

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 FORMATTING INSTRUCTIONS FOR ICLR 2026 CONFERENCE SUBMISSIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

3D Gaussian Splatting (3DGS) has emerged as a powerful and efficient 3D representation for novel view synthesis. This paper extends 3DGS capabilities to inpainting, where masked objects in a scene are replaced with new contents that blend seamlessly with the surroundings. Unlike 2D image inpainting, 3D Gaussian inpainting (3DGI) is challenging in effectively leveraging complementary visual and semantic cues from multiple input views, as occluded areas in one view may be visible in others. To address this, we propose a method that measures the visibility uncertainties of 3D points across different input views and uses them to guide 3DGI in utilizing complementary visual cues. We also employ uncertainties to learn a semantic concept of scene without the masked object and use a diffusion model to fill masked objects in input images based on the learned concept. Finally, we build a novel 3DGI framework, VISTA, by integrating VISibility-uncerTainty-guided 3DGI with scene conceptuAl learning. VISTA generates high-quality 3DGS models capable of synthesizing artifact-free and naturally inpainted novel views. Furthermore, our approach extends to handling dynamic distractors arising from temporal object changes, enhancing its versatility in diverse scene reconstruction scenarios. We demonstrate the superior performance of our method over state-of-the-art techniques using two challenging datasets: the SPIIn-NeRF dataset, featuring 10 diverse static 3D inpainting scenes, and an underwater 3D inpainting dataset derived from UTB180, including fast-moving fish as inpainting targets.

1 INTRODUCTION

3D representation effectively models a scene and has the ability to synthesize new views of the scene (Barron et al., 2021; Mildenhall et al., 2021; Wang et al., 2021; Kerbl et al., 2023). 3D Gaussian splatting (3DGS) methods have been demonstrated as efficient and effective ways to represent the scene from a set of images taken from different viewpoints (Kerbl et al., 2023; Tang et al., 2023; Wu et al., 2024). Further, enabling editability of 3D scene representations is a cornerstone of technologies like augmented reality and virtual reality Tewari et al. (2022). 3D Gaussian inpainting task is one of the key editing techniques, aiming to replace specified objects with new contents that blend seamlessly with the surroundings. This capability allows us to: (1) *Remove objects from static scenes*: given multi-view images, we can create a 3D representation that generates novel views with specific objects removed and believably filled in (Figure 1 (Upper)). (2) *Clean up dynamic scenes*: for scenes with moving elements like fish in water (see Figure 1 (Bottom)), we can build a 3D representation excluding these transient objects, enabling clear, consistent novel view synthesis.

However, such an important task is non-trivial and the key challenge is how to leverage the complementary visual and semantic cues from multiple input views. Intuitively, for a synthesized view, the ideal approach is to replace the targeted erasure region with the occluded content, which naturally completes the inpainting. The key information for this process lies within the other view images, where the obscured areas may be visible from different angles. However, how to utilize multi-view information effectively is still an open question. State-of-the-art works first remove the targeted erasure region-related Gaussians and fill the regions via 2D image inpainting method (Ye et al., 2024; Wang et al., 2024), which, however, neglects the complementary cues from other views. The latest work (Liu et al., 2024) leverages depth maps of different views to involve the cross-view complementary cues implicitly. However, depth maps cannot fully represent complementary cues, such

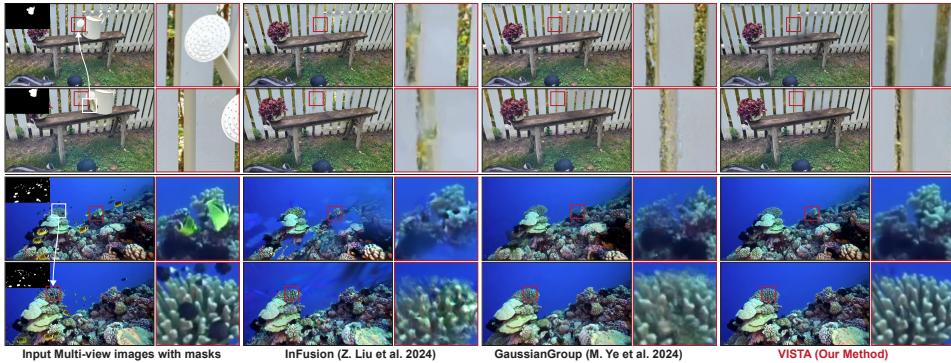


Figure 1: Two examples demonstrating two state-of-the-art methods, namely InFusion (Liu et al., 2024) and GaussianGroup (Ye et al., 2024), alongside our proposed method for 3D Gaussian inpainting to fill masked static and dynamic objects, respectively. The red boxes highlight the advantages of our method and are enlarged on the right side of each image for better visibility. The white boxes and arrows indicate complementary visual cues between two different viewpoints

as the texture pattern from adjacent perspectives, and the depth project can hardly get high-quality depth maps when moving objects across different views. As the two cases shown in Figure 1, InFusion synthesizes new views with obvious artifacts.

In this work, we propose VISibility-uncerTainty-guided 3DGI via scene conceptual Learning (VISTA), a novel framework for 3D Gaussian inpainting that leverages complementary visual and semantic cues. Our approach begins by measuring the visibility of 3D points across different views to generate visibility uncertainty maps for each input image. These maps indicate which pixels are most valuable for the inpainting task, based on the principle that pixels visible and consistent from multiple views contribute more significantly. We then integrate these visibility uncertainty maps into the 3D Gaussian splatting (3DGS) process. This enables the resulting Gaussian model to synthesize new views where masked regions are seamlessly filled with visual information from complementary perspectives. To address scenarios where large masked regions lack complementary visual cues from other views, we propose learning the concept of the scene without the masked objects. This conceptual learning is guided by the prior inpainting mask and the visibility uncertainty maps derived from the input multi-view images. The learned concept is then utilized to refine the input images, effectively filling the masked objects through a pre-trained Diffusion model. Furthermore, we implement an iterative process alternating between visibility-uncertainty-guided 3DGI and scene conceptual learning, progressively refining the 3D representation. As illustrated in Figure 1 (Upper), our method successfully reconstructs high-quality 3D representations of static scenes, naturally filling masked object regions with contextually appropriate content. Additionally, VISTA demonstrates its versatility by effectively removing distractors in dynamic scenes (see Figure 1 (Bottom) for examples).

We demonstrate the superior performance of our method over state-of-the-art techniques using two challenging datasets: the SPI-NeRF dataset, featuring 10 diverse static 3D in-painting scenes, and an underwater 3D inpainting dataset derived from UTB180, which includes fast-moving fish as inpainting targets. In summary, the contributions of our work are as follows:

1. We propose VISibility-uncerTAinty-guided 3D Gaussian inpainting (VISTA-GI) that explicitly leverages multi-view information through visibility uncertainty, achieving 3D Gaussian inpainting for more coherent and accurate scene completions.
2. We propose VISibility-uncerTAinty-guided scene conceptual learning (VISTA-CL) and leverage it for diffusion-based inpainting. VISTA-CL fills masked regions in input images using learned scene concepts, addressing the inpainting task at its core. This approach enhances the fundamental understanding of the scene, leading to more accurate and contextually appropriate inpainting results.
3. We introduce VISTA (VISibility-uncerTainty-guided 3D gaussian inpainTing via scene conceptual learning), a novel framework that iteratively combines VISTA-GI and VISTA-CL. This approach simultaneously leverages complementary visual and semantic cues, enhancing 3D Gaussian inpainting with geometric and conceptual information.

108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161

4. We extend VISTA to handle dynamic distractor removal in 3D Gaussian splatting, significantly improving its performance on scenes with temporal variations and outperforming state-of-the-art methods.

2 RELATED WORK

2.1 NERF AND 3D GAUSSIAN SPLATTING

Reconstructing 3D scenes from 2D images for novel view synthesis remains a fundamental challenge in computer vision and graphics (Lombardi et al., 2019; Kutulakos & Seitz, 2000). Recent advancements are driven by two contrasting paradigms: NeRF’s implicit neural representations (Mildenhall et al., 2021) and 3DGS’s explicit Gaussian modeling (Kerbl et al., 2023).

