which ensures geometric consistency but often lacks generative expressiveness, or rely on 2D amodal completion, which provides strong appearance priors but does not guarantee reliable 3D structure. This raises a key question: *how can we achieve both generative plausibility and geometric coherence in amodal 3D modeling?* To answer this question, we introduce GENA3D (**GEN**erative **A**modal 3D), a framework that integrates learned 2D generative priors with explicit 3D geometric reasoning within a conditional 3D generation paradigm. The 2D priors enable the model to plausibly infer diverse occluded content, while the 3D representation enforces multi-view consistency and spatial validity. Our design incorporates a novel View-Wise Cross-Attention for multi-view alignment and a Stereo-Conditioned Cross-Attention to anchor generative predictions in 3D relationships. By combining generative imagination with structural constraints, GENA3D generates complete and coherent 3D objects from limited observations without sacrificing geometric fidelity. Experiments demonstrate that our method outperforms existing approaches in both synthetic and real-world amodal scenarios, highlighting the effectiveness of bridging 2D priors and 3D coherence in generating plausible and geometrically consistent 3D structures in complex environments.

Keywords: Amodal 3D Object Generation · Generative Models

1 Introduction

Generating complete 3D objects under partial occlusions is a fundamental but challenging task in computer vision. In real-world scenarios, objects are often captured from arbitrary viewpoints, with certain portions hidden by surrounding structures or other objects. Such incomplete observations make the problem of generating 3D objects in the scene inherently amodal: a system must infer the geometry and appearance of regions that are never directly observed in the input images. This ability is crucial for object-level 3D scene understanding and




Fig. 1: Comparison of 3D object reconstruction/generation from sparse and partial occluded inputs. From left to right: input views, results from (a) MVS-based reconstruction (*e.g.*, VGGT [61]), (b) 2D amodal completion lifted to 3D (TRELLIS [68]), and (c) **GENA3D (ours)**. Our model faithfully generates complete, occlusion-free 3D geometry consistent with input partially visible constraints.

downstream applications that are critically dependent on such capability, such as robotic interaction, AR/VR content creation, and embodied AI perception.

Recent advances in 3D object generation have been driven by diffusion-based and autoregressive models that learn powerful structural and semantic priors from large-scale 3D datasets. These approaches can synthesize diverse and realistic shapes from text prompts, single images, or even sparse inputs. However, most existing 3D generative models heavily depend on learned strong geometric priors and implicitly assume high-quality observations and canonical object poses to produce high-quality results. When confronted with sparse and partially occluded real-world inputs, these models often struggle to effectively leverage the limited visible evidence to plausibly infer the missing regions. The inherent complexity of the 3D domain and the scarcity of corresponding training data further exacerbate this issue, leading to inconsistent geometry, implausible structures, loss of fine-grained details, or even catastrophic generation failures.

To address this challenge, we introduce **GENA3D**, a generative framework for amodal 3D object generation from sparse, unposed views. Our approach treats this problem as a conditional generation process that bridges 2D priors and 3D coherence. Specifically, we first perform 2D object-centric amodal completion that are rich in generative ability on each view to recover occluded regions. The features of completed but inconsistent views are then aggregated through a *View-Wise Cross Attention* module that produces model latents fused from sparse-view information. To ensure the fidelity of visible regions and 3D coherence, we condition the generation on partial geometry (point clouds) estimated from an MVS model. This information is incorporated through a *Stereo-Conditioned Cross Attention* module equipped with a geometry-guided gating mechanism,

which modulates fused latent based on 3D spatial cues. After processing the generated sparse structure through another pretrained Flow Transformer, the resulting structured latent representation is then decoded to produce a complete, occlusion-free 3D model that faithfully respects to the observed geometry while plausibly generating the hidden regions. As shown in Fig. 1 (c), our method produces coherent and complete 3D objects even when large portions of the geometry are never observed or the 2D completions are not consistent.

Extensive experiments on both synthetic and real-world datasets demonstrate that **GENA3D** achieves superior generation quality compared with existing sparse-view and amodal 3D generation baselines. Our method improves geometric faithfulness while maintaining realistic appearance, achieving consistent and superior performance from single-view to multi-view inputs. These results highlight the potential of generative modeling for practical 3D object-level scene synthesis under real-world occlusion and view sparsity.

The contribution of this work is threefold:

- We introduce **GENA3D**, a generative amodal 3D modeling framework that recovers complete geometry and appearance from sparse and occluded inputs by bridging 2D priors and 3D coherence during generation.
- We propose two principled conditioning mechanisms for amodal 3D generation: (1) *View-Wise Cross Attention*, which introduces multi-view reasoning to mitigate view dominance and geometric drift; and (2) *Stereo-Conditioned Cross Attention*, which integrates partial stereo geometry through logit-level gating, transforming geometric cues into explicit attention regulators for structurally coherent generation.
- Comprehensive experiments validate that GENA3D achieves superior generation fidelity, completeness, and consistency compared with existing sparse-view and amodal generation methods.

2 Related Work

3D Generative Models. Recent advances in 3D generative models have opened new possibilities for structure-aware reconstruction. Methods leveraging Gaussian Splatting [28] and NeRF [45] as these representations excel at novel view synthesis [59, 76], but struggle to segregate object-level entities, hindering object-level 3D scene understanding. Diffusion-based generative models further enhance visual fidelity, controllability, and scalability by leveraging pretrained 2D diffusion priors [9, 13, 24] or 3D-aware latent representations [35, 41, 58, 68, 83]. These methods typically rely on massive optimization iterations or decently posed input images for proper results. Some works attempt to mitigate missing regions using inpainting-augmented generation [11, 20, 84, 85]. Other scene-level generative pipelines [10, 25, 39, 44, 75] synthesize entire 3D scenes from a single image, but they rely on carefully curated data and pose-calibrated views, which limits their effectiveness under real-world, unposed scenarios.

