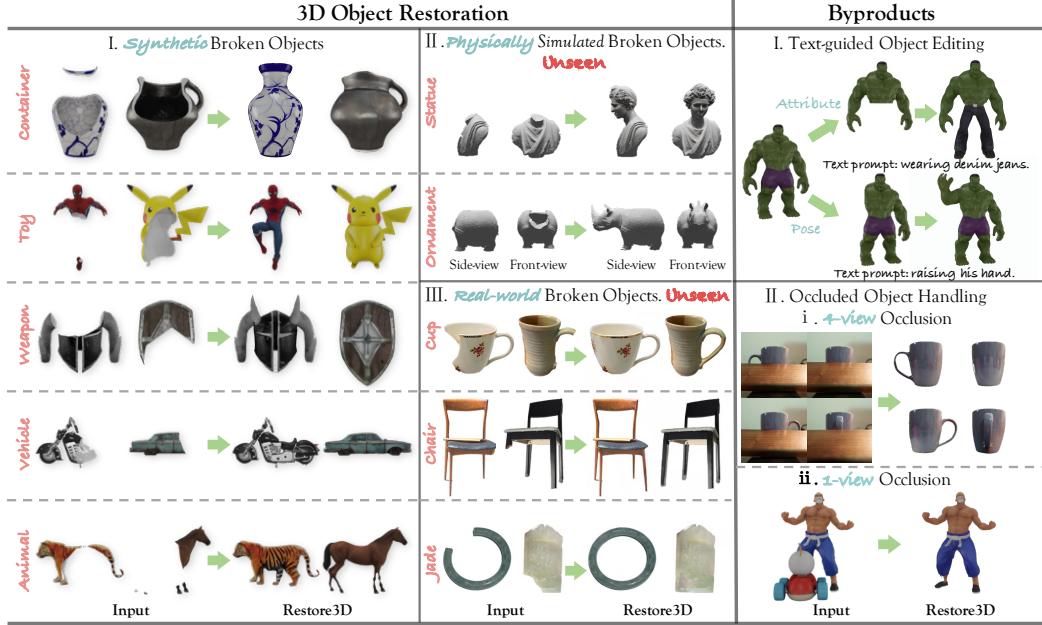


000 RESTORE3D: BREATHING LIFE INTO BROKEN 001 OBJECTS WITH SHAPE AND TEXTURE RESTORATION 002

003 **Anonymous authors**

004 Paper under double-blind review



027 Figure 1: Our **Restore3D** is among the first to simultaneously restore the shape and texture of
028 relatively complex and diverse objects, producing highly plausible and realistic results.
029

030 ABSTRACT

031 Restoring incomplete or damaged 3D objects is crucial for cultural heritage preser-
032 vation, occluded object reconstruction, and artistic design. Existing methods pri-
033 marily focus on geometric completion, often neglecting texture restoration and
034 struggling with relatively complex and diverse objects. We introduce Restore3D,
035 a novel framework that simultaneously restores both the shape and texture of bro-
036 ken objects using multi-view images. To address limited training data, we de-
037 velop an automated data generation pipeline that synthesizes paired incom-
038 plete samples from large-scale 3D datasets. Central to Restore3D is a multi-
039 view model, enhanced by a carefully designed Mask Self-Perceiver module with
040 a Depth-Aware Mask Rectifier. The rectified masks, learned through the self-
041 perceiver, facilitate an image integration and enhancement phase that preserves
042 shape and texture patterns of incomplete objects and mitigates the low-resolution
043 limitations of the base model, yielding high-resolution, semantically coherent,
044 and view-consistent multi-view images. A coarse-to-fine reconstruction strategy
045 is then employed to recover detailed textured 3D meshes from refined multi-
046 view images. Comprehensive experiments show that Restore3D produces vis-
047 ually and geometrically faithful 3D textured meshes, outperforming existing meth-
048 ods and paving the way for more robust 3D object restoration. Project Page:
049 <https://iclr-subx.github.io/Restore3D/>

050 1 INTRODUCTION

051 Recent advances in 3D generation and reconstruction techniques (Cheng et al., 2023b; Poole et al.,
052 2022; Lin et al., 2023; Li et al., 2023; Tang et al., 2024) have demonstrated impressive capabil-

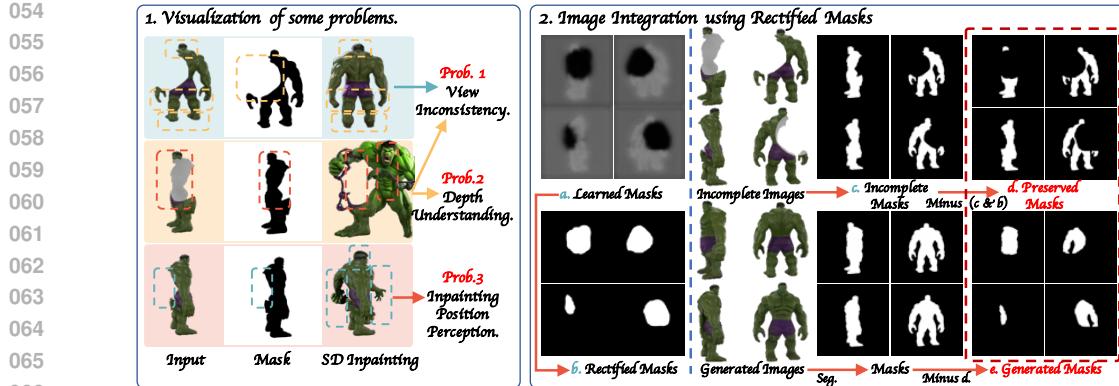


Figure 2: **The importance of masks.** In single-view inpainting, user-provided masks define the regions requiring inpainting. However, in a multi-view context, manually creating consistent masks across all views is impractical. Directly inverting object masks to serve as inpainting masks inevitably causes issues (see Prob. 1 & 3). Moreover, manually adjusting masks based on depth information (see Prob. 2) is labor-intensive and time-consuming. As shown in the right figure (a), our mask self-perceiver can automatically indicate the regions that need to be completed. By leveraging both preserved and generated masks (d & e), our approach retains the incomplete object’s patterns, ensuring accurate and consistent multi-view inpainting. These masks are also used for the image enhancement stage to yield high-resolution restored images (see Fig. 5).

ties, paving the way for innovative applications across diverse fields. Despite these strides, a significant gap remains in the comprehensive restoration of both shape and texture for broken or incomplete 3D objects. This challenge is particularly relevant for some applications such as cultural heritage preservation, occluded objects reconstruction, and artistic creation, where high-fidelity restoration/completion is crucial.

In this study, we aim to develop a robust framework that can simultaneously restore the shape and texture of incomplete 3D objects while handling complex and diverse data types. Key challenges in achieving this goal include: *i) Data Collection.* Existing 3D datasets (Chang et al., 2015; Dai et al., 2017; Rao et al., 2022) focus primarily on shape completion, often neglecting the equally critical aspect of texture restoration. Furthermore, these datasets typically contain simple objects. Creating a diverse, high-quality dataset remains labor-intensive and time-consuming. *ii) Complexity of Object Completion.* Addressing the intricacies of restoring complex and general objects requires a robust framework, as simpler methods typically work only for limited categories of simple objects, but when applied to more complex cases, they often produce inconsistent or incomplete results. The synthesized regions fail to align with the original parts, or even worse, parts of the original structure are overwritten or discarded during the restoration process. *iii) Consistency Preservation of Broken Parts.* Incomplete objects may exhibit varying degrees of degradation in shape and texture. Therefore, preserving the integrity of original components, including consistent color, style, and structural coherence, is crucial for realistic restoration.

To address these challenges, we propose several complementary solutions: **i) Synthetic Data Generation.** To overcome the limitations of existing datasets, we propose to synthesize paired broken and complete data. **ii) Leveraging Foundation Models.** Recent advancements in foundation models (Hong et al., 2023; Shi et al., 2023; Rombach et al., 2022; Oquab et al., 2023; Kirillov et al., 2023; Yang et al., 2024) have demonstrated exceptional generalizability, due to their extensive architectures, large-scale datasets, and adaptability through fine-tuning. We incorporate foundation models to provide prior knowledge, enabling our framework to effectively handle complex and diverse cases. **iii) Task-Specific Structures.** While foundation models offer valuable priors, task-specific components are necessary to tailor their application. Motivated by studies (Zhang et al., 2023b; Ye et al., 2023; Mou et al., 2023), we guide these models toward optimal probability distributions with specialized modules, achieving more accurate and contextually appropriate restorations.

Concretely, we first produce an automatic pipeline to construct paired data, which uses the Boolean modifier in Blender. It offers diverse and large-scale data that are difficult to acquire manually. Second, we propose an innovative framework named **Restore3D**, comprising two key components, *i.e.*, **multi-view image inpainting and reconstruction**. There are several foundational models (Shi

108 et al., 2023; Liu et al., 2023a; Xu et al., 2024) in these two components that we can leverage prior
 109 knowledge to further handle more diverse incomplete objects effectively. However, simply applying
 110 foundational models to multi-view images introduces several **challenges**, as shown in Fig. 2,
 111 including: 1) *View Inconsistency*: Generated results often differ across views, leading to visual in-
 112 coherence. 2) *Depth Understanding*: Existing models often lack robust depth perception, resulting
 113 in failures to recognize occlusions and spatial relationships. 3) *Inpainting Position Perception*: Acc-
 114 curately identifying regions requiring inpainting can be difficult, especially for large masks.