Neural Radiance Fields (NeRF) reconstruct scenes by learning continuous volumetric radiance fields through multi-view image optimization, achieving photorealistic view synthesis (Barron et al., 2021). However, its computational intensity in both training and rendering (Barron et al., 2022; 2023) limits practical deployment. In contrast, 3D Gaussian Splatting (3DGS) employs explicit geometric primitives - anisotropic Gaussian blobs with position, opacity, covariance, and color attributes (Lu et al., 2024). This approach achieves real-time rendering of high-fidelity novel views through direct spatial optimization of discrete 3D elements.

2.2 2D INPAINTING AND 3D INPAINTING

Image inpainting, a fundamental task in image generation, has evolved from traditional patch-based approaches (Ružić & Pižurica, 2014) to GAN-driven methods (Goodfellow et al., 2014; Yu et al., 2018), which handle regular regions but struggle with complex occlusions. Diffusion models now dominate this field (Ho et al., 2020; Sohl-Dickstein et al., 2015; Song et al., 2020), demonstrating superior capability in generating semantically coherent content for large missing areas (Lugmayr et al., 2022; Suvorov et al., 2022; Li et al., 2022).

Extending inpainting to 3D neural representations presents unique challenges. While NeRF-based approaches (Liu et al., 2022; Mirzaei et al., 2023; Weder et al., 2023) achieve partial success with static volumetric representations, their effectiveness remains constrained by inherent architectural limitations. Explicit 3DGS inpainting methods like Gaussian Grouping (Ye et al., 2024), InFusion (Liu et al., 2024), and GaussianEditor (Wang et al., 2024) primarily address static scene editing, overlooking dynamic interference during scene capture. Recent works attempt new directions: GScream (Wang et al., 2025) integrates monocular depth estimation with cross-attention for object removal, while SpotLessSplats (Sabour et al., 2024) addresses dynamic distractors through mask propagation - though both struggle with reconstructing heavily occluded regions.

3 PRELIMINARIES: 3D GAUSSIAN SPLATTING AND INPAINTING

3.1 3D GAUSSIAN SPLATTING

Given a set of images $\mathcal{I} = \{\mathbf{I}_i\}_{i=1}^N$ captured from various viewpoints and timestamps, 3D Gaussian splatting (3DGS) aims to learn a collection of anisotropic Gaussian splats $\mathcal{G} = \{\mathbf{g}_j\}_{j=1}^M$ from these multi-view images. Each splat \mathbf{g}_j is characterized by a Gaussian function with mean μ_j , a positive semi-definite covariance matrix Σ_j , an opacity α_j , and view-dependent color coefficients \mathbf{c}_j . Once the parameters of the 3D Gaussian splats \mathcal{G} are determined, novel view synthesis can be achieved through alpha-blending: $\hat{\mathbf{I}}^p = \text{Render}(\mathcal{G}, \mathbf{p})$. We can use \mathcal{I} to supervise the optimization of \mathcal{G}

$$\arg \min_{\mathcal{G}} \lambda_1 \sum_{i=1}^N \|(\mathbf{I}_i - \hat{\mathbf{I}}^{p_i})\|_1 + \lambda_2 \sum_{i=1}^N \text{D-SSIM}(\mathbf{I}_i, \hat{\mathbf{I}}^{p_i}), \quad (1)$$

where \mathbf{p}_i denotes the camera perspective of the image \mathbf{I}_i , $\hat{\mathbf{I}}^{p_i} = \text{Render}(\mathcal{G}, \mathbf{p}_i)$, and $\lambda_1 + \lambda_2 = 1$. For novel view synthesis, given a camera perspective \mathbf{p} , the process involves the following steps: projecting each 3D Gaussian onto a 2D image plane, sorting the Gaussians by depth along the view direction, and blending the Gaussians from front to back for each pixel. A key advantage of 3DGS

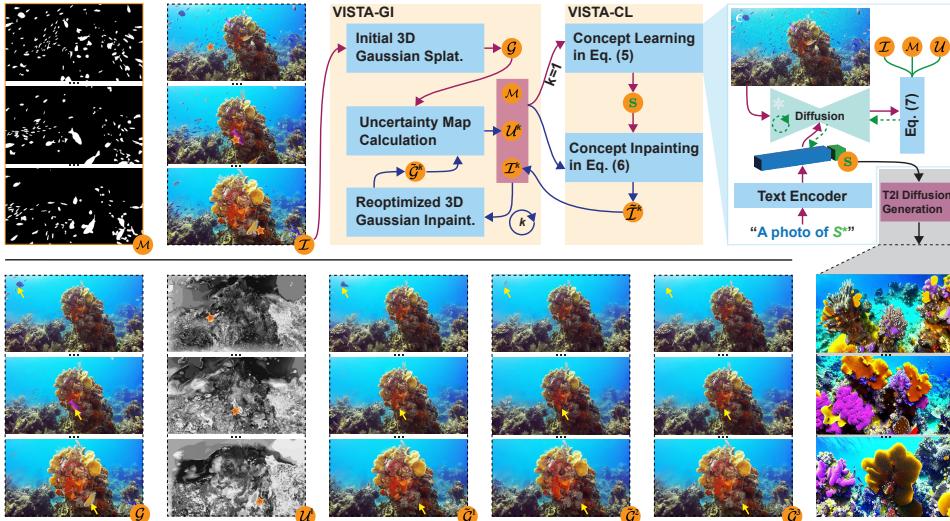


Figure 2: The framework of VISTA comprises two modules: VISTA-GI (described in Section 4.1) and VISTA-CL (detailed in Section 4.2). Results from three views are displayed for key variables in the framework. Note that \mathcal{G} , $\tilde{\mathcal{G}}^1$, $\tilde{\mathcal{G}}^2$, and $\tilde{\mathcal{G}}^3$ are 3DGS representations, and the displayed examples are rendered from these representations. The last column shows generated images derived from the learned scene concept. In the uncertainty map, we use \star to highlight areas of high uncertainty, which denote points (e.g., dynamic fishes) visible from only a few views. Yellow arrows demonstrate the progressive improvement in inpainting quality achieved by our method.

(Kerbl et al., 2023) is its ability to synthesize a new view in a single pass, whereas NeRF requires pixel-by-pixel rendering. This efficiency makes 3DGS particularly well-suited for time-sensitive 3D representation applications, offering a significant performance boost over NeRF.

3.2 3D GAUSSIAN INPAINTING

Given a set of captured images $\mathcal{I} = \{\mathbf{I}_i\}_{i=1}^N$ and corresponding binary mask maps $\mathcal{M} = \{\mathbf{M}_i\}_{i=1}^N$ delineating objects for removal (See Figure 1), 3D Gaussian Inpainting (3DGI) constructs a new 3D Gaussian splatting (3DGS) representation. This representation eliminates specified objects and replaces them with content that integrates with the environment. The resulting 3DGS representation can synthesize arbitrary views where the specified objects are imperceptibly absent, maintaining visual coherence across viewpoints while effectively 'erasing' targeted objects. We can use the segment anything model (SAM) (Kirillov et al., 2023) with few manual annotations to generate mask maps, aligning with methods like (Ye et al., 2024) for precise object delineation.

SOTA methods and limitations. An intuitive approach to 3D Gaussian Inpainting (3DGI) involves deriving a 3D mask for the specified objects based on the provided 2D masks. The process of new view synthesis then follows a two-step procedure: first, generating the specified view and its corresponding mask, and then applying existing 2D image inpainting techniques to achieve the desired 3DGI effect. This methodology has been adopted in recent works by Wang et al. (2024) and Ye et al. (2024). However, this approach does not leverage the complementary information available across multiple viewpoints during the inpainting process. A key example is the failure to utilize information from regions that may be occluded in one view but visible in another. Consequently, this method struggles to maintain consistency with the surrounding environment, particularly when dealing with large masked regions. This limitation underscores the need for more sophisticated techniques to effectively integrate and synthesize information from multiple perspectives to achieve more coherent and realistic 3D inpainting results. Beyond this solution, the latest work Liu et al. (2024) utilizes the cross-view complementary cues through depth perception. It formulates the 3D Gaussian inpainting as two tasks, *i.e.*, 2D image inpainting and depth inpainting, and the complementary cues in multiple views are implicitly utilized via depth projection. However, depth maps cannot fully represent complementary cues, such as the texture pattern from adjacent perspectives, and the depth project can hardly get high-quality depth maps when moving objects across different views. As case 2 shown in Figure 1, InFusion synthesizes new views with obvious artifacts.