In contrast to prior approaches, GENA3D is a generative 3D reconstruction model tailored for sparse, unposed inputs. By fusing generative priors with

partial geometric conditioning, it produces reconstructions that are both geometrically faithful and semantically coherent, remaining consistent with sparse-view constraints while plausibly inferring unseen object regions.

Amodal Completion in 2D/3D. Amodal completion aims to infer the hidden or occluded parts of objects beyond visible evidence. Early works address this through explicit amodal segmentation and completion tasks [32, 34, 37, 46, 50], focusing primarily on 2D perception and annotation-based benchmarks. These methods rely on supervised mask prediction or latent shape priors to hallucinate occluded object regions, but lack in fidelity and geometric consistency.

Generative 2D inpainting models [54, 78] have since become common solutions in image-level amodal completion, synthesizing realistic missing content based on surrounding context. Nowadays, researchers further enhance this capacity by introducing specialized amodal reasoning modules [1, 16, 30] or by fine-tuning diffusion models for occlusion-aware reconstruction [43, 49, 71]. Extending beyond single-view imagery, multi-view inpainting models [3, 65] attempt to maintain cross-view coherence but often assume calibrated cameras or dense view coverage. More geometry-aware approaches [17, 33, 55, 57, 79] leverage 3D representation priors for scene-level completion; however, they remain limited in handling object-level amodal generation due to alignment errors and weak semantic reasoning. The most relevant work is Amodal3R [67], which trains a native 3D generation model from occluded objects, but it is constrained by data scale and limited occlusion reasoning ability.

In sparse and unposed scenarios, 2D inpainting models can recover locally plausible structures but often introduce cross-view inconsistencies that lead to geometric artifacts in 3D reconstruction. Our method integrates strong 2D amodal priors into a 3D generative framework, enforcing global geometric coherence via view-wise and stereo-conditioned attention. Compared with single-image 3D amodal methods, GENA3D enables more robust reconstruction under severe occlusion by aggregating multi-view information and enhancing cross-view geometric reasoning, while preserving realistic image-level completion.

Multi-View Stereo (MVS). Multi-view stereo (MVS) [19, 27, 31, 60, 62, 63, 72, 73, 77, 80, 82] has been studied for decades as a core technique for 3D reconstruction from multiple input views. Recent advances [52, 56, 66] in *dense correspondence learning* have enhanced MVS accuracy and robustness by jointly reasoning about geometry and radiance. These methods achieve fine-grained alignment but remain limited when it comes to calibrated settings and fixed depth ranges. To improve scalability and avoid per-scene optimization, *feed-forward extensions* [7, 14, 26, 40] introduce transformer-based architectures and specific techniques for zero-shot multi-view reconstruction. These designs collectively enhance scalability and inference efficiency.

Despite substantial progress, MVS methods are primarily designed for *reconstruction* rather than *generation*. In contrast, we do not seek to improve MVS. Instead, we leverage partial MVS geometry as an explicit structural prior within a generative 3D framework. By conditioning on stereo-derived point clouds through our Stereo-Conditioned Cross Attention, GENA3D converts incomplete




Fig. 2: Overview of GENA3D. Given sparse images S , visibility masks M_{vis} , and occlusion masks M_{occ} indicating the occluded object O , GENA3D first generates a sparse structure (stage 1) by aggregating multi-view information in Sec. 3.1 and infers the complete geometric structure from the partial stereo point cloud $\mathbf{P}_O$ in Sec. 3.2. We then employ a pretrained amodal SLAT Transformer to generate details with M_{vis} and M_{occ} , and then generate the structured latent. This structured latent is decoded into an occlusion-free 3D object with high-quality geometry and appearance.

geometric cues into attention regulators that guide plausible and geometrically coherent completion. This effectively bridges feed-forward MVS geometry with conditional 3D generation for object-level amodal modeling.

3 Method

For an arbitrary object O in the given K sparse views $S = \{s^1, s^2, \dots, s^K\}$ (typically 1-4 views), we can obtain the visibility masks of O and the corresponding text prompt automatically or manually using vision foundation models, *e.g.*, SAM [29] and Florence 2 [69]. After that, we follow the occlusion analysis and 2D amodal completion strategies from [1] with minor modifications (using the above obtained visibility mask and prompt) to first perform generative image-level amodal completion across all views using [1] (denoted as OAAC). Altogether, we can obtain visual clues of the desired object O at all K input views: the visibility masks $M_{vis} = \{m_v^1, m_v^2, \dots, m_v^K\}$ indicating the visible part of O ; the inconsistent 2D completions of O : $I = \{i^1, i^2, \dots, i^K\}$; and the occlusion masks $M_{occ} = \{m_o^1, m_o^2, \dots, m_o^K\}$ representing the potential occlusion regions.

We illustrate the overall model structure in Fig. 2. Our design prioritizes the Sparse Structure Transformer, *i.e.*, stage 1. Given object-centric sparse views I processed using 2D amodal completion methods, we apply the proposed **View-Wise Cross Attention** (Sec. 3.1) to aggregate multi-view features and fuse their latents. The fused representation is further refined through the **Stereo-Conditioned Cross Attention** (Sec. 3.2) for unobserved structure inference, guided by observed partial stereo geometry and a geometry-aware gating mechanism. The resulting Sparse Structure is decoded and then passed to the pretrained Structure Latent Transformer from [67] for Structured Latent generation. Implementation details and training data preparation are described in Sec. 3.3.




Fig. 3: The original model shows a strong bias and results in artifact accumulation when conditioned with multi-view.




Fig. 4: A detailed illustration of our proposed View-Wise Cross Attention and Stereo-Conditioned Cross Attention modules.