115 To address these issues, we propose a **multi-view** base model combined with a specially de-
 116 signed **mask self-perceiver** module incorporating a **depth-aware mask rectifier**. This module
 117 autonomously perceives and reconstructs missing components, preserving the integrity of original
 118 broken regions and ensuring consistent results across multiple views. Additionally, by leveraging
 119 the preserved and generated masks predicted by the self-perceiver, we can develop an image in-
 120 tegration and enhancement pipeline (see Fig. 2 & 5), yielding high-quality and consistent results.
 121 To convert high-quality multi-view images into 3D objects, we employ large reconstruction mod-
 122 els (LRMs)(Hong et al., 2023; Xu et al., 2024), which offer efficient single- and multi-view object
 123 reconstruction capabilities. To overcome the limitation of coarse outputs from these models, we
 124 adopt a coarse-to-fine refinement approach. Leveraging recent advances in surface normal predic-
 125 tion models(Bae & Davison, 2024; Ye et al., 2024), we inject normal priors to progressively enhance
 126 geometric quality, and refine texture based on updated geometry by using enhanced images. This
 127 ensures that our refined shapes and textures maintain high fidelity, even for complex scenarios.
 128

129 We conduct extensive experiments on Objaverse (Deitke et al., 2023), GSO (Downs et al., 2022),
 130 Breaking Bad Dataset (Sellán et al., 2022), Fantastic Breaks (Lamb et al., 2023) and OmniObject3D
 131 (Wu et al., 2023) to validate the quality of inpainting and reconstruction. The results demon-
 132 strate that our inpainting method significantly outperforms previous approaches (Lugmayr et al., 2022;
 133 Zhang et al., 2023b; Rombach et al., 2022), e.g., $\uparrow 13$ in PSNR compared to Nerfiller (Weber et al.,
 134 2024). By carefully designing a mask self-perceiver, our method can alleviate view inconsis-
 135 tency, understand depth concepts, and capture inpainting regions, achieving consistent structure and
 136 texture styles without requiring user-provided masks to indicate inpainting regions. For reconstruc-
 137 tion, our approach enhances both geometric and texture quality as shown in Fig. 1, indicating that our
 138 proposed framework is capable of producing complete shapes and textures with relatively high fidelity
 139 compared to baseline methods (He & Wang, 2023; Xu et al., 2024; Xiang et al., 2024). Overall, our
 140 contributions are summarized as follows,
 141

- 142 • To the best of our knowledge, we are among the first to explore the completion of relatively com-
 143 plex shapes and textures. To support this task, we introduce an automated data synthesis pipeline
 144 that generates paired incomplete and complete shapes and textures, providing a rich source of
 145 training data named RestoreIt-3D.
- 146 • We propose Restore3D, a novel framework to tackle shape and texture completion through a
 147 combination of multi-view image inpainting and reconstruction. In multi-view image inpainting,
 148 we design a mask self-perceiver with a depth-aware mask rectifier for autonomous perception and
 149 reconstruction of missing components, ensuring preservation of original features. Moreover, we
 150 introduce an image integration and enhancement pipeline to restore fine details. We refine coarse
 151 meshes by using normal priors and enhanced images.
- 152 • Comprehensive experiments validate the effectiveness of Restore3D, demonstrating its ability to
 153 produce complete and high-quality textured meshes.

2 RELATED WORK

154 **2D Inpainting and Generation models** 2D inpainting methods are designed to complete missing
 155 content in an image using a given image and mask. LaMa (Suvorov et al., 2021) utilizes fast Fourier
 156 convolutions, a large receptive field, and extensive training masks to effectively fill large missing
 157 areas, producing plausible inpainting results. Recent advancements in image generation (Rombach
 158 et al., 2022; Zhang et al., 2023b) have demonstrated superior performance and can be adapted for
 159 inpainting tasks with high-quality outcomes. RePaint (Lugmayr et al., 2022) modifies the diffu-
 160 sion generation process, allowing it to be used for inpainting. NeRFiller (Weber et al., 2024) uses
 161 grid priors to make the 2D diffusion model produce more consistent multi-view inpainting results.
 Instant3Dit (Barda et al., 2025) employs a multi-view inpainting model combined with a large recon-

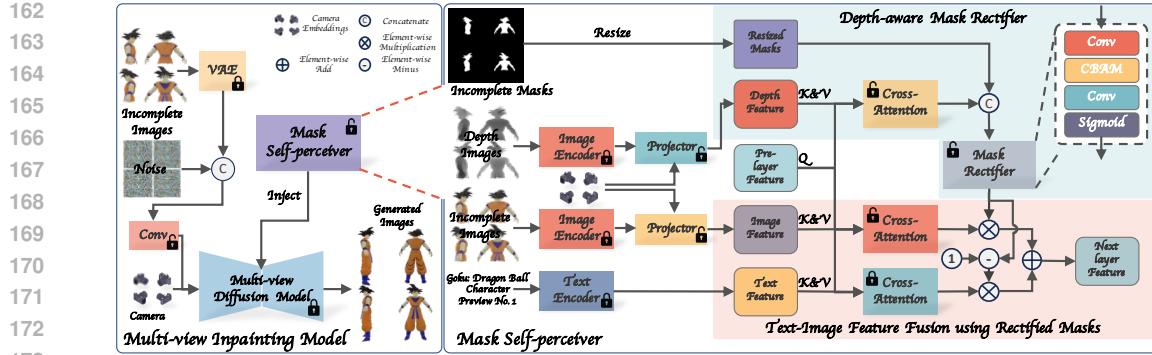


Figure 3: **An overview of multi-view image inpainting.** We carefully design a mask self-perceiver based on a multi-view diffusion model that composes the image and text features with a spatial mask predicted by a depth-aware mask rectifier, therefore the model can automatically perceive the missing part and further generate it meanwhile preserving the original parts.

struction model to enable rapid editing of 3D objects. However, these methods require a user-defined mask to specify the regions that need inpainting.

3D Generation and Completion Recent 3D generation models (Wang et al., 2023b; Lin et al., 2023; Chen et al., 2023c) showcase promising results. DreamFusion (Poole et al., 2022) and SJC (Wang et al., 2023a) are first proposed to generate 3D assets from text using the strong 2D text-to-image generation model (Rombach et al., 2022). As 2D diffusion models easily lead to 3D inconsistency, some works (Liu et al., 2023a; Zhou & Tulsiani, 2023; Tang et al., 2023; Szymanowicz et al., 2023; Tewari et al., 2023; Xu et al., 2023) focus on consistent multi-view image diffusion models. MVDream (Shi et al., 2023) uses 3D self-attention and camera embedding to achieve multi-view text-to-image generation. Considering the time-consuming nature of SDS-based methods, there are some works (Face, 2023; Long et al., 2022; Li et al., 2023; Long et al., 2023; Tang et al., 2024; Wu et al., 2024a; Lu et al., 2024) that use multi-view diffusion models and reconstruction models. Another line for 3D generation is that directly train 3D generative models using 3D representations like point cloud (Nichol et al., 2022; Zeng et al., 2022; Luo & Hu, 2021), meshes (Liu et al., 2023b; Gao et al., 2022), neural fields (Kim et al., 2023; Anciuvevičius et al., 2023; Müller et al., 2023; Jun & Nichol, 2023; Zhang et al., 2023a; Erkoç et al., 2023; Chen et al., 2023b). In addition to 3D generation, recent 3D shape completion works (Kasten et al., 2023; Zhang et al., 2021; Dai & Nießner, 2019; Mittal et al., 2022; Pan et al., 2021; Cheng et al., 2023b; Chu et al., 2023) usually use different types of 3D representations and networks to model global and local structures, *e.g.*, point cloud, sdf, GAN, VAE, and diffusion models. However, they all learn models on small-scale datasets, therefore the modeling capacity is limited compared with some 3D generation models trained on large-scale datasets (*e.g.*, Objavese (Deitke et al., 2023)). Moreover, these works do not consider the texture.

Texture Generation. Several texture generation works (Richardson et al., 2023; Cao et al., 2023; Chen et al., 2023a) use an iteratively texturing strategy based on the pre-trained depth-to-image diffusion models, yielding high-quality texture. However, these methods tend to inherit lighting from training data. Paint3D (Zeng et al., 2023) proposes a shape-aware UV Inpainting and a shape-aware UVHD diffusion model to alleviate this situation. There is another line to learn texture. Texturify (Siddiqui et al., 2022) employs texture maps on the surface of meshes and uses StyleGAN (Karras et al., 2019) to predict texture. Mesh2Tex (Bokhovkin et al., 2023) incorporates an implicit texture field for texture prediction. These methods are lacking in global information modeling. PointUV (Yu et al., 2023) first trains a diffusion model specifically for mesh texture generation, and the proposed coarse-to-fine framework allows it to enjoy the efficiency of 2D representation while enhancing 3D consistency. Other approaches like AUV-net (Chen et al., 2022), LTG (Yu et al., 2021), and TUVF (Cheng et al., 2023a) learn to generate UV-Maps for 3D shapes. However, they typically focus on the texture generation starting from a complete shape.