216 **4 METHODOLOGY**
 217

218 This section details the proposed framework called VISibility-uncerTainty-guided 3D Gaussian in-
 219 painting via scene concepTional learning (VISTA). The core principle is to identify the visibility
 220 of 3D points across different views and utilize this information to guide the use of complementary
 221 visual and semantic cues for 3D Gaussian inpainting.

222 To elucidate this concept, we introduce the visibility-uncertainty-guided 3D Gaussian inpainting
 223 (VISTA-GI) in Section 4.1, where we define the visibility uncertainty of 3D points and employ it
 224 to guide the use of complementary visual cues for 3DGI. In Section 4.2, we propose leveraging
 225 the visibility uncertainty to learn the semantic concept of the scene without specified objects. We
 226 then perform concept-driven Diffusion inpainting to process the input images, harnessing comple-
 227 mentary semantic cues. To fully utilize complementary visual and semantic cues, we propose in
 228 Section 4.3 an iterative combination of VISTA-GI and VISTA-CL. Finally, in Section 4.4, we ex-
 229 tend our VISTA framework to address the challenge of dynamic distractors in captured images. This
 230 extension excludes transient objects, resulting in clearer and more consistent novel view synthesis.

231
 232 **4.1 VISTA-GI: VISIBILITY-UNCERTAINTY-GUIDED 3D GAUSSIAN INPAINTING**
 233

234 **Initial 3D Gaussian Splatting.** Given the input images $\mathcal{I} = \{\mathbf{I}_i\}_{i=1}^N$, we employ the original 3DGS
 235 method in Section 3.1 and Equation (1) to construct a 3D representation \mathcal{G} . This representation can
 236 then be utilized to render novel views. However, as illustrated in Figure 2, this initial representation
 237 fails to exclude dynamic objects (such as fish) and exhibits noticeable artifacts, including blurring.

238 **Visibility uncertainty of 3D Points.** We define a set of adjacent camera views (indicating the top-K
 239 nearest cameras, determined by the distance of camera centers), denoted as $\mathcal{P} = \{\mathbf{p}_v\}_{v=1}^V$, where V
 240 is the number of adjacent views. For a 3D point \mathbf{X} in the scene, we can project it to different camera
 241 perspectives in \mathcal{P} via the built 3DGS \mathcal{G} and get their colors under V views, *i.e.*, $\{\mathbf{x}_v\}_{v=1}^V$. Then, we
 242 calculate the variations of the colors of the point under different views

243
$$u_{\mathbf{x}} = \text{var}(\{\mathbf{x}_v\}_{v=1}^V), \quad (2)$$

 244

245 where $\text{var}(\cdot)$ is the variation function. We denote the result $u_{\mathbf{x}}$ as the *visibility uncertainty* of the 3D
 246 point \mathbf{X} . Intuitively, $u_{\mathbf{x}}$ represents the visibility and consistency of the point across the V views.
 247 For example, if the point \mathbf{X} can be seen at all views, the colors under different views are consistent
 248 and $u_{\mathbf{x}}$ is small. If the point can be only seen by a few views or its color deviates between different
 249 views, the visibility uncertainty tends to be significantly high.

250 **Reoptimized 3D Gaussian inpainting.** With the 3D point’s visibility uncertainty, we aim to cal-
 251 culate the visibility uncertainty map of the input image and measure the visibility of each pixel at
 252 other views. Specifically, for an image \mathbf{I}_i in \mathcal{I} , we first calculate its depth map \mathbf{D}_i based on the
 253 \mathcal{G} . Then, we project each pixel of \mathbf{I}_i to a 3D point and calculate its visibility uncertainty via Equa-
 254 tion (2) under V adjacent views. Then, we obtain a pixel-wise visibility uncertainty map, which is
 255 normalized by dividing each pixel’s uncertainty value by the standard deviation computed across all
 256 uncertainty values. The resulting normalized map is denoted as \mathbf{U}_i . For the N input images, we
 257 have N visibility uncertainty maps $\mathcal{U} = \{\mathbf{U}_i\}_{i=1}^N$. Then, we use them to update the original mask
 258 maps \mathcal{M} and uncertainty maps \mathcal{U} by

259
$$\mathbf{M}'_i = \mathbf{U}_i \odot (1 - \mathbf{M}_i) + \vartheta \cdot \mathbf{M}_i, \quad (3)$$

260 where the first term weights the unmasked regions via the visibility uncertainty map: the points other
 261 views cannot see should be assigned low weights during optimization. The ϑ controls the constraint
 262 degrees of the original masks. Then, we obtain the finer mask maps $\{\mathbf{M}'_i\}_{i=1}^N$ and re-optimize the
 263 3D representation by adding the guidance of mask maps to the objective function in Equation (1):

264
$$\arg \min_{\mathcal{G}} \lambda_1 \sum_{i=1}^N \|(1 - \mathbf{M}'_i) \odot (\mathbf{I}_i - \hat{\mathbf{I}}^{p_i})\|_1 + \lambda_2 \sum_{i=1}^N \text{D-SSIM}(\mathbf{I}_i, \hat{\mathbf{I}}^{p_i}, 1 - \mathbf{M}'_i), \quad (4)$$

 265

266 where we have $\hat{\mathbf{I}}^{p_i} = \text{Render}(\mathcal{G}, \mathbf{p}_i)$ and $\lambda_1 + \lambda_2 = 1$. Intuitively, the objective function is to ignore
 267 the mask and high-uncertainty regions during the optimization. As a result, we get an updated
 268 counterpart $\tilde{\mathcal{G}}$. Similar strategies have been also adopted in recent works (Sabour et al., 2024; 2023).

270 Intuitively, with the visibility uncertainty maps, we can exclude the pixels that other views cannot
 271 see to build the 3D representation, which explicitly leverages the complementary visual cues. As
 272 the \mathcal{U} shown in Figure 2 (Bottom), the pixels with high uncertainty denote the corresponding points
 273 (e.g., dynamic fishes) visible from only a few views. This is reasonable since the dynamic fishes
 274 are at different locations across different views. We also display the updated 3D representation $\tilde{\mathcal{G}}^1$,
 275 showing that the dynamic objects and some artifacts are removed.
 276

277 4.2 VISTA-CL: VISIBILITY-UNCERTAINTY-GUIDED SCENE CONCEPTUAL LEARNING

279 VISTA-GI can reconstruct masked objects when complementary visual information is available from
 280 alternative viewpoints. However, for masked regions lacking such cues, we need a more sophisticated
 281 approach to comprehend the scene holistically and generate plausible new content to fill these
 282 gaps. To achieve this, we propose to learn a conceptual representation \mathbf{s} of the scene through textual
 283 inversion (Gal et al., 2022; Zhu et al., 2024), which can be formulated as
 284

$$\mathbf{s} = \text{ConceptLearn}(\mathcal{I}, \mathcal{U}, \mathcal{M}), \quad (5)$$

285 The learned concept \mathbf{s} is a token and encapsulates the scene’s essence without the masked objects.
 286 We then leverage \mathbf{s} to process the input images, eliminating the masked objects

$$\tilde{\mathbf{I}}_i = \text{ConceptInpaint}(\mathbf{s}, \mathbf{I}_i, \mathcal{U}, \mathcal{M}), \forall \mathbf{I}_i \in \mathcal{I}, \quad (6)$$