3.1 View-Wise Cross Attention

Existing multi-view conditioning strategies in generative models typically adopt either sequential cross-attention (conditioning on one view per sampling step) or naïve feature concatenation across views. As shown in Fig. 3, under sparse and unposed amodal settings, these strategies introduce two critical issues: (1) view dominance, where later-conditioned views overwrite earlier structural hypotheses, and (2) geometric drift, where inconsistencies in independently completed 2D views accumulate into unstable 3D structure.

To address this, we propose View-Wise Cross Attention, a parallelized multi-view conditioning mechanism that enforces simultaneous and balanced cross-view reasoning at every generation step. Since the base 3D generation model consistently produces objects (latents) in canonical poses regardless of the conditions’ camera poses, we exploit this pose-invariant property as a prior in designing our multi-view fusion strategy. We denote the intermediate latent as $z \in \mathbb{R}^{L \times D}$, where L is the number of tokens and D is the hidden dimension of the tokens. In the standard Cross Attention (denoted as $\mathcal{CA}$), the latent z interacts with the encoded image feature c ,

$$c = \text{DINO}(i), \quad z' = \mathcal{CA}(z, c). \quad (1)$$

Here, DINOv2 [48] serves as the image encoder, extracting features for input image i . When multiple images are available, for each 2D completion view i^n , we extract its DINO feature c^n , forming a set of encoded features $C = [c^1, c^2, \dots, c^K]$. We then apply Cross Attention in parallel for each view:

$$z'_n = \mathcal{CA}(z, c^n), \quad \forall n \in 1, 2, \dots, K. \quad (2)$$

The resulting view-wise latents $Z' = [z'_1, z'_2, \dots, z'_K]$ are fused into a single latent z' through a visibility-weighted average:

$$z' = \frac{1}{K} \sum_{n=1}^K w_n z'_n, \quad w_n = \frac{\tau_n}{\sum_{j=1}^K \tau_j}, \quad \tau_n = \frac{m_v^n}{m_o^n + m_v^n}. \quad (3)$$

τ_n represents the visibility ratio, defined as the proportion of possibly visible region in view i^n . A final LayerNorm is applied after fusion for numerical stability.

By structurally decoupling per-view reasoning and explicitly rebalancing contributions using geometric visibility priors, View-Wise Cross Attention converts inconsistent 2D amodal completions into a coherent shared latent space, significantly improving cross-view stability under sparse observations.

3.2 Stereo-Conditioned Cross Attention Module

While prior generative models may incorporate geometric signals (e.g., depth maps or point clouds) as additional tokens, they typically treat geometry as passive conditioning features. This often results in weak structural enforcement, particularly when the geometry is partial, noisy, or unposed. To address this, we introduce Stereo-Conditioned Cross Attention, a geometry-aware attention mechanism that actively modulates the attention distribution using encoded stereo structure. Instead of merely concatenating geometric tokens, Stereo-Conditioned Cross Attention injects geometric priors directly into the attention logits via a geometry-guided gating function.

Given sparse observations $S = \{s^1, s^2, \dots, s^K\}$, we obtain the stereo point cloud $\mathbf{P} = \{p^1, p^2, \dots, p^K\} = \text{MVS}(S)$, where each p^n corresponds to a point cloud reconstructed from view s^n , by lifting each pixel to the 3D space using the depth value. We extract the partial point cloud of the target object by selecting points projected within the visibility masks:

$$\mathbf{P}_O = \bigcup_{n=1}^K p_v^n, \quad p_v^n = \{p \in p^n \mid m_v^n(p) = 1\}. \quad (4)$$

The resulting $\mathbf{P}_O$ is voxelized into active voxels $\mathbf{V}_O$ with grid size φ , then encoded by the Sparse Structure VAE, and finally transformed into geometric feature tokens using the same patchify and linear layer for input transformation:

$$\begin{aligned} \mathbf{V}_O &= \text{Voxelize}(\mathbf{P}_O, \varphi), \quad f_{\mathbf{V}_O} = \text{Enc}(\mathbf{V}_O), \\ c_{\text{geo}} &= \text{Linear}(\text{Patchify}(f_{\mathbf{V}_O})). \end{aligned} \quad (5)$$

We further introduce a **Geometry-Guided Gating MLP** that modulates attention using the encoded geometry:

$$\begin{aligned} g &= \sigma(\text{MLP}_{\text{gating}}(c_{\text{geo}})), \\ \alpha &= \text{softmax}\left(\frac{QK^\top}{\sqrt{D}} + \log(g + \varepsilon)\right), \end{aligned} \quad (6)$$

where α denotes the attention weights computed from query $Q = W_q z'$ and key $K = W_k c_{\text{geo}}$, and ε is a small constant for logarithmic numerical stability.

By embedding 3D structure directly into the attention computation, Stereo-Conditioned Cross Attention establishes a tighter coupling between generative imagination and geometric plausibility, enforcing structural coherence even when large object regions remain unobserved.

3.3 Data Engine & Training/Inference Proxy

Data Engine. We leverage object-centric 3D datasets from [68] and convert them into watertight meshes using Manifold [23]. To simulate realistic occlusions, we generate 3D-consistent occlusion masks by randomly selecting a seed face on each mesh and iteratively expanding to its neighboring faces. This process continues outward until reaching a randomly sampled coverage ratio, producing spatially coherent and natural occlusion patterns. We then render object views and their corresponding occlusion masks, discarding camera poses with extreme angles or insufficient occlusion. Finally, completed object views are synthesized by treating the occlusion masks as inpainting regions and applying controlled mask distortions for greater diversity (see supplementary for more details).

Training Proxy. During training, the number of sparse views per object is randomly sampled from $\{1, 2, 3, 4\}$ for each batch. After sampling, stereo point clouds are generated online based on the selected views. The resulting partial point clouds are centralized and normalized to ensure training stability. Thanks to the synthesized occluded object inputs and visibility masks, modern multi-view stereo models can still recover plausible partial geometry, providing reliable geometric supervision throughout training.