216

3 METHOD

217

3.1 DATA PREPARATION & TASK DEFINITION

218 **Motivation.** We browse the datasets of related tasks and find that the existing datasets (Chang et al.,
 219 2015; Deitke et al., 2023; Collins et al., 2022) are not sufficient to handle the shape and texture
 220 completion of broken objects, which suggests the need to construct specific broken and complete
 221 paired data. However, collecting large-scale paired data in the real world is *time-consuming and*
 222 *labor-intensive*. Thus we propose to *synthesize* broken and complete paired data.
 223

224 **Data Collection.** We select the recent dataset, G-objaverse (Qiu et al., 2023) that has *more diverse*
 225 and *general objects*, and sample about 83K 3D objects from this dataset.
 226

227 **Synthesis Pipeline.** Specifically, we propose an automatic data processing technique using Boolean
 228 operations (*i.e.*, Difference and Intersect) of Blender. Additionally, we equip the dataset with text
 229 captions using Cap3D (Luo et al., 2023). Subsequently, we normalize and merge the prepared 3D
 230 data. The use of Boolean operations requires the introduction of another object. Therefore, we use
 231 an ico sphere or cube with random size and rotation angle and then randomly place them inside the
 232 3D bounding box of the prepared 3D data to ensure that the objects can be realistically segmented.
 233 After that, it is essential to render this processed data in the format of RGB images to facilitate
 234 model learning. We execute the rendering at a resolution of 256×256 . The camera settings include
 235 a randomly chosen elevation between -10° and 30° . Additionally, the azimuth values are uniformly
 236 rendered from 0° to 360° with a randomly sampled start view, producing a total of 32 images per
 237 object. The Fov of the camera is randomly from 35° to 45° and the distance is always 2.
 238

239 **Task Definition.** The 3D object restoration task aims to reconstruct a complete 3D mesh with texture
 240 from multi-view images of a damaged object. Given **multi-view images** $\{I_1, I_2, \dots, I_n\}$ capturing
 241 a damaged object from different angles and corresponding **camera parameters** (K_i, E_i) , the model
 242 will output a complete 3D mesh $M = (V, F, T)$: Vertices, Faces, and Textures.
 243

244

3.2 MULTI-VIEW IMAGE INPAINTING

245 **Motivation.** Traditional single-view image inpainting methods (Suvorov et al., 2021; Rombach
 246 et al., 2022; Zhang et al., 2023b) rely on the user-provided masks that indicate the areas to be
 247 inpainted. While this approach works well in the context of single-view images, it presents significant
 248 challenges when extended to multi-view contexts as shown in Fig. 2. *1. View inconsistency*. In a
 249 multi-view scenario, the user is required to manually provide a mask for each of the views (*e.g.*, four
 250 views in our case). This also introduces the risk of errors, as the mask needs to be accurately aligned
 251 across different perspectives to maintain 3D consistency. *2. Uncertainty Regarding Inpainting Areas*.
 252 These models cannot autonomously perceive the regions that require inpainting when a large
 253 mask is applied. Additionally, they do not incorporate depth perception, limiting their understand-
 254 ing of occlusion and spatial relationships. To address these challenges, we propose an innovative
 255 approach that enables the model to *ensure view consistency* and *self-perceive the mask*. Concretely,
 256 we design the following two parts.
 257

258 **Mask Self-perceiver.** We propose a mask self-perceiver module based on a multi-view image gen-
 259 eration model as shown in Fig. 3. It has two projectors that consist of transformer-based blocks
 260 and camera modulation layers, which project the depth and image features (f_d, f_r) extracted from
 261 CLIP (Radford et al., 2021) to the diffusion feature space. The camera modulation helps the model
 262 to discriminate the feature under different cameras. Then these projected features (p_d, p_r) will be
 263 fed to the respective cross-attention blocks as key and value $(\mathbf{K}_d, \mathbf{K}_r, \mathbf{V}_d, \mathbf{V}_r)$. The process can be
 264 formulated as follows, where f_* can be depth or image features, p_* is the projected features of them.
 265

$$p_* = \mathbf{Proj}(f_*, c) = \mathbf{Trans}(\mathbf{Mod}(f_*, c)) \quad (1)$$

$$s_* = \mathbf{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right)\mathbf{V}_* \quad (2)$$

266 Similarly, s_* , \mathbf{K}_* and \mathbf{V}_* are the results of p_* via cross-attention and linear layers. \mathbf{Q} originates
 267 from the pre-layer features in the diffusion model.
 268

269 **Depth-aware Mask Rectifier.** Since depth effectively captures the incomplete shape while dis-
 270 regarding texture information, the rectifier can focus solely on identifying the regions that require
 271

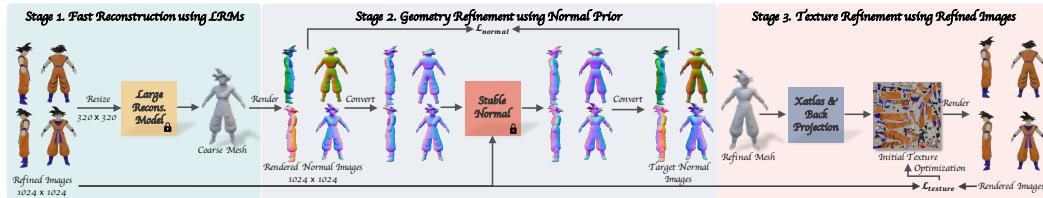


Figure 4: **Geometry and Texture Refinement.** We separately refine the geometry and texture of the coarse results inferred by LRM (Xu et al., 2024).

generation and preservation. Moreover, the depth can help the model understand the spatial relation and occlusion. Specifically, This module leverages depth features obtained after the cross-attention layer, along with incomplete masks, and inputs them into a mask rectifier. The rectifier then outputs a mask indicating where needs to be generated *i.e.*, leveraging the text features and where needs to be preserved *i.e.*, using the image features. The process can be formulated as follows,

$$\mathcal{M}_r = \text{Sigmoid}(\text{Conv}(\text{CBAM}(\text{Conv}[s_d, \mathcal{M}_o]))) \quad (3)$$

$$f_n = (1 - \mathcal{M}_r)s_t + \mathcal{M}_r s_r \quad (4)$$

Conv is convolution layers, **CBAM** is Convolutional Block Attention Module (Woo et al., 2018).

Training objectives Given training samples, including incomplete images \mathcal{I} , depth images \mathcal{D} , incomplete masks \mathcal{M} , text prompts \mathcal{P} and camera embedding \mathcal{C} , the multi-view inpainting loss can be formulated as follows,

$$\mathcal{L} = \min_{\theta} \mathbb{E}_{z, \epsilon \sim \mathcal{N}(\mathbf{o}, \mathbf{I}), t} \|\epsilon - \epsilon_{\theta}(z_t; t, \mathcal{I}, \mathcal{D}, \mathcal{M}, \mathcal{P}, \mathcal{C})\|_2^2. \quad (5)$$

3.3 IMAGE INTEGRATION AND ENHANCEMENT

Motivation. The input resolution of multi-view model is 256 x 256, which is subsequently encoded to 32 x 32 using a Variational Autoencoder. As a result, *local details are compressed, leading to a loss of clarity in both the original and generated regions of the image*. This compression often causes the inpainted part to be unclear, and the reconstructed image may lose fine details that are essential for achieving high-quality results. Moreover, *high-quality images will help the next reconstruction stage to give accurate and detailed textured meshes*. To address these challenges, we propose a pipeline that enables the model to *restore local details and preserve the original patterns*.

Enhancement Models. We explore two types of enhancement models. *Real-ESRGAN* (Wang et al.) is effective at preserving the patterns of low-resolution images with minimal misalignment, making it ideal for recovering the overall structure. *ControlNet-Tile* (Zhang et al., 2023b) offers advanced capabilities for enhancing image details, but will modify the original pattern when a high denoising step is used. Based on these properties, we design the following enhancement pipeline. 1. *Input resolution alignment using Real-ESRGAN*. Before integrating with the original images, we need to align the resolution. Using Real-ESRGAN effectively preserves the overall structure and does not introduce content that is not related to the original style.

2. *Integration of generated and original parts using rectified masks.* As depicted in Fig. 5, this procedure infers the preserved and generated masks used to compose the images, which preserves the original parts as soon as possible. However, this procedure inevitably leads to some artifacts, *e.g.*, inconsistent color transitions. To address these artifacts, we leverage the mentioned property of ControlNet-Tile to enhance the images. 3. *Image harmonizing using ControlNet-Tile with a blending strategy.* Directly using ControlNet-Tile will alter the original pattern and destroy the integration step. Inspired by previous works (Avrahami et al., 2022; Lugmayr et al., 2022), we incorporate a mask blending technique within the diffusion process. This technique helps maintain the original patterns, eliminates any gaps caused by integration in image space, and enhances the image quality.

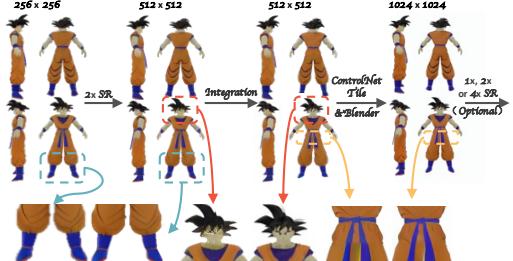


Figure 5: **Image Integration and Enhancement Pipeline using Rectified Masks.**

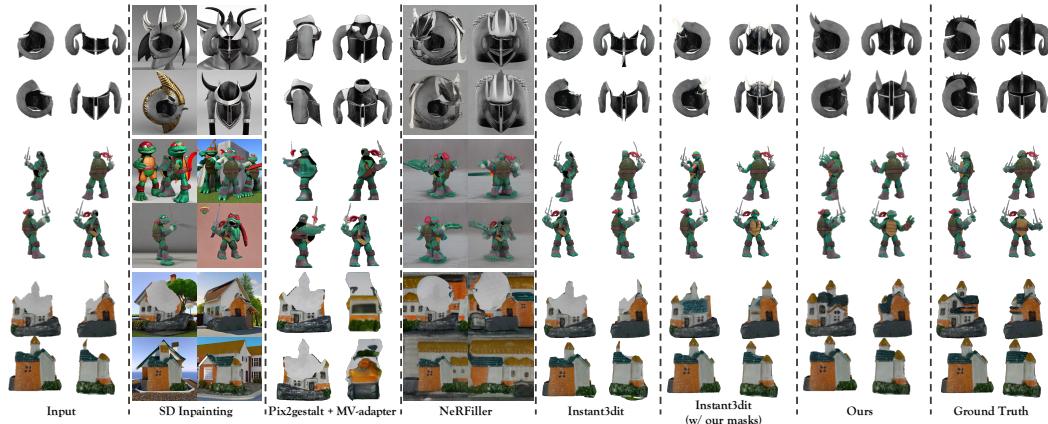


Figure 6: Visual comparison with inpainting methods.