289 **Scene conceptual learning.** We formulate the scene conceptual learning, *i.e.*, as the personalization
 290 text-to-image problem (Ruiz et al., 2023) based on textual inversion (Gal et al., 2022), and we add
 291 the guidance of the visibility uncertainty maps in Section 3.2. Specifically, we have a pre-trained
 292 text-2-image diffusion model containing an image autoencoder with ϕ and ϕ^{-1} as encoder and
 293 decoder, a text encoder φ , and a conditional diffusion model ϵ_θ at latent space. Then, we learn the
 294 scene concept \mathbf{s} by optimizing the following objective function
 295

$$\mathbf{s} = \arg \min_{\mathbf{s}^*} \mathbb{E}_{\mathbf{I}_i \in \mathcal{I}, \mathbf{z} = \phi(\mathbf{I}_i), \mathbf{y}, \epsilon \in \mathcal{N}(0, 1), t} (\|(1 - \mathbf{M}'_i) \odot (\epsilon_\theta(\mathbf{z}_t, t, \Upsilon(\varphi(\mathbf{y}), \mathbf{s}^*)) - \epsilon)\|_2^2), \quad (7)$$

296 where \mathbf{y} is a fixed text (*i.e.*, ‘a photo of S^* ’) and the function $\Upsilon(\Gamma(\mathbf{y}), \mathbf{s}^*)$ is to replace the token
 297 of ‘ S^* ’ within $\Gamma(\mathbf{y})$ with \mathbf{s}^* . The tensor \mathbf{M}'_i is calculated via Equation (3) based on the visibility
 298 uncertainty map and the given mask map. Intuitively, we use the Equation (7) to force the learned
 299 concept to mainly contain the unmasked scene regions. To validate the learned concept, we can
 300 feed ‘a photo of S^* ’ to the T2I diffusion model to generate images about the learned concept. As
 301 shown in Figure 2, the images in the lower right are created directly by the T2I diffusion model and
 302 illustrate a concept similar to the original scene without any dynamic objects.

303 **Scene conceptual-guided inpainting.** We use the learned concept \mathbf{s} to inpaint all input images
 304 through the pre-trained T2I diffusion model. Given one image \mathbf{I} from \mathcal{I} , we can extract its latent
 305 code by $\mathbf{z} = \phi(\mathbf{I})$. Then, we perform the forward diffusion process by iteratively adding Gaussian
 306 noise to the \mathbf{z} over T timesteps, obtaining a sequence of noisy latent codes, *i.e.*, $\mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_T$,
 307 where $\mathbf{z}_0 = \mathbf{z}$. At the t th step, the latent is obtained by
 308

$$q(\mathbf{z}_t | \mathbf{z}_0) = \sqrt{\bar{\alpha}_t} \mathbf{z}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \mathbb{I}), \quad (8)$$

309 where $\bar{\alpha}_t = \prod_{\tau=1}^t (1 - \beta_\tau)$. $\mathcal{N}(0, \mathbb{I})$ represents the standard Gaussian distribution. As we set the time
 310 step as T , the complete forward process can be expressed as $\mathbf{z}_T \sim q(\mathbf{z}_{1:T} | \mathbf{z}_0) = \prod_{t=1}^T q(\mathbf{z}_t | \mathbf{z}_{t-1})$.
 311

312 At the reverse denoising process, we follow the strategy of RePaint (Lugmayr et al., 2022) but embed
 313 the guidance of visibility uncertainty maps and the learned concept \mathbf{s} . Intuitively, at the time step
 314 $t > 1$ during denoising, we only denoise the masked regions conditioned on the scene concept \mathbf{s}
 315 while maintaining the unmasked regions with the same content in Equation (8), that is, we have
 316

$$\tilde{\mathbf{z}}_{t-1} = (1 - \mathbf{m}') \odot \mathbf{z}_{t-1} + \mathbf{m}' \odot \hat{\mathbf{z}}_{t-1}, \quad (9)$$

317 where $\mathbf{z}_{t-1} \sim q(\mathbf{z}_t | \mathbf{z}_0)$ and \mathbf{m}' is the downsampled $\mathbf{M}' \in \{\mathbf{M}'_i\}_{i=1}^N$ calculated by Equation (3) and
 318 has the exact resolution as the latent code \mathbf{z}_{t-1} . $\hat{\mathbf{z}}_{t-1}$ is denoised from the $\tilde{\mathbf{z}}_t$ with the guidance of
 319 the learned concept \mathbf{s} , that is,

$$\hat{\mathbf{z}}_{t-1} = \frac{1}{\sqrt{\alpha_t}} (\tilde{\mathbf{z}}_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\tilde{\mathbf{z}}_t, t, \mathbf{s})) + \sigma_t \xi, \text{s.t., } \xi \sim \mathcal{N}(0, \mathbb{I}), \quad (10)$$

320 If $t = 1$, $\tilde{\mathbf{z}}_0 = (1 - \mathbf{m}') \odot \mathbf{z} + \mathbf{m}' \cdot \hat{\mathbf{z}}_0$. Then, we can get the inpainted image via decoder $\tilde{\mathbf{I}} = \phi^{-1}(\tilde{\mathbf{z}}_0)$.
 321 We can use the above ConceptInpaint to process each image within \mathcal{I} and get a new image set $\tilde{\mathcal{I}}$.
 322

324 4.3 VISTA: COMBINING VISTA-GI AND VISTA-CL
325

326 Given the input images \mathcal{I} and their corresponding mask maps \mathcal{M} , VISTA-GI generates the uncer-
327 tainty maps \mathcal{U} as the visual cues and refines the 3DGS representation $\tilde{\mathcal{G}}$. VISTA-CL takes \mathcal{I} , \mathcal{U} ,
328 and \mathcal{M} as inputs and produces processed input images $\tilde{\mathcal{I}}$ as the semantic cues. Intuitively, we can
329 combine the raw images \mathcal{I} and $\tilde{\mathcal{I}}$, feed them back into VISTA-GI, where $\tilde{\mathcal{I}}$ serve as better views.
330 This allows for iterative process between VISTA-GI and VISTA-CL. We denote the k -th iteration's
331 3D representation from VISTA-GI as $\tilde{\mathcal{G}}^k$ and the processed images from VISTA-CL as $\tilde{\mathcal{I}}^k$.
332

333 In practice, three iterations are typically sufficient to achieve smooth convergence of the training
334 metrics. The hyperparameter ϑ is initialized by 0 and increases by 0.1 with each iteration. We show
335 an example in Figure 2. The synthetic views $\tilde{\mathcal{G}}^1$, $\tilde{\mathcal{G}}^2$, and $\tilde{\mathcal{G}}^3$ gradually contain fewer distractors,
336 and the results of the final iteration $\tilde{\mathcal{G}}^3$ demonstrate clean and clear views, which means better 3D
337 inpainting under the guidance of the visual and semantic cues.
338

339 4.4 VISTA FOR DYNAMIC DISTRACTOR REMOVAL

340 VISTA could be easily extended to remove dynamic distractors across multi-view images \mathcal{I} by
341 identifying the dynamic regions in \mathcal{I} and obtaining the mask maps \mathcal{M} . In our implementation,
342 we use the tracking method and MASA (Li et al., 2024) to automatically get the mask maps for
343 dynamic objects in the scene. MASA is an open-vocabulary video detection and segmentation model
344 introducing coarse pixel-level information to our method. This plays a similar role as DEVA (Cheng
345 et al., 2023) used in Gaussian Grouping (Ye et al., 2024). However, the masks used in Gaussian
346 Grouping are limited to static objects, while we mask static and dynamic objects that need to be
347 inpainted. For dynamic objects, the uncertainty map can complement the coarse mask that excludes
348 those dynamic distractors from the reconstruction. As shown in Figure 2, the synthetic view \mathcal{G}
349 obtained without masks fairly removes those fish moving greatly but ignores those objects without
350 significant movement. The semantic information in the coarse masks \mathcal{M} identifies these distractors,
351 which the uncertainty map \mathcal{U} cannot detect, and then these distractors can be eliminated by VISTA-
352 CL. As a result, VISTA can remove both static and dynamic distractors in the scene by combining
353 these two mask maps in Equation (3).
354