We optimize the model using the standard conditional flow matching (CFM) [38] objective during training:

$$\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t, z_0, \epsilon} \|v_\theta(z, t) - (\epsilon - z_0)\|_2^2, \quad (7)$$

where θ denotes the network parameters, z_0 and ϵ are data and noise samples, t is the timestep, and $z = (1 - t)x_0 + t\epsilon$ represents the interpolated state.

Inference Proxy. During inference, our method can process both scene-level and object-level images. For scene-level images, we first generate the stereo point clouds for the entire scene with S , and then extract the partial point clouds for each object using the corresponding visibility masks M_{vis} . This approach reduces computational overhead when generating multiple objects within a single scene. All other procedures remain consistent with what is done in the training process. Fig. 2 shows an example of scene-level inference.

4 Experiments

4.1 Experimental Settings

Dataset. We use 3D-FUTURE [18] and Amazon Berkeley Objects (ABO) [12] for 3D consistent occluded object simulation. 3D-FUTURE contains 9472 objects, and ABO contains 4485. We train our network with all objects combined.




Fig. 5: Amodal 3D object generation results on the GSO dataset. We give examples generated with 1, 2, and 4 views, respectively. Note that for FreeSplatter [70], we adopt the Hunyuan-MV [74] to generate multiple views under the input of single image.

For each object, we render 150 views while filtering out camera elevation angles lower than -75 degrees or higher than 75 degrees, which results in around 90 views per object that are close to real-world scenarios and captures. We select views whose occluded regions take up from 20% to 60% during training. For evaluation, we use the Google Scanned Object (GSO) [15] dataset for object-level test. In addition, we sample from Hypersim [53] dataset to simulate *in-the-scene* scenarios; COCO [36] and Mip-NeRF 360 [4] dataset to evaluate our *in-the-wild* performance in Sec. 4.4.

Implementation Details. We make the entire sparse structure generation stage (stage 1) of the model trainable and use 8 RTX 6000 GPUs in training. We use π^3 [63] as the multi-view stereo model and normalize the coordinates of partial point clouds. Then we perform voxelization with a grid size of 64^3 . We adopt the AdamW [42] optimizer with a learning rate of $5e-5$, and Classifier-Free-Guidance [22] with a drop rate of 0.1 during training. We train the model for 12 epochs with a global batch size of 32, which takes approximately 16 hours to complete. Please refer to the supplementary for more details in the training, inference, and data engine settings.

4.2 Main Results

Comparison Methods We compare our method with various 3D object generation methods equipped with the latest 2D completion methods. First, to validate the 3D amodal completion ability with unposed sparse views of GENA3D, we employ TRELIS [68], Amodal3R [67], FreeSplatter [70], and the most recent SAM3D [8] (implemented MV-SAM3D for sparse-view scenarios) for comparison. We utilize [1] to perform 2D amodal completion at sparse views. Second,




Fig. 6: Additional qualitative results, together with occluded input images after 2D amodal completion.




Fig. 7: We demonstrate that GENA3D can resist a certain degree of errors in the MVS reconstructed point clouds, *e.g.*, misalignment and overlapping geometry from different views.




Fig. 8: Visual comparisons from the ablation studies showing the impact of each proposed model design.

to demonstrate that GENA3D is also effective and can generate high-quality 3D assets in real-world scenarios, we show *in-the-wild* and *in-the-scene* qualitative results, compared with Amodal3R [67] and SceneGen [44], respectively.

Quantitative Results. We present quantitative results tested on the GSO dataset. We first render 3D consistent masks and then follow the inference proxy (Sec. 3.3) to generate all 1030 objects in the dataset. For a fair comparison, we select those views whose visible parts occupy from 30% to 70% of the whole rendering. To validate the effectiveness and generalizability of our GENA3D, we compare our method with others on different numbers of sparse view settings, *i.e.*, 1, 2, and 4 views. We note that SAM3D is evaluated solely under the single-view setting, consistent with its original design configuration and the official implementation. As shown in Tab. 1, we employ Fréchet Inception Distance (FID) [21], Kernel Inception Distance (KID) [5] to evaluate the quality of rendered images; CLIP-Score [51] to measure the feature similarity; Minimum

Table 1: Quantitative results on amodal 3D object generation, validated on the GSO dataset. We compare under the settings of 1, 2, and 4 views. The best and second-best results are **bold** and underlined, respectively. For sparse-view settings, we either apply 2D amodal completion to each view using [1] (OAAC), or perform 2D amodal completion on a single view followed by multi-view diffusion generation (OAAC + MV)

Method	View Num.	2D Amodal	FID ↓	KID(%) ↓	CLIP(%) ↑	MMD(‰) ↓	COV(%) ↑	PCS(%) ↑
TRELLIS [68]	1 view	OAAC	49.68	3.21	76.79	6.08	33.37	21.1
FreeSplat [†] [70]		OAAC + MV	72.88	4.68	78.14	7.24	35.78	25.5
Amodal3R [67]		N/A	<u>39.05</u>	0.44	<u>80.06</u>	<u>5.52</u>	38.64	46.3
SAM3D [8]		N/A	34.68	<u>0.46</u>	82.16	5.50	<u>39.03</u>	<u>46.5</u>
Ours		OAAC	33.91	<u>0.46</u>	82.23	5.68	39.19	61.9
TRELLIS [68]	2 views	OAAC	46.23	3.28	77.01	6.27	34.12	25.0
FreeSplat [†] [70]		OAAC	91.32	6.77	74.24	10.44	27.74	12.4
Amodal3R [67]		N/A	<u>35.52</u>	0.44	<u>81.84</u>	<u>5.52</u>	<u>38.92</u>	<u>47.1</u>
Ours		OAAC	32.12	<u>0.45</u>	82.33	5.50	39.08	66.3
TRELLIS [68]	4 views	OAAC	45.41	2.79	78.17	5.94	34.42	23.2
TRELLIS [‡] [68]		OAAC + MV	44.85	1.94	76.24	6.71	34.24	28.6
FreeSplat [†] [70]		OAAC	94.10	4.64	75.58	10.48	32.31	8.4
Amodal3R [67]		N/A	<u>35.15</u>	0.43	<u>82.04</u>	<u>5.51</u>	<u>38.83</u>	<u>46.2</u>
Ours		OAAC	30.73	0.43	82.53	5.48	39.48	67.0

Table 2: The generation faithfulness to the observed visual content (regions within visibility masks M_{vis}). We evaluate with 1, 2, and 4 input views.