3.4 MULTI-VIEW IMAGE RECONSTRUCTION

Fast Reconstruction using Large Reconstruction Models (LRMs). Recent advancements in LRM (Hong et al., 2023; Tang et al., 2024; Xu et al., 2024), which leverage sophisticated architectures, large-scale datasets, and extensive model parameters, have demonstrated impressive capabilities in 3D object reconstruction from single or sparse-view images. These models are particularly well-suited for tasks requiring fast mesh reconstruction. However, while LRM can produce initial reconstructions efficiently, the results are often *coarse and lack the fine details* necessary for high-quality 3D representations. To address this limitation, we adopt a coarse-to-fine schema and refine the shapes and textures of the outputs generated by LRM, separately, as shown in Fig. 4.

Geometry Refinement using Normal Prior. A key component in optimizing shape structure is to obtain high-quality surface normals. Recent surface normal estimation methods (Ye et al., 2024) have demonstrated the ability to predict relatively accurate normals for in-the-wild monocular images or videos. Therefore, we can employ an *off-the-shelf* normal estimation model to provide normal priors and then use it to optimize the shape structure of 3D objects. Since these models are primarily trained on monocular images or videos, the predicted normals are typically in camera space. Thus we need to convert these normals into world space using camera extrinsic parameters. Specifically, we select StableNorm, a model that accepts coarse rendered normals and RGB images as inputs to predict refined normal outputs. The consistency of the rendered normals contributes to the stability and accuracy of the predicted normals, allowing for more precise geometry refinement.

Texture Refinement using High-quality Images. Since the current shape differs from the coarse shape, the original texture no longer aligns with the updated geometry. Thus we propose to learn the textures that better match the optimized shape. Concretely, we can use Xatlas to obtain UV coordinates, enabling us to back-project the colors from the inpainted images onto the UV textures. After that, we treat the UV textures as parameters and use the high-quality images to optimize it.

Training Objectives. We apply a normal loss \mathcal{L}_{normal} based on the rendered normals \mathcal{I}_n and the target normals $\hat{\mathcal{I}}_n$. Additionally, we apply a mask loss \mathcal{L}_{mask} to ensure that the optimization regions are correctly aligned. The loss function is defined as follows,

$$\mathcal{L}_{shape} = \mathcal{L}_{normal} + \mathcal{L}_{mask} = \|\mathcal{I}_n - \hat{\mathcal{I}}_n\|_2^2 + \|\mathcal{M} - \hat{\mathcal{M}}\|_2^2. \quad (6)$$

To optimize the texture, we use a RGB loss \mathcal{L}_{rgb} on the rendered images \mathcal{I}_{rgb} and enhanced images $\hat{\mathcal{I}}_{rgb}$. The mask loss \mathcal{L}_{mask} is also applied. Moreover, the SSIM \mathcal{L}_{ssim} loss is introduced to improve the texture quality. The loss functions are defined as follows, where λ is a weight parameter.

$$\mathcal{L}_{tex} = \mathcal{L}_{rgb} + \mathcal{L}_{mask} + \lambda \mathcal{L}_{ssim} = \|\mathcal{I}_{rgb} - \hat{\mathcal{I}}_{rgb}\|_2^2 + \|\mathcal{M} - \hat{\mathcal{M}}\|_2^2 + \lambda \text{SSIM}(\mathcal{I}, \hat{\mathcal{I}}), \quad (7)$$

4 EXPERIMENTS

Dataset. For model training, we sample approximately 83K data from the G-objaverse dataset (Qiu et al., 2023) and process them using our proposed pipeline. For model testing, we sample approximately 350 data from the GSO (Downs et al., 2022), Omniobject (Wu et al., 2023), and

378 Table 1: **Comparison with inpainting and reconstruction methods.** \triangle means using Depth-
 379 Anything (Yang et al., 2024) to obtain the depth images. \clubsuit means using MV-adapter(Huang et al.,
 380 2024). \heartsuit means using our model’s predicted masks as inpainting masks.

(a) Inpainting.				(b) Reconstruction.				
Method	PSNR \uparrow	LPIPS \downarrow	FID \downarrow	Method	PSNR \uparrow	LPIPS \downarrow	CD \downarrow	
Repaint	10.55	0.31	69.57	0.76	Open-LRM	16.90	0.15	0.011
SD	12.58	0.22	61.15	0.83	InstantMesh	20.60	0.11	0.006
ControlNet	10.66	0.30	69.91	0.76	Unique3D	22.00	0.14	0.005
Pix2gestalt \clubsuit	16.43	0.21	75.08	0.86	Direct3D	-	-	0.006
NeRFiller	12.03	0.25	65.20	0.82	Trellis	21.78	0.12	0.005
Instant3dit	19.40	0.10	48.03	0.94	Hunyuan3D-2	21.31	0.14	0.006
Instant3dit \heartsuit	22.37	0.07	36.08	0.95	Amodal3R	19.37	0.15	0.008
Ours \triangle	25.29	0.07	32.05	0.95	Ours	23.35	0.09	0.005
Ours	25.50	0.06	31.82	0.95				

391 Table 2: **Generalization Ability.**

(a) Fantastic Breaks Dataset.				(b) Breaking Bad Dataset.			
Method	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	Method	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow
SD	12.59	0.72	0.40	SD	12.02	0.74	0.53
Controlnet	15.63	0.55	0.56	ControlNet	14.50	0.59	0.71
Nerfiller	18.94	0.52	0.81	NeRFiller	17.66	0.52	0.79
Instant3dit	23.11	0.14	0.96	Instant3dit	22.27	0.15	0.95
Ours	26.91	0.09	0.97	Ours	25.09	0.10	0.95

400 Objaverse (Deitke et al., 2023) datasets. We also test our model on the Breaking Bad Dataset (Sellán
 401 et al., 2022) and Fantastic Breaks (Lamb et al., 2023), which include physically simulated and real-
 402 world broken objects, to evaluate its generalizability.

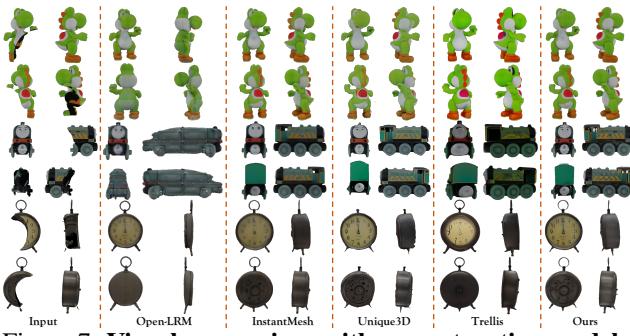
403 **Metrics.** To assess image quality, we choose Peak Signal-to-Noise Ratio (PSNR), Frechet Inception
 404 Distance (FID), Learned Perceptual Image Patch Similarity (LPIPS), and Structural Similarity Index
 405 Measure (SSIM). We evaluate geometry quality using Chamfer Distance (CD) and F-scores.

407 4.1 INPAINTING RESULTS.

409 **Baselines.** We compare our method with single-view image inpainting, *i.e.*, Repaint(Lugmayr et al.,
 410 2022), Stable-Diffusion (Rombach et al., 2022), Controlnet (Zhang et al., 2023b), *i.e.*, Pix2gestalt
 411 + MV-adater (Ozguroglu et al., 2024; Huang et al., 2024) and multi-view inpainting methods, *i.e.*,
 412 Nerfiller (Weber et al., 2024) and Instant3dit (Barda et al., 2025). **Note that we do not use the image**
 413 **integration and enhancement pipeline for a fair evaluation.**

414 **Qualitative Comparison.** As shown in Fig. 6, the results demonstrate that our model produces
 415 plausible and coherent inpainting outcomes. Previous methods require user-provided masks to guide
 416 the model in generating missing parts. When given a relatively large mask, these methods struggle
 417 to capture the inherent structure of the objects, leading to less accurate and coherent inpainting.
 418 In contrast, our approach does not
 419 require predefined inpainting masks.
 420 It autonomously perceives and recon-
 421 structures missing regions, capturing the
 422 underlying structure of the object.
 423 This capability allows our method to
 424 produce high-quality and structurally
 425 consistent inpainting results.

426 **Quantitative Comparison.** As illus-
 427 trated in Table 1a, we observe the fol-
 428 lowing: **1)** Our approach achieves the
 429 best performance in restoring shape
 430 and texture. **2)** When applying depth images predicted by Depth-Anything (Yang et al., 2024),
 431 our method yields results comparable to those obtained with ground truth depths. **3)** The compared
 432 methods produce noticeably inferior results in terms of inpainting quality.