355 5 EXPERIMENTS

356 5.1 DATASETS AND METRICS
357

358 To evaluate our method, we conduct experiments on the SPIn-NeRF Dataset and the Underwater 3D
359 Inpainting Dataset for scene repair in different scenarios.
360

361 **Underwater 3D inpainting dataset.** This
362 dataset is derived from the underwater ob-
363 ject tracking dataset UTB180 (Alawode et al.,
364 2022), from which we selected multiple videos
365 for resampling, ultimately forming 10 under-
366 water scene datasets. We resample the video at
367 a certain FPS to fulfill the motion requirements
368 of the initial reconstruction. Each scene con-
369 tains dozens of images from various viewpoints,
370 and the initial Structure from Motion point cloud
371 and camera intrinsics are obtained via COLMAP
372 (Schonberger & Frahm, 2016). Each viewpoint
373 image undergoes object detection using the open-source method MASA (Li et al., 2024) to obtain
374 rough object masks.
375

376 **SPIn-NeRF dataset.** The SPIn-NeRF dataset
377 was proposed in Mirzaei et al. (2023). It con-
378 tains 10 general 3D inpainting scenes, divided
379 into 3 indoor and 7 outdoor scenes. Each
380 scene includes 100 images from various view-
381 points, along with corresponding masks. In
382 these datasets, the ratio of the training set to the
383

Method	UCIQE \uparrow	URanker \uparrow	CLIP Score \uparrow	Rank-1 \uparrow
SPIn-NeRF	0.49	1.59	0.70	15.89 %
InFusion	0.50	1.52	0.71	3.39 %
SpotLess	0.50	1.59	0.70	9.63 %
Ours	0.51	1.64	0.72	71.09 %

384 Table 1: Results of dynamic inpainting on the Un-
385 derwater 3D Inpainting Dataset.
386

387 and the initial Structure from Motion point cloud
388 and camera intrinsics are obtained via COLMAP
389 (Schonberger & Frahm, 2016). Each viewpoint
390 image undergoes object detection using the open-source method MASA (Li et al., 2024) to obtain
391 rough object masks.
392

Method	LPIPS \downarrow	Fid \downarrow	PSNR \uparrow	SSIM \uparrow
Masked Gaussians	0.594	278.32	10.77	0.29
SPIn-NeRF	0.465	156.64	15.80	0.46
Gaussians Grouping	0.454	123.48	14.86	0.27
GScream	0.422	114.41	15.98	0.59
InFusion	0.567	118.26	15.59	0.53
Ours	0.418	113.58	16.48	0.59

393 Table 2: Static inpainting for SPIn-NeRF Dataset.
394

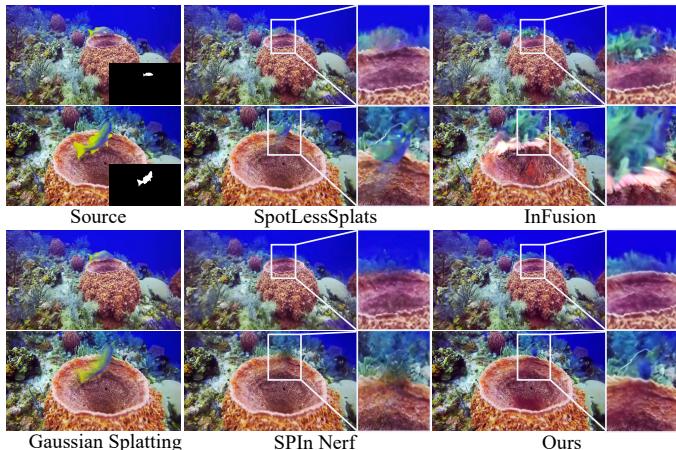
378 testing set is 6 to 4. We compare our method with other approaches using the provided camera
 379 intrinsics and an initialized SfM point cloud.
 380

381 **Metrics.** Following SPIIn-NeRF, we evaluate results using two quantitative protocols: (1) for static
 382 scenes with ground truth, we adopt PSNR, SSIM, LPIPS, and FID for reference-based image quality
 383 assessment (IQA); (2) for dynamic scenes without ground truth, we employ UCIQE (Yang &
 384 Sowmya, 2015), URanker (Guo et al., 2023), and CLIP Score (Hessel et al., 2021) for underwater
 385 non-reference IQA. Consistent with SPIIn-NeRF (Mirzaei et al., 2023) and RefFusion (Mirzaei et al.,
 386 2024), LPIPS and FID are computed within the masked region defined by its bounding box. UCIQE
 387 combines chroma, saturation, and contrast to quantify color casts, blurriness, and low contrast;
 388 URanker is a transformer-based underwater image quality metric; CLIP Score measures image–text
 389 alignment, using the caption “*An underwater scene without fish*” to assess fish removal. To address
 390 the insensitivity of non-reference metrics, we conducted a user study on 48 images randomly sam-
 391 pled from six scenes. For each case, participants compared renderings from all methods and selected
 392 the one with the highest visual quality. The *Rank-1* column reports the proportion of cases in which
 393 a method was chosen as the top-performing approach.
 394

395 5.2 EXPERIMENTAL RESULTS

396 We compare our method with several state-of-the-art open-source 3D inpainting methods, such as
 397 Infusion (Liu et al., 2024), SPIIn-NeRF (Mirzaei et al., 2023), Gaussian Grouping (Ye et al., 2024),
 398 and SpotLessSplats (Sabour et al., 2024). SpotLessSplats is only designed for scenarios with dy-
 399 namic distractors, while others are the latest static inpainting methods. Infusion (Liu et al., 2024)
 400 and GScream (Wang et al., 2025) are retrained using their publicized code. Further ablation studies
 401 and detailed discussions are presented in Appendix A.

402 **Results on the underwater**
 403 **3D inpainting dataset.** Our
 404 method demonstrates superior
 405 underwater inpainting perfor-
 406 mance in Figure 3, achieving
 407 artifact-free 3D consistency
 408 across dynamic objects where
 409 baselines fail: SpotLessSplats
 410 incompletely removes moving
 411 fish Gaussians; Infusion induces
 412 viewpoint distortions despite
 413 localized clarity; SPIIn-NeRF
 414 maintains 3D coherence but
 415 exhibits artifacts. Quantitative
 416 validation in Table 1 reveals our
 417 approach surpasses competitors
 418 in UCIQE and URanker metrics
 419 through uncertainty-guided
 420 reconstruction, while superior CLIP
 421 Scores confirm effective target object
 422 removal through optimized
 423 semantic-geometric integration. The 71.09% rank-1 preference—significantly higher than
 424 baselines—aligns with visual evidence (e.g., Figure 3) and underscores our method’s effectiveness
 425 in perceptually critical areas. While standard non-reference metrics indicate modest gains, human
 426 evaluation confirms that our method delivers the best reconstruction in structural coherence and
 427 detail recovery.



428 Figure 3: Visualization for dynamic inpainting.
 429



430 Figure 4: Visualized examples of static inpainting on SPIIn-NeRF.
 431

432 Results on SPIn-NeRF dataset. Figure 4 depicts an example scene from the SPIn-NeRF Dataset
 433 masking a stationary box that requires inpainting. The results of Gaussian Grouping are fairly re-
 434 alistic at the 2D image level, but there are significant inconsistencies between perspectives, such as
 435 distortion at the edges of stairs. The results of InFusion appear more realistic from a certain per-
 436 spective. Still, its approach of optimizing one single view compromises the performance of other
 437 perspectives, leading to unpredictable artifacts in those views. As shown in Table 2, our method ben-
 438 efits from an iterative progressive optimization approach, ensuring consistency across perspectives
 439 through multiple inpainting and reconstruction, resulting in more stable outcomes.

440 Ablation study on VISTA-GI and
441 VISTA-CL. We conducted the ablation

Method	UCIQE \uparrow	URanker \uparrow	CLIP Score \uparrow
Ours w/o VISTA-GI	0.48	1.52	0.70
Ours w/o VISTA-CL	0.50	1.59	0.69
Ours	0.51	1.64	0.72

Table 3: Ablation study of VISTA-GI and VISTA-CL on the Underwater 3D Inpainting Dataset.