Methods	Num. of Views	SSIM↑	PSNR↑	LPIPS↓
TRELLIS [68]	1 view	0.721	14.15	0.325
SAM3D [8]		0.797	15.54	0.226
Amodal3R [67]		0.814	16.41	0.241
Ours		0.832	16.52	0.231
TRELLIS [68]	2 views	0.686	13.98	0.346
Amodal3R [67]		0.819	16.52	0.238
Ours		0.837	16.99	0.212
TRELLIS [68]	4 views	0.674	13.66	0.357
Amodal3R [67]		0.832	16.64	0.236
Ours		0.838	16.98	0.208

Matching Distance (MMD) and Coverage (COV) [47] to evaluate the geometry quality. Moreover, we evaluate the Perceptual Coherence Score (denoted as PCS) of the generated 3D objects using a VLM [2]. See supplementary material for detailed explanations of applied metrics.

Qualitative Results. We present qualitative results on the GSO dataset with simulated occlusions for amodal 3D object generation. As shown in Fig. 5, our model outperforms other methods in rendering quality and geometry consistency. Moreover, by bridging 2D priors and 3D coherence, GENA3D shows better potential in uncommon and complex amodal object completion. We showcase more visualization results in Fig. 6, together with the 2D amodal completion results to demonstrate GENA3D’s ability in fixing image-level structural inconsistencies in the 3D space.

Table 3: Ablation study on different proposed modules. Red indicators represent performance **drops** compared to the full model.

Model Structure	FID ↓	MMD(‰) ↓	COV(%)↑
Ours full	32.12	5.50	39.08
- Gating MLP	32.87 ^{+0.55}	5.58 ^{+0.08}	37.68 ^{-1.40}
- View-Wise CA	38.64 ^{+6.42}	5.76 ^{+0.26}	38.34 ^{-0.74}
- Stereo-Cond. CA	42.92 ^{+10.80}	5.91 ^{+0.41}	35.26 ^{-3.82}
- all proposed	46.23 ^{+14.11}	5.94 ^{+0.44}	34.42 ^{-4.66}

Table 4: Ablation study on GENA3D’s ability with number of input views, we test up to 20 views. The red/green indicator represents performance gain/loss, respectively.

Num. of Views	FID ↓	CLIP-Score(%) ↑	MMD(‰) ↓
1	33.91	82.23	5.52
2	32.12 ^{-1.79}	82.33 ^{+0.10}	5.50 ^{-0.02}
4	30.73 ^{-1.39}	82.53 ^{+0.20}	5.48 ^{-0.02}
10	29.56 ^{-1.17}	82.98 ^{+0.45}	5.45 ^{-0.03}
20	29.43 ^{-0.13}	82.95 ^{-0.03}	5.45 ^{0.00}

Generation Faithfulness. We also compute the average of visibility-masked SSIM [64], LPIPS [81], and PSNR over observable regions at their corresponding views, focusing on the generation fidelity to the visible parts of sparse inputs in amodal scenarios. As shown in Tab. 2, compared to previous SoTA methods and baselines, our method achieve superior performance in terms of preserving observed visual details within the generated 3D objects.

4.3 Ablation Study

Structure Design. We conduct ablation studies to assess the effectiveness of different module designs. Tab. 3 reports quantitative results, and Fig. 8 shows visual comparisons. With *View-Wise Cross Attention*, our method performs notably better under sparse views, demonstrating better rendering quality with more reasonable sparse structure as prior for texture generation. Moreover, incorporating *Stereo-Conditioned Cross Attention* for explicit geometry guidance yields large benefits in the geometric metrics (MMD). As shown in Fig. 8, MVS conditioning helps GENA3D produce more faithful and structurally consistent geometry, even when the 2D completions are inconsistent or contain artifacts.

Number of Views. We investigate how the number of input views influence our method. As shown in Tab. 4, the provided metrics remain stable or exhibit slight improvements, indicating that our method can handle semi-dense settings with more input images, which demonstrates the adaptability of our work.

2D Amodal Completion. We conduct experiments on how our model performs with different 2D amodal completion methods. As shown in Tab. 5, we can see that for different methods, our GENA3D is able to produce consistent and robust results that are comparable in numerical metrics. Additionally, it is intuitive that stronger completion methods produce better results, which shows the potential and adaptability of our GENA3D with more advanced techniques in the future.




Fig. 9: In-the-wild 3D amodal generation results. Single-view examples are from the COCO [36] dataset; sparse-view examples are from the Mip-NeRF 360 [4].

Table 5: Ablation study on different 2D Amodal Completion methods. We present results generated with 2 views.

2D Amodal Comp.	FID ↓	CLIP-Score(%) ↑	COV(%) ↓
pix2gestalt [49]	38.61	82.07	37.26
Flux inpainting [6]	33.07	84.26	39.49
OAAC [1]	32.12	82.33	39.08

Errors in MVS. In recent MVS models, the reconstructed point clouds occasionally exhibit misaligned or failing structures. Therefore, we analyze how such MVS errors affect our generation pipeline. As shown in Fig. 7, when the reconstructed stereo is geometrically implausible, our model can be less influenced by the stereo conditioning and perform self-correction with the 3D generative prior.