Figure 7: **Visual comparison with reconstruction models.**

432 **Table 3: Ablation studies for multi-view inpainting and reconstruction.**
 433 (a) **Inpainting.** (b) **Reconstruction.**

Method	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	Method	PSNR \uparrow	LPIPS \downarrow	CD \downarrow	F-Score \uparrow
IF	22.65	0.14	0.90	Baseline	20.60	0.11	0.006	0.321
IF + Conv	26.53	0.08	0.94	GR	-	-	0.005	0.389
IF + Conv + DMR	29.44	0.06	0.95	GR + TR	23.35	0.09	0.005	0.389

439 **Generalization Ability.** 1. *Physically simulated broken objects.* As shown in Fig. 1 and Table 2b,
 440 we further test our model on the Breaking Bad Dataset (Sellán et al., 2022), synthesized by a phys-
 441 ically based method that simulates the natural destruction process of geometric objects. 2. *Real-
 442 world broken objects.* As shown in Fig. 1 and Table 2a, we also evaluate our model on Fantastic
 443 Breaks (Lamb et al., 2023). These experiments demonstrate the generalization ability of our model
 444 to both **unseen** real-world scenarios and physically simulated cases, validating its robustness and
 445 practical applicability, despite being trained solely on synthetic data.

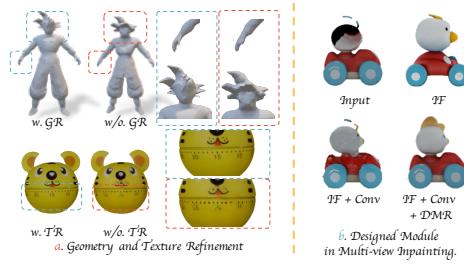
446 4.2 RECONSTRUCTION RESULTS.

447 **Baselines.** We compare our method against both single-view and multi-view LRMs, including LRM
 448 (He & Wang, 2023; Hong et al., 2023) and InstantMesh (Xu et al., 2024), Unique3D (Wu et al.,
 449 2024a). We also compare our method with image-to-3D generation methods, Direct3D (Wu et al.,
 450 2024b), and Trellis (Xiang et al., 2024). For single-view baselines, we input the front-view image.
 451 [All of the methods use our inpainted and enhanced images as input for a fair comparison.](#)

452 **Quantitative & Qualitative Comparison.** As shown in Table 1b, our method achieves superior
 453 rendered image quality and geometry accuracy, with a substantial improvement over baseline meth-
 454 ods. In Fig. 7, it is evident that our approach delivers clearer details and the most accurate geometry
 455 among the compared methods. **Training time.** Our approach is highly efficient, requiring 20 sec-
 456 onds per object for geometry and texture refinements.

457 4.3 ABLATION STUDY

458 **Multi-view Inpainting.** We conduct ablation studies on the proposed multiview Inpainting module
 459 in the following components: 1) **IF.** Only inputting incomplete images into the cross-attention layers.
 460 2) **Conv.** Concatenating noise and incomplete im-
 461 ages to a learnable convolutional layer. 3) **DMR.**
 462 Adding the designed Depth-aware Mask Rectifier.
 463 As shown in Table 3a, the results improve progres-
 464 sively with each added component, and using all de-
 465 signed components achieves the highest results. As
 466 shown in Fig. 8b, 1) IF Only: the model captures the
 467 general style of the object but lacks an understand-
 468 ing of spatial relationships and structure. 2) IF + Conv:
 469 This enables the model to capture spatial position-
 470 ing and understand object structure. However, it is
 471 still prone to color inaccuracies, especially in areas like the head (blended with error black color).
 472 Additionally, the region that needs to be preserved is changed. 3) IF + Conv + DMR: This allows
 473 the model to improve its ability to handle occlusions and spatial relationships, producing the best
 474 inpainting quality, with coherent colors and well-preserved spatial structure.



475 **Figure 8: Visualization of ablation studies.**

476 **Reconstruction.** We evaluate the impact of the following components: 1) Geometry Refine-
 477 ment (GR), and 2) Texture Refinement (TR). In Table 3b and Fig. 8a, incorporating GR leads to
 478 substantial improvements in geometry quality. TR improves the visual quality of rendered images.

481 5 CONCLUSION

482 In this paper, we propose a novel framework named Restore3D, consisting of multi-view image
 483 inpainting and reconstruction, to simultaneously complete both the shape and texture of broken
 484 3D objects. To facilitate this task, we develop an automated data processing pipeline that collects

486 pair-wise data from a large-scale dataset (Deitke et al., 2023). In the multi-view image inpainting,
 487 we design a mask self-perceiver with a depth-aware mask rectifier. This component autonomously
 488 identifies and reconstructs missing regions while preserving the original patterns. To address the low
 489 resolution resulting from the base model (Shi et al., 2023), we implement an image integration and
 490 enhancement pipeline, allowing for seamless integration and detail enhancement by learned masks.
 491 For the reconstruction stage, we employ an LRM to quickly generate a coarse result, followed by
 492 separate geometry refinement using normal priors and texture refinement using enhanced images.
 493 Through this designed framework, our model produces coherent completions of broken objects as
 494 illustrated in Fig. 1. Moreover, our designed framework can also handle simple 3D object editing
 495 and occluded objects.

496

497 REFERENCES

498 Titas Anciukevičius, Zexiang Xu, Matthew Fisher, Paul Henderson, Hakan Bilen, Niloy J Mitra, and
 499 Paul Guerrero. Renderdiffusion: Image diffusion for 3d reconstruction, inpainting and generation.
 500 In *CVPR*, 2023.

501

502 Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of
 503 natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern
 504 Recognition (CVPR)*, pp. 18208–18218, June 2022.

505

506 Gwangbin Bae and Andrew J. Davison. Rethinking inductive biases for surface normal estimation.
 507 In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.

508

509 Amir Barda, Matheus Gadelha, Vladimir G. Kim, Noam Aigerman, Amit H. Bermano, and Thibault
 510 Groueix. Instant3dit: Multiview inpainting for fast editing of 3d objects, 2025.

511

512 Alexey Bokhovkin, Shubham Tulsiani, and Angela Dai. Mesh2tex: Generating mesh textures from
 513 image queries. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*,
 pp. 8918–8928, 2023.

514

515 Tianshi Cao, Karsten Kreis, Sanja Fidler, Nicholas Sharp, and Kangxue Yin. Texfusion: Synthe-
 516 sizing 3d textures with text-guided image diffusion models. In *Proceedings of the IEEE/CVF
 517 International Conference on Computer Vision*, pp. 4169–4181, 2023.

518

519 Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li,
 520 Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d
 521 model repository. *arXiv preprint arXiv:1512.03012*, 2015.

522

523 Dave Zhenyu Chen, Yawar Siddiqui, Hsin-Ying Lee, Sergey Tulyakov, and Matthias Nießner.
 524 Text2tex: Text-driven texture synthesis via diffusion models. *arXiv preprint arXiv:2303.11396*,
 525 2023a.

526

527 Hansheng Chen, Jiatao Gu, Anpei Chen, Wei Tian, Zhuowen Tu, Lingjie Liu, and Hao Su. Single-
 528 stage diffusion nerf: A unified approach to 3d generation and reconstruction. In *ICCV*, 2023b.

529

530 Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and
 531 appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873*, 2023c.

532

533 Zhiqin Chen, Kangxue Yin, and Sanja Fidler. Auv-net: Learning aligned uv maps for texture transfer
 534 and synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern
 535 Recognition*, pp. 1465–1474, 2022.

536

537 An-Chieh Cheng, Xuetong Li, Sifei Liu, and Xiaolong Wang. Tuvf: Learning generalizable texture
 538 uv radiance fields. *arXiv preprint arXiv:2305.03040*, 2023a.

539

540 Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sd-
 541 fusion: Multimodal 3d shape completion, reconstruction, and generation. In *CVPR*, 2023b.

542

543 Ruihang Chu, Enze Xie, Shentong Mo, Zhenguo Li, Matthias Nießner, Chi-Wing Fu, and Jiaya Jia.
 544 Diffcomplete: Diffusion-based generative 3d shape completion, 2023.

540 Jasmine Collins, Shubham Goel, Kenan Deng, Achleshwar Luthra, Leon Xu, Erhan Gundogdu,
 541 Xi Zhang, Tomas F Yago Vicente, Thomas Dideriksen, Himanshu Arora, Matthieu Guillaumin,
 542 and Jitendra Malik. Abo: Dataset and benchmarks for real-world 3d object understanding. *CVPR*,
 543 2022.

544 Angela Dai and Matthias Nießner. Scan2mesh: From unstructured range scans to 3d meshes, 2019.

545 Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-
 546 predictor cnns and shape synthesis. In *Proceedings of the IEEE conference on computer vision*
 547 *and pattern recognition*, pp. 5868–5877, 2017.

548 Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig
 549 Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of anno-
 550 tated 3d objects. In *CVPR*, 2023.

551 Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann,
 552 Thomas B. McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of
 553 3d scanned household items, 2022. URL <https://arxiv.org/abs/2204.11918>.

554 Ziya Erkoç, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. Hyperdiffusion: Generat-
 555 ing implicit neural fields with weight-space diffusion. *arXiv preprint arXiv:2303.17015*, 2023.

556 Hugging Face. One-2-3-45. [https://huggingface.co/spaces/One-2-3-45/](https://huggingface.co/spaces/One-2-3-45/One-2-3-45)
 557 One-2-3-45, 2023.

558 Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan
 559 Gojcic, and Sanja Fidler. Get3d: A generative model of high quality 3d textured shapes learned
 560 from images. *NeurIPS*, 2022.

561 Zexin He and Tengfei Wang. Openlrm: Open-source large reconstruction models. <https://github.com/3DTopia/OpenLRM>, 2023.

562 Yicong Hong, Kai Zhang, Juxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli,
 563 Trung Bui, and Hao Tan. Lrm: Large reconstruction model for single image to 3d. *arXiv preprint*
 564 *arXiv:2311.04400*, 2023.

565 Zehuan Huang, Yuan-Chen Guo, Haoran Wang, Ran Yi, Lizhuang Ma, Yan-Pei Cao, and Lu Sheng.
 566 Mv-adapter: Multi-view consistent image generation made easy, 2024. URL <https://arxiv.org/abs/2412.03632>.