442
443
444
445
446
447 The specific results are shown in Table 3.
448 Our experiments reveal two main findings:
449 1) Using only a 2D generative model with-
450 out VISTA-GI results in poor metrics, confirming that VISTA-GI’s uncertainty guidance reduces
451 multi-view inconsistencies in 3D reconstruction, leading to better outputs. 2) Omitting VISTA-CL
452 maintains image quality similar to methods like SplotLess and SPIn-NeRF, but significantly lowers
453 CLIP-Score metrics. This shows that without conceptual constraints, the inpainting process is
454 visually plausible but semantically inconsistent with the scene context.

455 Impact of noise reduction ratios in
456 diffusion. We analyze noise reduction
457 strategies during diffusion model infer-
458 ence by implementing fixed-ratio iter-
459 ative noise reduction from 1.0, evaluating
460 four ratios $\{0.1, 0.2, 0.3, 0.4\}$ on recon-
461 struction quality. As shown in Figure 5
462 (a), all ratios improve PSNR iteratively,
463 but the 0.2 ratio achieves optimal con-
464 vergence fastest. Empirical validation
465 confirms 0.2 as our method’s optimal re-
466 duction ratio.

467 Impacts of ϑ in equation 3. We analyze mask prior constraint control via ϑ increase strate-
468 gies (0.1-0.5), initialized at $\vartheta = 0$ with 0.1 per-iteration increments. Figure 5 (b) demon-
469 strates iterative PSNR improvements across all ratios. While higher ratios enhance early-stage recon-
470 struction, values bigger than 0.1 trigger premature inpainting confidence, causing insufficient VISTA-GI/CL
471 module interaction and subsequent performance degradation.

472 Comparisons of different methods in extreme cases.

To validate our advantages in the extreme case with
 473 large viewpoint differences (sparse views), we quanti-
 474 tatively evaluate various methods for the extreme case,
 475 and the results are as Table 4. Our method still outper-
 476 forms existing methods in removing dynamic distrac-
 477 tors under such extreme conditions.

6 CONCLUSION

In this work, we presented VISTA, a novel framework for 3D Gaussian inpainting that leverages complementary visual and semantic cues from multiple input views. By introducing visibility uncer-
 478 tainty maps and integrating visibility-uncertainty-guided inpainting (VISTA-GI) with scene concep-
 479 tual learning (VISTA-CL), our method addresses key challenges in both static and dynamic 3D scene
 480 editing. Experiments on the SPIn-NeRF and UTB180-derived datasets demonstrate that VISTA out-
 481 performs state-of-the-art methods, producing high-quality 3D reconstructions with seamlessly filled
 482 masked regions and effective distractor removal. Its ability to handle complex inpainting scenarios
 483 and dynamic objects makes it a versatile tool for augmented and virtual reality, advancing the pursuit
 484 of seamless, realistic 3D scene editing.

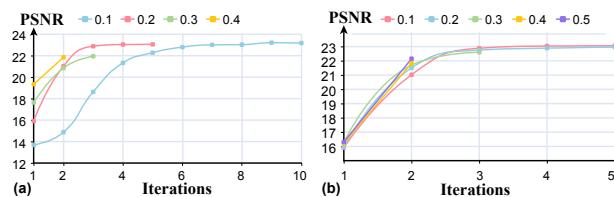


Figure 5: (a) Relationship between 3DGS rendering quality and noise reduction ratios in diffusion. (b) Relationship between 3DGS rendering quality and increasing ratio of ϑ .

Method	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow
InFusion	0.23	19.34	0.78
SPIn-NeRF	0.15	23.33	0.82
SpotLess	0.14	24.75	0.84
Ours	0.10	26.38	0.86

Table 4: Comparison of different methods in the extreme case.

486 REFERENCES
487

488 Basit Alawode, Yuhang Guo, Mehnaz Ummar, Naoufel Werghi, Jorge Dias, Ajmal Mian, and Sajid
489 Javed. Utb180: A high-quality benchmark for underwater tracking. In *Proceedings of the Asian
490 Conference on Computer Vision*, pp. 3326–3342, 2022.

491 Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and
492 Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields.
493 In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5855–5864,
494 2021.

495 Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf
496 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF conference
497 on computer vision and pattern recognition*, pp. 5470–5479, 2022.

498 Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-nerf:
499 Anti-aliased grid-based neural radiance fields. In *Proceedings of the IEEE/CVF International
500 Conference on Computer Vision*, pp. 19697–19705, 2023.

502 Ho Kei Cheng, Seoung Wug Oh, Brian Price, Alexander Schwing, and Joon-Young Lee. Tracking
503 anything with decoupled video segmentation. In *Proceedings of the IEEE/CVF International
504 Conference on Computer Vision*, pp. 1316–1326, 2023.

505 Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel
506 Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual
507 inversion. *arXiv preprint arXiv:2208.01618*, 2022.

509 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
510 Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information
511 processing systems*, 27, 2014.

512 Chunle Guo, Ruiqi Wu, Xin Jin, Linghao Han, Weidong Zhang, Zhi Chai, and Chongyi Li. Under-
513 water ranker: Learn which is better and how to be better. In *Proceedings of the AAAI conference
514 on artificial intelligence*, volume 37, pp. 702–709, 2023.

515 Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A
516 reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021.

518 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
519 neural information processing systems*, 33:6840–6851, 2020.

520 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splat-
521 ting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.

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

526 Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. *International journal
527 of computer vision*, 38:199–218, 2000.

529 Siyuan Li, Lei Ke, Martin Danelljan, Luigi Piccinelli, Mattia Segu, Luc Van Gool, and Fisher Yu.
530 Matching anything by segmenting anything. In *Proceedings of the IEEE/CVF Conference on
531 Computer Vision and Pattern Recognition*, pp. 18963–18973, 2024.

532 Wenbo Li, Zhe Lin, Kun Zhou, Lu Qi, Yi Wang, and Jiaya Jia. Mat: Mask-aware transformer for
533 large hole image inpainting. In *Proceedings of the IEEE/CVF conference on computer vision and
534 pattern recognition*, pp. 10758–10768, 2022.

535 Hao-Kang Liu, I Shen, Bing-Yu Chen, et al. Nerf-in: Free-form nerf inpainting with rgb-d priors.
536 *arXiv preprint arXiv:2206.04901*, 2022.

538 Zhiheng Liu, Hao Ouyang, Qiuyu Wang, Ka Leong Cheng, Jie Xiao, Kai Zhu, Nan Xue, Yu Liu,
539 Yujun Shen, and Yang Cao. Infusion: Inpainting 3d gaussians via learning depth completion from
diffusion prior. *arXiv preprint arXiv:2404.11613*, 2024.

540 Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser
 541 Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint*
 542 *arXiv:1906.07751*, 2019.

543 Tao Lu, Mulin Yu, Lining Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-gs:
 544 Structured 3d gaussians for view-adaptive rendering. In *Proceedings of the IEEE/CVF Conference*
 545 *on Computer Vision and Pattern Recognition*, pp. 20654–20664, 2024.

546 Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool.
 547 Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the*
 548 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 11461–11471, 2022.

549 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
 550 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications*
 551 *of the ACM*, 65(1):99–106, 2021.

552 Ashkan Mirzaei, Tristan Amentado-Armstrong, Konstantinos G Derpanis, Jonathan Kelly, Mar-
 553 cus A Brubaker, Igor Gilitschenski, and Alex Levinstein. Spin-nerf: Multiview segmentation
 554 and perceptual inpainting with neural radiance fields. In *Proceedings of the IEEE/CVF Confer-*
 555 *ence on Computer Vision and Pattern Recognition*, pp. 20669–20679, 2023.

556 Ashkan Mirzaei, Riccardo De Lutio, Seung Wook Kim, David Acuna, Jonathan Kelly, Sanja Fidler,
 557 Igor Gilitschenski, and Zan Gojcic. Reffusion: Reference adapted diffusion models for 3d scene
 558 inpainting. *arXiv preprint arXiv:2404.10765*, 2024.