4.4 Generation In-the-wild & In-the-Scene

We further evaluate our model in both *in-the-wild* and *in-the-scene* settings to assess its robustness and generalization ability. The *in-the-wild* cases correspond to real-world environments with complex interactions between objects, while the *in-the-scene* setting mainly involves structured indoor layouts, such as furniture arrangements or tabletop scenes. As shown in Fig. 9, our GENA3D achieves realistic 3D amodal generation across various *in-the-wild* conditions. In Fig. 10, we show *in-the-scene* generation results compared with [67] and [44], indicating the potential of our model in object-level scene generation under the sparse-view and severe occlusion setting. These results demonstrate that with strong 2D




Fig. 10: In-the-scene 3D amodal generation results. Input images are from the Hyper-sim [53] dataset. Red and Green icons indicate different target objects.

priors and 3D coherence, our method can generalize across different domains, retain visible partial information, and faithfully hallucinate unseen parts.

5 Conclusion

In this paper, we propose GENA3D, a 3D object generator capable of generating amodal 3D objects from arbitrary sparse and partial occluded observations. Using only object-level occlusion training data and simulated 2D amodal completion, GENA3D can seamlessly integrate with 2D amodal completion methods and multi-view stereo models, enabling generalization across both synthetic and real-world scenarios and paving the way for high-quality 3D generation in diverse applications. While GENA3D relies on priors learned from generative models to infer missing regions, which may look plausible but not always match the ground-truth geometry, the approach still represents a meaningful step forward in 3D generation, demonstrating the potential of bridging strong 2D generative priors and 3D geometry coherence under real-world amodal scenarios.

References

1. Ao, J., Jiang, Y., Ke, Q., Ehinger, K.A.: Open-world amodal appearance completion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6490–6499 (2025)
2. Bai, S., Cai, Y., Chen, R., Chen, K., Chen, X., Cheng, Z., Deng, L., Ding, W., Gao, C., Ge, C., et al.: Qwen3-vl technical report. arXiv preprint arXiv:2511.21631 (2025)
3. Barda, A., Gadelha, M., Kim, V.G., Aigerman, N., Bermano, A.H., Groueix, T.: Instant3dit: Multiview inpainting for fast editing of 3d objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16273–16282 (2025)
4. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mipnerf 360: Unbounded anti-aliased neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
5. Bińkowski, M., Sutherland, D.J., Arbel, M., Gretton, A.: Demystifying mmd gans. In: International Conference on Learning Representations (ICLR) (2018)
6. Black Forest Labs: Flux.1 kontext: Flow matching for in-context image generation and editing in latent space. arXiv preprint arXiv:2506.15742 (2025)
7. Cabon, Y., Stoffl, L., Antsfeld, L., Csurka, G., Chidlovskii, B., Revaud, J., Leroy, V.: Must3r: Multi-view network for stereo 3d reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2025)
8. Chen, X., Chu, F.J., Gleize, P., Liang, K.J., Sax, A., Tang, H., Wang, W., Guo, M., Hardin, T., Li, X., et al.: Sam 3d: 3dfy anything in images. arXiv preprint arXiv:2511.16624 (2025)
9. Chen, Y., Pan, Y., Yang, H., Yao, T., Mei, T.: Vp3d: Unleashing 2d visual prompt for text-to-3d generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2024)
10. Chen, Y., Ni, J., Jiang, N., Zhang, Y., Zhu, Y., Huang, S.: Single-view 3d scene reconstruction with high-fidelity shape and texture. In: International Conference on 3D Vision (3DV). pp. 1456–1467 (2024)
11. Chen, Y., Wang, T., Wu, T., Pan, X., Jia, K., Liu, Z.: Comboverse: Compositional 3d assets creation using spatially-aware diffusion guidance. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 128–146. Springer (2024)
12. Collins, J., Goel, S., Deng, K., Luthra, A., Xu, L., Gundogdu, E., Zhang, X., Vicente, T.F.Y., Dideriksen, T., Arora, H., 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 (CVPR). pp. 21126–21136 (2022)
13. Ding, L., Dong, S., Huang, Z., Wang, Z., Zhang, Y., Gong, K., Xu, D., Xue, T.: Text-to-3d generation with bidirectional diffusion using both 2d and 3d priors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2024)
14. Dong, Y., Li, Y., Huang, Z., Bian, W., Liu, J., Bao, H., Cui, Z., Li, H., Zhang, G.: A global depth-range-free multi-view stereo transformer network with pose embedding. In: Advances in Neural Information Processing Systems (NeurIPS) (2024)
15. Downs, L., Francis, A., Koenig, N., Kinman, B., Hickman, R., Reymann, K., McHugh, T.B., Vanhoucke, V.: Google scanned objects: A high-quality dataset