567 Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions, 2023.

568 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative
 569 adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*
 570 *recognition*, pp. 4401–4410, 2019.

571 Yoni Kasten, Ohad Rahamim, and Gal Chechik. Point-cloud completion with pretrained text-to-
 572 image diffusion models, 2023.

573 Seung Wook Kim, Bradley Brown, Kangxue Yin, Karsten Kreis, Katja Schwarz, Daiqing Li, Robin
 574 Rombach, Antonio Torralba, and Sanja Fidler. Neuralfield-ldm: Scene generation with hierarchi-
 575 cal latent diffusion models. In *CVPR*, 2023.

576 Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete
 577 Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick.
 578 Segment anything. *arXiv:2304.02643*, 2023.

579 Nikolas Lamb, Cameron Palmer, Benjamin Molloy, Sean Banerjee, and Natasha Kholgade Baner-
 580 jee. Fantastic breaks: A dataset of paired 3d scans of real-world broken objects and their com-
 581 plete counterparts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
 582 *Recognition (CVPR)*, pp. 4681–4691, June 2023.

583 Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan
 584 Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3d: Fast text-to-3d with sparse-view gen-
 585 eration and large reconstruction model, 2023.

594 Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten
 595 Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content
 596 creation. In *CVPR*, 2023.

597

598 Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick.
 599 Zero-1-to-3: Zero-shot one image to 3d object. In *ICCV*, 2023a.

600

601 Zhen Liu, Yao Feng, Michael J Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu.
 602 Meshdiffusion: Score-based generative 3d mesh modeling. In *ICLR*, 2023b.

603

604 Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. Sparseneus: Fast gen-
 605 eralizable neural surface reconstruction from sparse views. In *European Conference on Computer
 606 Vision*, pp. 210–227. Springer, 2022.

607

608 Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma,
 609 Song-Hai Zhang, Marc Habermann, Christian Theobalt, and Wenping Wang. Wonder3d: Single
 610 image to 3d using cross-domain diffusion, 2023.

611

612 Yuanxun Lu, Jingyang Zhang, Shiwei Li, Tian Fang, David McKinnon, Yanghai Tsin, Long Quan,
 613 Xun Cao, and Yao Yao. Direct2.5: Diverse text-to-3d generation via multi-view 2.5d diffusion.
 614 *Computer Vision and Pattern Recognition (CVPR)*, 2024.

615

616 Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van
 617 Gool. Repaint: Inpainting using denoising diffusion probabilistic models, 2022. URL <https://arxiv.org/abs/2201.09865>.

618

619 Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceed-
 620 ings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2837–2845,
 621 2021.

622

623 Tiange Luo, Chris Rockwell, Honglak Lee, and Justin Johnson. Scalable 3d captioning with pre-
 624 trained models, 2023. URL <https://arxiv.org/abs/2306.07279>.

625

626 Paritosh Mittal, Yen-Chi Cheng, Maneesh Singh, and Shubham Tulsiani. Autosdf: Shape priors for
 627 3d completion, reconstruction and generation. In *Proceedings of the IEEE/CVF Conference on
 628 Computer Vision and Pattern Recognition*, pp. 306–315, 2022.

629

630 Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, Ying Shan, and
 631 Xiaohu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image
 632 diffusion models. *arXiv preprint arXiv:2302.08453*, 2023.

633

634 Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulo, Peter Kortscheder, and
 635 Matthias Nießner. Diffrr: Rendering-guided 3d radiance field diffusion. In *CVPR*, 2023.

636

637 Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system
 638 for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022.

639

640 Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov,
 641 Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao
 642 Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran,
 643 Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Ar-
 644 mand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision,
 645 2023.

646

647 Ege Ozguroglu, Ruoshi Liu, Dídac Surís, Dian Chen, Achal Dave, Pavel Tokmakov, and Carl
 648 Vondrick. pix2gestalt: Amodal segmentation by synthesizing wholes, 2024. URL <https://arxiv.org/abs/2401.14398>.

649

650 Liang Pan, Xinyi Chen, Zhongang Cai, Junzhe Zhang, Haiyu Zhao, Shuai Yi, and Ziwei Liu. Vari-
 651 ational relational point completion network, 2021.

652

653 Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d
 654 diffusion. *arXiv preprint arXiv:2209.14988*, 2022.

648 Lingteng Qiu, Guanying Chen, Xiaodong Gu, Qi Zuo, Mutian Xu, Yushuang Wu, Weihao Yuan,
 649 Zilong Dong, Liefeng Bo, and Xiaoguang Han. Richdreamer: A generalizable normal-depth
 650 diffusion model for detail richness in text-to-3d, 2023.

651

652 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agar-
 653 wal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya
 654 Sutskever. Learning transferable visual models from natural language supervision, 2021. URL
 655 <https://arxiv.org/abs/2103.00020>.

656

657 Yuchen Rao, Yinyu Nie, and Angela Dai. Patchcomplete: Learning multi-resolution patch priors for
 658 3d shape completion on unseen categories, 2022.

659

660 Elad Richardson, Gal Metzer, Yuval Alaluf, Raja Giryes, and Daniel Cohen-Or. Texture: Text-
 661 guided texturing of 3d shapes. *arXiv preprint arXiv:2302.01721*, 2023.

662

663 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 664 resolution image synthesis with latent diffusion models, 2022.

665

666 Silvia Sellán, Yun-Chun Chen, Ziyi Wu, Animesh Garg, and Alec Jacobson. Breaking bad: A
 667 dataset for geometric fracture and reassembly, 2022. URL <https://arxiv.org/abs/2210.11463>.

668

669 Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multi-view
 670 diffusion for 3d generation. *arXiv preprint arXiv:2308.16512*, 2023.

671

672 Yawar Siddiqui, Justus Thies, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. Textu-
 673 rify: Generating textures on 3d shape surfaces. In *European Conference on Computer Vision*, pp.
 674 72–88. Springer, 2022.

675

676 Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha,
 677 Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempit-
 678 sky. Resolution-robust large mask inpainting with fourier convolutions. *arXiv preprint
 679 arXiv:2109.07161*, 2021.

680

681 Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Viewset diffusion:(0-) image-
 682 conditioned 3d generative models from 2d data. *arXiv preprint arXiv:2306.07881*, 2023.

683

684 Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm:
 685 Large multi-view gaussian model for high-resolution 3d content creation, 2024.

686

687 Shitao Tang, Fuyang Zhang, Jiacheng Chen, Peng Wang, and Yasutaka Furukawa. Mvdiffusion: En-
 688 abling holistic multi-view image generation with correspondence-aware diffusion. *arXiv preprint
 689 arXiv:2307.01097*, 2023.

690

691 Ayush Tewari, Tianwei Yin, George Cazenavette, Semon Rezhikov, Joshua B Tenenbaum, Frédo
 692 Durand, William T Freeman, and Vincent Sitzmann. Diffusion with forward models: Solving
 693 stochastic inverse problems without direct supervision. *arXiv preprint arXiv:2306.11719*, 2023.

694

695 Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian
 696 chaining: Lifting pretrained 2d diffusion models for 3d generation. In *CVPR*, 2023a.

697

698 Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind
 699 super-resolution with pure synthetic data. In *International Conference on Computer Vision Work-
 700 shops (ICCVW)*.

701

702 Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolif-
 703 dicrimer: High-fidelity and diverse text-to-3d generation with variational score distillation. *arXiv
 704 preprint arXiv:2305.16213*, 2023b.

705

706 Ethan Weber, Aleksander Holynski, Varun Jampani, Saurabh Saxena, Noah Snavely, Abhishek Kar,
 707 and Angjoo Kanazawa. Nerfiller: Completing scenes via generative 3d inpainting. In *CVPR*,
 708 2024.

702 Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block
 703 attention module, 2018. URL <https://arxiv.org/abs/1807.06521>.

704

705 Kailu Wu, Fangfu Liu, Zhihan Cai, Runjie Yan, Hanyang Wang, Yating Hu, Yueqi Duan, and
 706 Kaisheng Ma. Unique3d: High-quality and efficient 3d mesh generation from a single image,
 707 2024a.

708 Shuang Wu, Youtian Lin, Feihu Zhang, Yifei Zeng, Jingxi Xu, Philip Torr, Xun Cao, and Yao Yao.
 709 Direct3d: Scalable image-to-3d generation via 3d latent diffusion transformer. *arXiv preprint*
 710 *arXiv:2405.14832*, 2024b.

711

712 Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Liang Pan Jiawei Ren, Wayne Wu, Lei Yang, Jiaqi
 713 Wang, Chen Qian, Dahua Lin, and Ziwei Liu. Omniobject3d: Large-vocabulary 3d object dataset
 714 for realistic perception, reconstruction and generation. In *IEEE/CVF Conference on Computer*
 715 *Vision and Pattern Recognition (CVPR)*, 2023.

716 Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin
 717 Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. *arXiv*
 718 *preprint arXiv:2412.01506*, 2024.

719 Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh:
 720 Efficient 3d mesh generation from a single image with sparse-view large reconstruction models.
 721 *arXiv preprint arXiv:2404.07191*, 2024.

722

723 Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli,
 724 Gordon Wetzstein, Zexiang Xu, and Kai Zhang. Dmv3d: Denoising multi-view diffusion using
 725 3d large reconstruction model, 2023.

726 Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang
 727 Zhao. Depth anything v2. *arXiv:2406.09414*, 2024.

728

729 Chongjie Ye, Lingteng Qiu, Xiaodong Gu, Qi Zuo, Yushuang Wu, Zilong Dong, Liefeng Bo, Yuliang
 730 Xiu, and Xiaoguang Han. Stablenormal: Reducing diffusion variance for stable and sharp normal.
 731 *ACM Transactions on Graphics (TOG)*, 2024.