559 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 560 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-*
 561 *ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.

562 Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman.
 563 Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Pro-*
 564 *ceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 22500–
 565 22510, 2023.

566 Tijana Ružić and Aleksandra Pižurica. Context-aware patch-based image inpainting using markov
 567 random field modeling. *IEEE transactions on image processing*, 24(1):444–456, 2014.

568 Sara Sabour, Suhani Vora, Daniel Duckworth, Ivan Krasin, David J Fleet, and Andrea Tagliasacchi.
 569 Robustnerf: Ignoring distractors with robust losses. In *Proceedings of the IEEE/CVF Conference*
 570 *on Computer Vision and Pattern Recognition*, pp. 20626–20636, 2023.

571 Sara Sabour, Lily Goli, George Kopanas, Mark Matthews, Dmitry Lagun, Leonidas Guibas, Alec
 572 Jacobson, David J. Fleet, and Andrea Tagliasacchi. SpotLessSplats: Ignoring distractors in 3d
 573 gaussian splatting. *arXiv:2406.20055*, 2024.

574 Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings*
 575 *of the IEEE conference on computer vision and pattern recognition*, pp. 4104–4113, 2016.

576 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
 577 learning using nonequilibrium thermodynamics. In *International conference on machine learn-*
 578 *ing*, pp. 2256–2265. PMLR, 2015.

579 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 580 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint*
 581 *arXiv:2011.13456*, 2020.

582 Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha,
 583 Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky.
 584 Resolution-robust large mask inpainting with fourier convolutions. In *Proceedings of the*
 585 *IEEE/CVF winter conference on applications of computer vision*, pp. 2149–2159, 2022.

586 Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative
 587 gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023.

594 Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Wang Yifan,
 595 Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. Advances
 596 in neural rendering. In *Computer Graphics Forum*, volume 41, pp. 703–735. Wiley Online
 597 Library, 2022.

598 Junjie Wang, Jiemin Fang, Xiaopeng Zhang, Lingxi Xie, and Qi Tian. Gaussianeditor: Editing
 599 3d gaussians delicately with text instructions. In *Proceedings of the IEEE/CVF Conference on*
 600 *Computer Vision and Pattern Recognition*, pp. 20902–20911, 2024.

601 Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus:
 602 Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv*
 603 *preprint arXiv:2106.10689*, 2021.

604 Yuxin Wang, Qianyi Wu, Guofeng Zhang, and Dan Xu. Learning 3d geometry and feature consistent
 605 gaussian splatting for object removal. In *European Conference on Computer Vision*, pp. 1–17.
 606 Springer, 2025.

607 Silvan Weder, Guillermo Garcia-Hernando, Aron Monszpart, Marc Pollefeys, Gabriel J Brostow,
 608 Michael Firman, and Sara Vicente. Removing objects from neural radiance fields. In *Proceedings*
 609 *of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16528–16538,
 610 2023.

611 Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian,
 612 and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings*
 613 *of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20310–
 614 20320, June 2024.

615 Miao Yang and Arcot Sowmya. An underwater color image quality evaluation metric. *IEEE Transactions*
 616 *on Image Processing*, 24(12):6062–6071, 2015.

617 Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit
 618 anything in 3d scenes. In *European Conference on Computer Vision*, 2024.

619 Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image
 620 inpainting with contextual attention. In *Proceedings of the IEEE conference on computer vision*
 621 *and pattern recognition*, pp. 5505–5514, 2018.

622 Jiayi Zhu, Qing Guo, Felix Juefei-Xu, Yihao Huang, Yang Liu, and Geguang Pu. Cosalpure: Learning
 623 concept from group images for robust co-saliency detection. In *Proceedings of the IEEE/CVF*
 624 *Conference on Computer Vision and Pattern Recognition*, pp. 3669–3678, 2024.

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648
649

A APPENDIX

650
651
652
653
654

Use of Large Language Models (LLMs) We acknowledge the use of a large language model
as a general-purpose assistive tool. Its role was limited to language polishing and minor stylistic
refinement of the manuscript. The LLM was not involved in research ideation, experimental design,
data analysis, or substantive writing. All scientific content, including methodology, results, and
conclusions, was conceived, implemented, and written entirely by the authors.

655
656
657
658
659
660
661

Experiment Setup Our 3D reconstruction and 2D inpainting method is implemented on a single
RTX 4090. We use the default parameters of 3DGS for reconstruction, generating a reconstructed
render every 10,000 iterations. Additionally, we employed the commonly used Stable Diffusion v1.5
(Rombach et al., 2022) as the base inpainting model, training it for 3,000 iterations using textual
inversion for scene representation. Our diffusion model inference consists of a 50-step denoising
process, initialized with a noise strength of 1.0 that is progressively reduced by a factor of 0.2 at
each iteration.

662
663
664
665
666
667

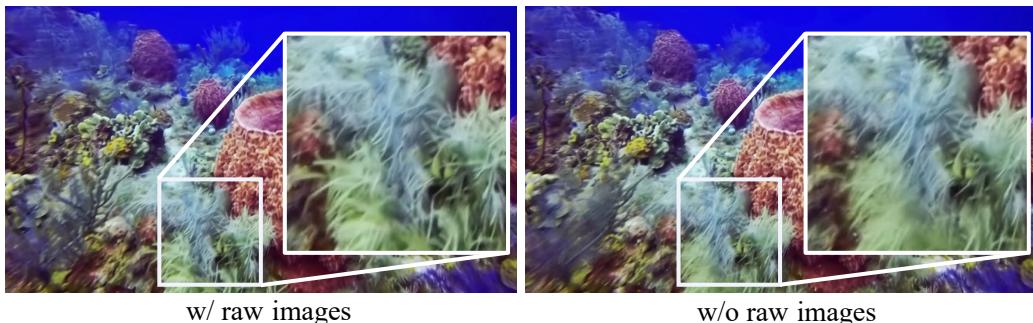
Time cost analysis and comparison. To quantitatively evaluate performance and computational
efficiency, we compare our method against baseline approaches (InFusion, SPIn-NeRF, and Spot-
Less) on the synthetic scene shown in Figure 9. This scene provides ground truth data, enabling
evaluation through reference-based metrics for both rendering quality and computational efficiency
during optimization.

668
669
670
671
672
673

Method	LPIPS ↓	PSNR ↑	Time Cost
InFusion	0.23	19.34	16m 34s
SPIn-NeRF	0.15	23.33	7h 32m 18s
SpotLess	0.14	24.75	30m 26s
Ours	0.10	26.38	33m 34s

674
675
Table 5: Quantitative results and time costs on the synthesis data.676
677
678
679
680
As shown in Table 5, while our method incurs additional computational overhead compared to
vanilla 3DGS due to the integration of iterations and diffusion models, it achieves superior rendering
quality while maintaining comparable efficiency to state-of-the-art 3DGS methods (e.g., SpotLess
(Sabour et al., 2024)). Furthermore, our approach demonstrates significantly better reconstruction
quality while being approximately 10x faster than leading NeRF-based methods such as SPIn-NeRF.681
682
683
684
685
686
687

**Reasons for combining raw images \mathcal{I} and $\tilde{\mathcal{I}}$ rather than substituting raw images \mathcal{I} with $\tilde{\mathcal{I}}$ in
Section 4.4** As shown in Figure 6, the reconstruction without raw images could not render the
seaweed without ambiguity. The accumulated error from two iterations, caused by 3DGS’s inability
to fit the scene fully and the uncertainty introduced by the generated model, deteriorates the image
quality. Raw images act as an “anchor” for our method, ensuring that the rendered images align
closely with the input images and do not deviate significantly.

700
701
Figure 6: Reconstruction results with and without raw images. Involving the raw images in our
method will improve the inpainting performance.