- of 3d scanned household items. In: International Conference on Robotics and Automation (ICRA). pp. 2553–2560 (2022)
16. Fan, H., Wang, L., Chen, H., Huang, Z., Wu, J., Sheng, L.: Multi-agent amodal completion: Direct synthesis with fine-grained semantic guidance. In: Proceedings of the 33rd ACM International Conference on Multimedia (2025)
 17. Feng, H., Liu, J., Tang, J., Wu, G., Chen, B., Lai, J., Wang, G.: Objfiller-3d: Consistent multi-view 3d inpainting via video diffusion models. arXiv preprint arXiv:2508.18271 (2025)
 18. Fu, H., Cai, B., Gao, L., Zhang, L.X., Wang, J., Li, C., Zeng, Q., Sun, C., Jia, R., Zhao, B., et al.: 3d-front: 3d furnished rooms with layouts and semantics. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10933–10942 (2021)
 19. Furukawa, Y., Hernández, C.: Multi-view stereo: A tutorial. *Foundations and Trends in Computer Graphics and Vision* **9**(1–2), 1–148 (2015)
 20. Han, H., Yang, R., Liao, H., Xing, J., Xu, Z., Yu, X., Zha, J., Li, X., Li, W.: Reparo: Compositional 3d assets generation with differentiable 3d layout alignment. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2025)
 21. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)
 22. Ho, J., Salimans, T.: Classifier-free diffusion guidance. In: NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications (2021)
 23. Huang, J., Su, H., Guibas, L.: Robust watertight manifold surface generation method for shapenet models. arXiv preprint arXiv:1802.01698 (2018)
 24. Huang, T., Zeng, Y., Zhang, Z., Xu, W., Xu, H., Xu, S., Lau, R.W.H., Zuo, W.: Dreamcontrol: Control-based text-to-3d generation with 3d self-prior. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2024)
 25. Huang, Z., Guo, Y.C., An, X., Yang, Y., Li, Y., Zou, Z.X., Liang, D., Liu, X., Cao, Y.P., Sheng, L.: Midi: Multi-instance diffusion for single image to 3d scene generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 23646–23657 (2025)
 26. Izquierdo, S., Sayed, M., Firman, M., Garcia-Hernando, G., Turmukhambetov, D., Civera, J., Aodha, O.M., Brostow, G., Watson, J.: Mvsanywhere: Zero-shot multi-view stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2025)
 27. Keetha, N., Müller, N., Schönberger, J., Porzi, L., Zhang, Y., Fischer, T., Knapitsch, A., Zauss, D., Weber, E., Antunes, N., Luiten, J., Lopez-Antequera, M., Bulò, S.R., Richardt, C., Ramanan, D., Scherer, S., Kotschieder, P.: Mapanything: Universal feed-forward metric 3d reconstruction. arXiv preprint arXiv:2509.13414 (2025)
 28. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4), 139–1 (2023)
 29. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 4015–4026 (2023)
 30. Lee, J.J., Benes, B., Yeh, R.A.: Tuning-free amodal segmentation via the occlusion-free bias of inpainting models. arXiv preprint arXiv:2503.18947 (2025)

31. Leroy, V., Cabon, Y., Revaud, J.: Grounding image matching in 3d with mast3r. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 71–91. Springer (2024)
32. Li, K., Malik, J.: Amodal instance segmentation. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 677–693 (2016)
33. Li, Z., Ye, W., Jiang, T., Huang, T.: 2d amodal instance segmentation guided by 3d shape prior. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 165–181 (2022)
34. Li, Z., Ye, W., Terven, J., Bennett, Z., Zheng, Y., Jiang, T., Huang, T.: Muva: A new large-scale benchmark for multi-view amodal instance segmentation in the shopping scenario. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 23504–23513 (October 2023)
35. Lin, J., Yang, X., Chen, M., Xu, Y., Yan, D., Wu, L., Xu, X., Xu, L., Zhang, S., Chen, Y.: Kiss3dgen: Repurposing image diffusion models for 3d asset generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2025)
36. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 740–755. Springer (2014)
37. Ling, H., Acuna, D., Kreis, K., Kim, S.W., Fidler, S.: Variational amodal object completion. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2020)
38. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: *International Conference on Learning Representations (ICLR)* (2023)
39. Liu, H., Zheng, Y., Chen, G., Cui, S., Han, X.: Towards high-fidelity single-view holistic reconstruction of indoor scenes. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2022)
40. Liu, T., Wang, G., Hu, S., Shen, L., Ye, X., Zang, Y., Cao, Z., Li, W., Liu, Z.: Mvsgaussian: Fast generalizable gaussian splatting reconstruction from multi-view stereo. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2024)
41. Long, X., Guo, Y.C., Lin, C., Liu, Y., Dou, Z., Liu, L., Ma, Y., Zhang, S.H., Habermann, M., Theobalt, C., et al.: Wonder3d: Single image to 3d using cross-domain diffusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 9970–9980 (2024)
42. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: *International Conference on Learning Representations (ICLR)* (2019)
43. Lu, R., Chen, Y., Liu, Y., Tang, J., Ni, J., Wan, D., Zeng, G., Huang, S.: Taco: Taming diffusion for in-the-wild video amodal completion. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 13638–13650 (2025)
44. Meng, Y., Wu, H., Zhang, Y., Xie, W.: Scenegen: Single-image 3d scene generation in one feedforward pass. *arXiv preprint arXiv:2508.15769* (2025)
45. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
46. Mohan, R., Valada, A.: Amodal panoptic segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 21023–21032 (June 2022)

47. Nichol, A., Jun, H., Dhariwal, P., Mishkin, P., Chen, M.: Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751* (2022)
48. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193* (2023)
49. Ozguroglu, E., Liu, R., Surá, D., Chen, D., Dave, A., Tokmakov, P., Vondrick, C.: pix2gestalt: Amodal segmentation by synthesizing wholes. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2024)
50. Qi, L., Jiang, L., Liu, S., Shen, X., Jia, J.: Amodal instance segmentation with kins dataset. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3014–3023 (2019)
51. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. *PmLR* (2021)
52. Ren, C., Xu, Q., Zhang, S., Yang, J.: Hierarchical prior mining for non-local multi-view stereo. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2023)
53. Roberts, M., Ramapuram, J., Ranjan, A., Kumar, A., Bautista, M.A., Paczan, N., Webb, R., Susskind, J.M.: Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 10912–10922 (2021)
54. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 10684–10695 (June 2022)
55. Salimi, A., Aumentado-Armstrong, T., Brubaker, M.A., Derpanis, K.G.: Geometry-aware diffusion models for multiview scene inpainting. *arXiv preprint arXiv:2502.13335* (2025)
56. Santo, H., Okura, F., Matsushita, Y.: Mvcpns-neus: Multi-view constrained photometric stereo for neural surface reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2024)
57. Shi, Z., Huo, D., Zhou, Y., Min, Y., Lu, J., Zuo, X.: Imfine: 3d inpainting via geometry-guided multi-view refinement. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 26694–26703 (2025)
58. Siddiqui, Y., Monnier, T., Kokkinos, F., Kariya, M., Kleiman, Y., Garreau, E., Gafni, O., Neverova, N., Vedaldi, A., Shapovalov, R., Novotny, D.: Meta 3d assetgen: Text-to-mesh generation with high-quality geometry, texture, and pbr materials. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2024)
59. Truong, P., Rakotosaona, M.J., Manhardt, F., Tombari, F.: Sparf: Neural radiance fields from sparse and noisy poses. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4190–4200 (2023)
60. Wang, F., Zhu, Q., Chang, D., Gao, Q., Han, J., Zhang, T., Hartley, R., Pollefeys, M.: Learning-based multi-view stereo: A survey. *arXiv preprint arXiv:2408.15235* (2024)
61. Wang, J., Chen, M., Karaev, N., Vedaldi, A., Rupperecht, C., Novotny, D.: Vggt: Visual geometry grounded transformer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5294–5306 (2025)