732 Hu Ye, Jun Zhang, Sibo Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt
 733 adapter for text-to-image diffusion models. 2023.

734

735 Rui Yu, Yue Dong, Pieter Peers, and Xin Tong. Learning texture generators for 3d shape collections
 736 from internet photo sets. In *British Machine Vision Conference*, 2021.

737 Xin Yu, Peng Dai, Wenbo Li, Lan Ma, Zhengzhe Liu, and Xiaojuan Qi. Texture generation on 3d
 738 meshes with point-uv diffusion. In *Proceedings of the IEEE/CVF International Conference on*
 739 *Computer Vision*, pp. 4206–4216, 2023.

740

741 Xianfang Zeng, Xin Chen, Zhongqi Qi, Wen Liu, Zibo Zhao, Zhibin Wang, Bin Fu, Yong Liu, and
 742 Gang Yu. Paint3d: Paint anything 3d with lighting-less texture diffusion models, 2023.

743

744 Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten
 745 Kreis. Lion: Latent point diffusion models for 3d shape generation. In *NeurIPS*, 2022.

746

747 Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape
 748 representation for neural fields and generative diffusion models. In *SIGGRAPH*, 2023a.

749

750 Junzhe Zhang, Xinyi Chen, Zhongang Cai, Liang Pan, Haiyu Zhao, Shuai Yi, Chai Kiat Yeo, Bo Dai,
 751 and Chen Change Loy. Unsupervised 3d shape completion through gan inversion. In *Proceedings*
 752 *of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1768–1777, 2021.

753

754 Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image
 755 diffusion models, 2023b.

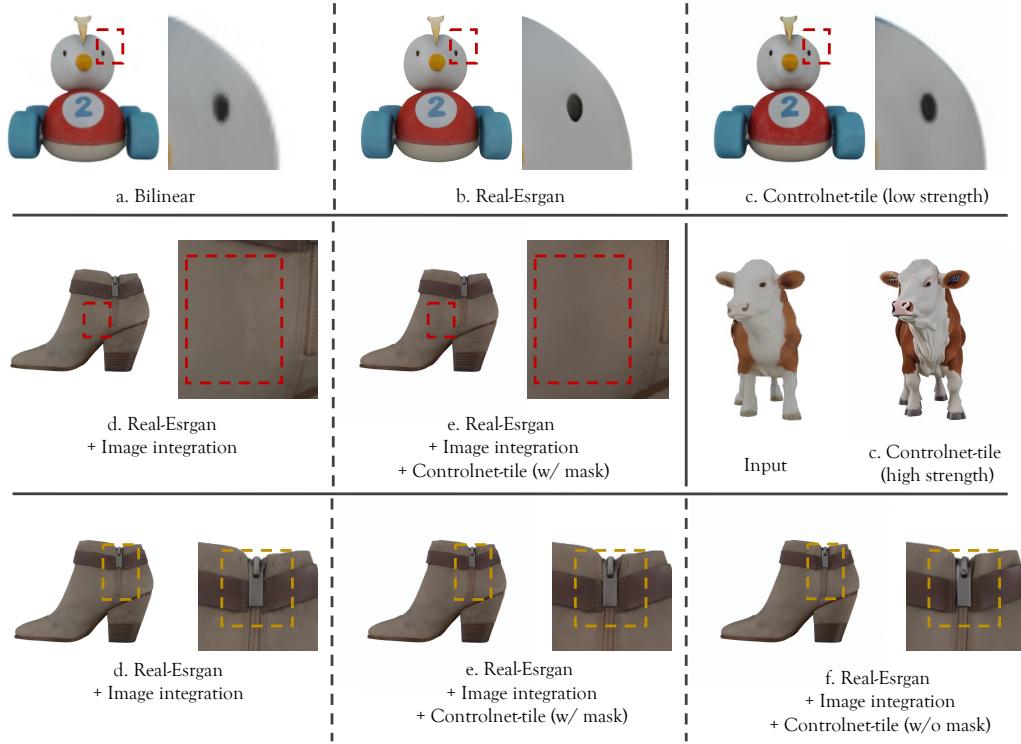
756

757 Zhizhuo Zhou and Shubham Tulsiani. Sparsefusion: Distilling view-conditioned diffusion for 3d
 758 reconstruction. In *CVPR*, 2023.

810 comes smaller. This is expected, as smaller missing regions provide the model with richer contextual
 811 information, enabling more accurate inference of the missing shape.

812 Moreover, when the missing region becomes very large, the network gains more flexibility in generating
 813 plausible content. In such cases, the output may deviate from the original ground truth, but this
 814 discrepancy should not necessarily be considered an "error." This is because extremely large
 815 missing regions often provide little or no contextual guidance, leading to inherently ambiguous re-
 816 constructions.

818 A.2 MORE DETAILS ABOUT IMAGE INTEGRATION AND ENHANCEMENT



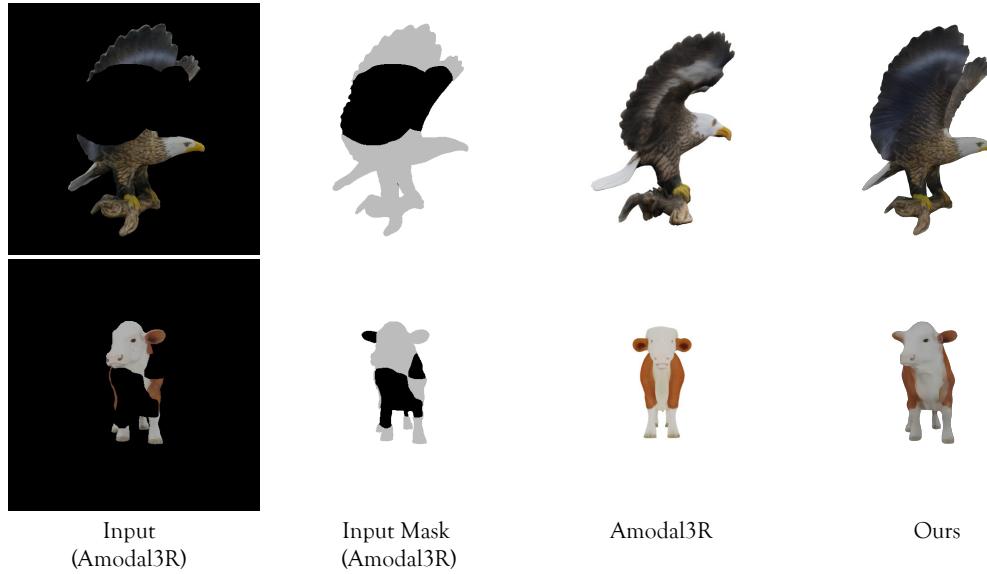
846 **Figure 10: Visualization of Image Integration and Enhancement**

847 **Table 4: Ablation studies for Image Integration and Enhancement.**

851 Method (256px to 1024px)	852 PSNR \uparrow	853 LPIPS \downarrow	854 SSIM \uparrow
855 Baseline (Bilinear Upsampling)	856 26.83	857 0.10	858 0.97
859 4x Real-ESRGAN	860 26.59	861 0.08	862 0.97
863 4x Controlnet-tile	864 26.56	865 0.08	866 0.96
867 Real-ESRGAN + Image Integration	868 27.13	869 0.06	870 0.97
871 Real-ESRGAN + Image Integration + Controlnet-tile (w/ mask blending)	872 26.94	873 0.06	874 0.97
875 Real-ESRGAN + Image Integration + Controlnet-tile (w/o mask blending)	876 26.55	877 0.07	878 0.97

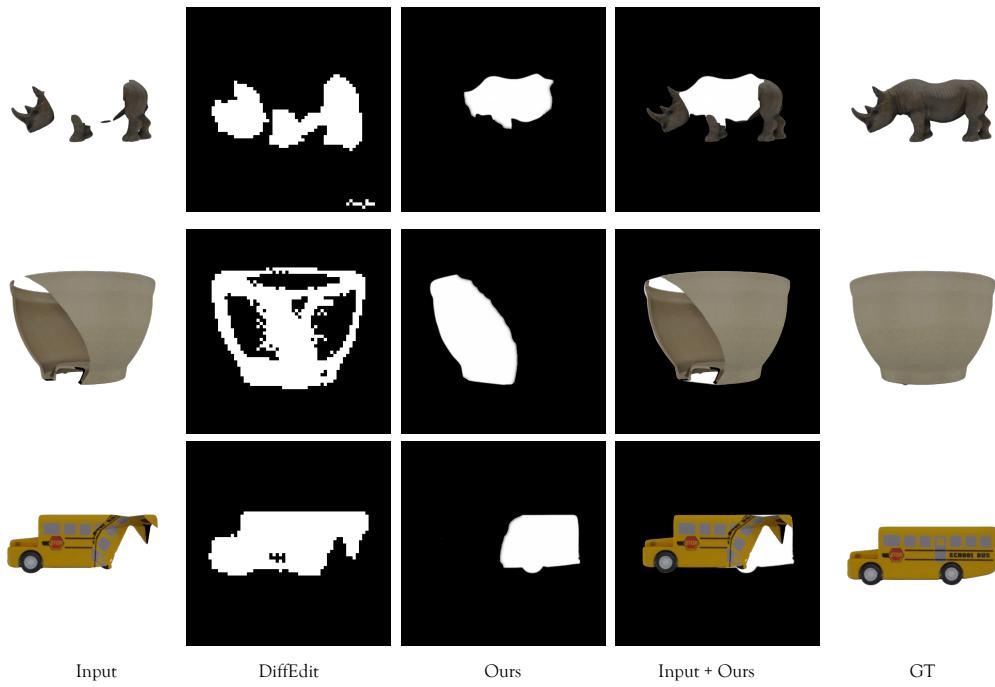
879 We conduct a more detailed ablation study as shown in the Table 4. The visualization results are
 880 shown in Fig. 10. We observed that: 1. Solely applying enhancement methods does not improve
 881 the quantitative metrics, but can improve visual quality. 2. The performance gains mainly originate
 882 from the image integration, which also validates that our rectified mask well indicates the regions
 883 requiring inpainting or preservation.