62. Wang, S., Leroy, V., Cabon, Y., Chidlovskii, B., Revaud, J.: Dust3r: Geometric 3d vision made easy. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20697–20709 (2024)
63. Wang, Y., Zhou, J., Zhu, H., Chang, W., Zhou, Y., Li, Z., Chen, J., Pang, J., Shen, C., He, T.: π^3 : Permutation-equivariant visual geometry learning. arXiv preprint arXiv:2507.13347 (2025)
64. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004)
65. Weber, E., Holynski, A., Jampani, V., Saxena, S., Snavely, N., Kar, A., Kanazawa, A.: Nerfiller: Completing scenes via generative 3d inpainting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
66. Wu, J., Li, R., Xu, H., Zhao, W., Zhu, Y., Sun, J., Zhang, Y.: Gomvs: Geometrically consistent cost aggregation for multi-view stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
67. Wu, T., Zheng, C., Guan, F., Vedaldi, A., Cham, T.J.: Amodal3r: Amodal 3d reconstruction from occluded 2d images. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2025)
68. Xiang, J., Lv, Z., Xu, S., Deng, Y., Wang, R., Zhang, B., Chen, D., Tong, X., Yang, J.: Structured 3d latents for scalable and versatile 3d generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 21469–21480 (2025)
69. Xiao, B., Wu, H., Xu, W., Dai, X., Hu, H., Lu, Y., Zeng, M., Liu, C., Yuan, L.: Florence-2: Advancing a unified representation for a variety of vision tasks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4818–4829 (2024)
70. Xu, J., Gao, S., Shan, Y.: Freesplatter: Pose-free gaussian splatting for sparse-view 3d reconstruction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 25442–25452 (2025)
71. Xu, K., Zhang, L., Shi, J.: Amodal completion via progressive mixed context diffusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9099–9109 (2024)
72. Yan, J., Wei, Z., Yi, H., Ding, M., Zhang, R., Chen, Y., Wang, G., Tai, Y.W.: Dense hybrid recurrent multi-view stereo net with dynamic consistency checking. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 674–689 (2020)
73. Yang, J., Sax, A., Liang, K.J., Henaff, M., Tang, H., Cao, A., Chai, J., Meier, F., Feiszli, M.: Fast3r: Towards 3d reconstruction of 1000+ images in one forward pass. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 21924–21935 (2025)
74. Yang, X., Shi, H., Zhang, B., Yang, F., Wang, J., Zhao, H., Liu, X., Wang, X., Lin, Q., Yu, J., et al.: Hunyuan3d 1.0: A unified framework for text-to-3d and image-to-3d generation. arXiv preprint arXiv:2411.02293 (2024)
75. Yao, K., Zhang, L., Yan, X., Zeng, Y., Zhang, Q., Xu, L., Yang, W., Gu, J., Yu, J.: Cast: Component-aligned 3d scene reconstruction from an rgb image. *ACM Transactions on Graphics (TOG)* **44**(4), 1–19 (2025)
76. Ye, B., Liu, S., Xu, H., Li, X., Pollefeys, M., Yang, M.H., Peng, S.: No pose, no problem: Surprisingly simple 3d gaussian splats from sparse unposed images. In: International Conference on Learning Representations (ICLR) (2025)

77. Yi, H., Wei, Z., Ding, M., Zhang, R., Chen, Y., Wang, G., Tai, Y.W.: Pyramid multi-view stereo net with self-adaptive view aggregation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 766–782 (2020)
78. Yu, T., Feng, R., Feng, R., Liu, J., Jin, X., Zeng, W., Chen, Z.: Inpaint anything: Segment anything meets image inpainting. arXiv preprint arXiv:2304.06790 (2023)
79. Zhan, G., Zheng, C., Xie, W., Zisserman, A.: Amodal ground truth and completion in the wild. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 28003–28013 (June 2024)
80. Zhang, J., Herrmann, C., Hur, J., Jampani, V., Darrell, T., Cole, F., Sun, D., Yang, M.H.: Monst3r: A simple approach for estimating geometry in the presence of motion. In: International Conference on Learning Representations (ICLR) (2025)
81. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 586–595 (2018)
82. Zhang, S., Wang, J., Xu, Y., Xue, N., Rupprecht, C., Zhou, X., Shen, Y., Wetzstein, G.: Flare: Feed-forward geometry, appearance and camera estimation from uncalibrated sparse views. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 21936–21947 (2025)
83. Zhao, Z., Lai, Z., Lin, Q., Zhao, Y., Liu, H., Yang, S., Feng, Y., Yang, M., Zhang, S., Yang, X., et al.: Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation. arXiv preprint arXiv:2501.12202 (2025)
84. Zhou, J., Li, X., Qi, L., Yang, M.H.: Coco4d: Comprehensive and complex 4d scene generation. arXiv preprint arXiv:2506.19798 (2025)
85. Zhou, J., Li, X., Qi, L., Yang, M.H.: Layout-your-3d: Controllable and precise 3d generation with 2d blueprint. In: International Conference on Learning Representations (ICLR) (2025)