864 Overall, the organization of this stage is flexible. The key ideas are: 1. Use ControlNet-Tile with
 865 a mask-blending strategy to eliminate color inconsistencies during image integration. 2. Upsample
 866 the image to the desired resolution using Real-ESRGAN, either before or after the integration step.
 867

Figure 11: **Comparison with Amodal3R.**

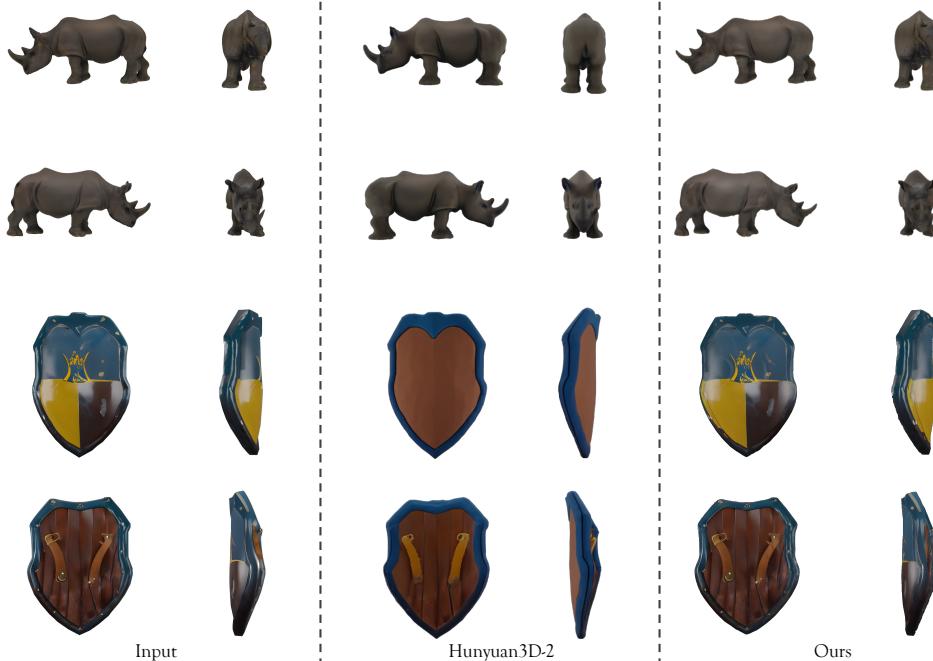
A.3 COMPARISON WITH AMODAL3R.

891 As shown in Figure 11, We find that the results of Amodal3R often misalign the conditioned images
 892 and masks. Furthermore, we notice that the base model used by Amodal3R (Trellis) also faced
 893 similar issues with misalignment and inconsistencies, which in turn affected its ability to generate
 894 accurate completions.

Figure 12: **Comparison with DiffEdit.**

918 A.4 COMPARISON WITH DIFFEDIT.
919

920 As shown in Figure 12, our method significantly outperforms DiffEdit in terms of mask quality.
 921 DiffEdit relies solely on the difference between the noise-conditioned and unconditioned text to infer
 922 the mask. However, this approach is suboptimal because it does not account for explicit image and
 923 depth information, which are crucial for guiding the model to generate more accurate, contextually
 924 appropriate masks in the object restoration task. In contrast, our method incorporates both the image
 925 and depth as conditions, significantly improving the quality of mask generation.

950 Figure 13: Comparison with Hunyuan3D-2.
951952 A.5 COMPARISON WITH HUNYUAN3D-2.
953

954 The shape-generation and texture-generation models in Hunyuan3D-2 are very large, so we use
 955 the fast version for inference. Even with the fast version, generating meshes still takes a long
 956 time—for example, shape generation alone often requires several minutes, while the texture-
 957 generation pipeline typically takes more than 30 minutes. In contrast, our model produces a fully
 958 textured mesh in only about one minute. As shown in Figure 13, the results of our model better align
 959 with the input images than Hunyuan3D-2.

960 A.6 THE ROLE OF COARSE METHERS INFERRED BY LRM
961

962 Without LRM, a typical alternative is to start from a simple primitive (e.g., a sphere) and opti-
 963 mize its shape using our geometry losses. As reported in Table 5, LRM provide a much better
 964 initialization, leading to faster convergence and improved reconstruction quality.

966 Table 5: The role of coarse methes inferred by LRM.
967

Method	CD \downarrow	F-score \uparrow
Sphere + geometry optimization	0.02	0.197
LRM + geometry optimization	0.005	0.389

972 A.7 COMPUTE BUDGETS.
973974 As shown in Table 6, our model is computationally efficient, runs on modest GPU memory (single
975 NVIDIA RTX 3090 GPU (24GB)), and delivers high-quality results.
976977 Table 6: Compute budgets.
978

Method	time
Inpainting	5s
Integration and enhancement	13s
Coarse mesh reconstruction	6s
Geometry and texture refinement	20s
Total	44s

985 A.8 VIEW-CONSISTENCY SCORING. & USER PREFERENCES.
986988 As shown in Table 7, we use MEt3R to measure the multiview inpainted images. The table shows
989 that our model outperforms other methods and is very close to the Ground Truth, further validating
990 its effectiveness. We also provide user studies to measure the reconstructed meshes as shown in
991 Table 8. 5 is the best score, 1 is the worst score. The results show our model outperforms other
992 methods.
993994 Table 7: View-consistency scoring.
995

Method	MEt3R
SD	0.44
Controlnet	0.53
Nerfiller	0.50
Pix2gestalt	0.41
Instant3dit	0.34
Ours	0.32
Ground Truth	0.29

1004 Table 8: User preferences.
1005

Method	geometry	texture
Open-LRM	1.9	2.1
InstantMesh	3.2	3.2
Unique3D	3.3	3.5
Direct3D	3.0	-
Trellis	3.1	3.0
Hunyuan3D	3.5	3.6
Ours	3.9	4.0

1015 A.9 COMPARISON WITH MVINPAINTER.
10161017 We include results for MVInpainter as shown in Table 9. Similar to other baselines, it is unable to
1018 accurately perceive the regions that require inpainting. This limitation is reasonable, as MVInpainter
1019 is specifically designed for object removal, which is inherently different from our task. Object
1020 removal typically involves eliminating an entire object, whereas our task focuses on completing
1021 partial regions of an object, leading to fundamentally different requirements and challenges.
10221023 A.10 HYPERPARAMETERS & ROBUSTNESS
10241025 Even though our pipeline includes a multi-stage process, our pipeline does not involve a large num-
ber of hyperparameters, making it relatively insensitive to hyperparameter choices.
1026

1026
1027
1028 **Table 9: Comparison with MVInpainter.**
1029
1030
1031

Method	PSNR	LPIPS	SSIM
MVInpainter	11.12	0.29	0.79
Ours	25.50	0.06	0.95

1032
1033 Across all experiments, we do not perform any instance-specific or object-specific hyperparameter
1034 tuning; instead, we simply adopt the default or officially recommended settings. This design choice
1035 enhances both the practicality and reproducibility of our method. Furthermore, the experimental re-
1036 sults in the ablation studies demonstrate that our model can consistently produce plausible outcomes
1037 in these standard settings.

1038 The key hyperparameters used in our pipeline are listed in Table 10. A higher Controlnet-tile
1039 Strength leads to inconsistent or misaligned results that do not match the preserved (visible) regions.
1040 In contrast, moderate Strength values (e.g., 0.25) reliably maintain alignment while enhancing detail
1041 quality.

1042
1043 **Table 10: Hyperparameters.**
1044

Method	Value
Multiview inpainting Inference timestep	50
Multiview inpainting CFG	5.0
Controlnet-tile Inference timestep	32
Controlnet-tile CFG	7.5
Controlnet-tile Strength	0.25
LRM	official setting
StableNormal	official setting

1053
1054

A.11 DISCUSSION FOR 3D CONSISTENCY

1055 Our method ensures 3D consistency in the inpainted regions through the following mechanisms:

1056 **Multi-View Constraints (Cross-View Attention).** Our approach enforces 3D consistency by lever-
1057 aging multi-view constraints, specifically cross-view attention, combined with the strong 3D prior
1058 of the base model, MVDream. The MVDream architecture utilizes a block structure that includes
1059 both cross-view attention and cross-attention mechanisms. Specifically, our Mask Self-perceiver is
1060 applied on the cross-attention layer to aggregate information from the incomplete images and the
1061 textual input. This information is then further processed through cross-view attention in subsequent
1062 blocks. The inpainting process thus implicitly requires the inpainted content across different views
1063 to agree in the latent space, which enforces 3D consistency.

1064 **Geometric Anchors in Incomplete Images.** While the masks do not contain geometric cues, the
1065 incomplete images themselves provide essential geometric context. These incomplete images act as
1066 geometric anchors, and the diffusion model synthesizes the missing regions in a way that ensures
1067 alignment with the visible portions across all views. If the inpainted region were to be inconsistent,
1068 it would contradict the visible regions from at least one viewpoint, which the model is trained to
1069 avoid such situations.

1070
1071
1072
1073
1074
1075
1076
1077
1078
1079