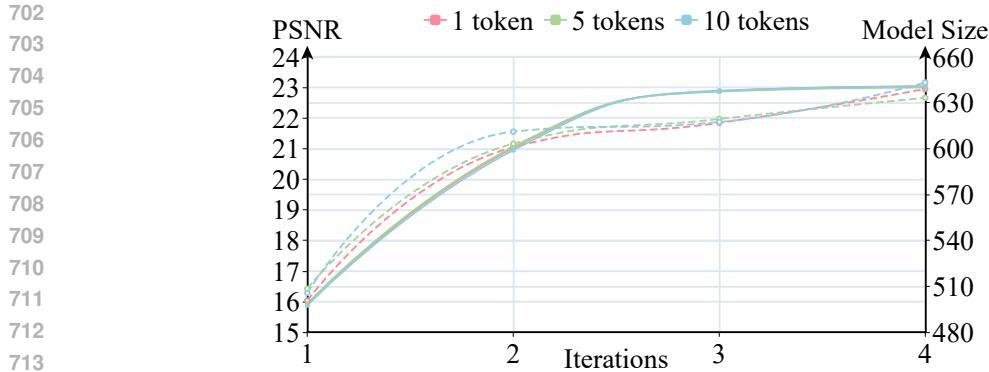


Figure 7: Relationship between model performance (PSNR) and model size (MB) with different token numbers. The dashed and solid lines represent the model size and performance variations, respectively. The model performance (solid lines) under different token numbers almost overlaps.

Effects of token numbers to depict one scene. As shown in Figure 7, our method balances scene depiction efficiency by optimizing token quantity (linked to color specificity) and textual inversion for coarse semantic learning, as excessive tokens inflate model size without performance gains. Training stabilizes at three iterations (determined by PSNR smoothing), beyond which redundant Gaussians overfit diffusion noise, further increasing model parameters.

A.1 IMPACT OF PRIOR (MASK / VIEWS) ON INPAINTING

Impact of the mask quality on Inpainting . Adding the prior (mask) information in our method will significantly improve the inpainting results, especially for those static objects. This is easy to understand because dynamic objects create inconsistencies during the reconstruction process, which our algorithm can detect. In contrast, static object inpainting necessitates the semantic information that the detection model identifies.

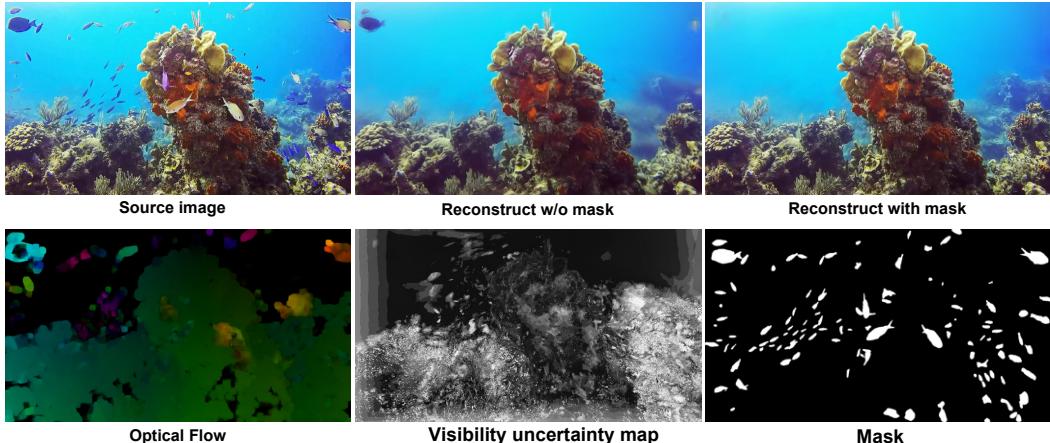


Figure 8: Impact of prior (mask) on inpainting results. Our method will improve inpainting performance by incorporating the mask information. To analyze the results, we also display the optical flow of the source image and the visibility uncertainty map.

For instance, in the top-left corner of Figure 8, the fish is retained while the others are removed. This is primarily because the fish remains stationary across different views (as evident in the optical flow map of Figure 8, where the top-left fish exhibits low flow values at its center). Consequently, it has a lower value in the visibility uncertainty map (see the corresponding map in Figure 8). Without using a mask to label this area for repair manually, the fish's geometric characteristics resemble those of a stationary object, such as a rock, making it indistinguishable from our uncertainty detection system.

Mask Quality	LPIPS ↓	PSNR ↑	SSIM ↑
100 %	0.10	26.38	0.86
50 %	0.10	25.73	0.81
0 %	0.12	25.06	0.79

Table 6: Quantitative results of large viewpoint differences.

In contrast, moving fish create significant geometric inconsistencies across viewpoints, enabling our uncertainty detection to flag them as anomalies. This leads to their removal through the inpainting process. To address these challenging scenarios, we introduced mask annotations for fish detection, providing semantic guidance for our inpainting method. As shown in the last column of Figure 8, incorporating the mask ensures the successful removal of the top-left fish.

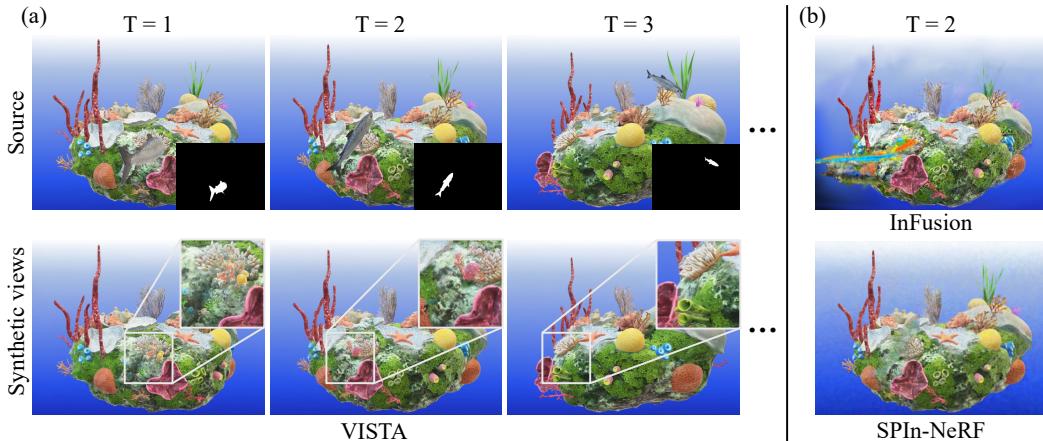


Figure 9: (a) **The figure of an artificial synthesis scene in extreme cases.** The original views of three adjacent cameras and the inpainting results of our method are demonstrated for comparison. (b) **The results of InFusion and SPIn-NeRF in extreme cases.** Their results are obtained by the camera from ‘T = 2’ in (a).

The ablation study of mask quality. Building upon the visualization of the missing mask prior presented earlier, we performed a rigorous quantitative analysis to validate the robustness of our proposed method. In particular, we conducted a controlled ablation study to systematically evaluate its performance under three distinct mask quality scenarios:

- **100% mask:** Complete masks from MASA detection as used in our method.
- **50% mask:** Randomly dropped 50% of the masked pixels from the complete masks.
- **0% mask:** No mask information provided.

We set three different mask qualities. Ablation studies (Table 6) show our method maintains stable performance when progressively removing object masks (100% → 50% → 0% mask). Metrics like PSNR dropped only 0.65 with 50% mask missing, demonstrating inherent robustness to mask absence. As shown in Figure 2 of the paper, this robustness is derived from iterative mask refinement utilizing the iterative refined uncertainty as an updated mask.

Ablation study of large viewpoint differences To evaluate the impact of viewpoint difference, we capture 34 images from continuously distributed viewpoints around a scene to create a ground truth (GT) 3DGS model. We systematically reduce the number of viewpoints by sampling them at different intervals {2, 3, 4, 5, 6, 7}, where larger intervals represent larger viewpoint differences. We reconstruct a 3DGS model per sampling interval and establish a quantitative analysis framework linking viewpoint differences to reconstruction quality by benchmarking rendered images

Sampling Interval	LPIPS ↓	PSNR ↑	SSIM ↑
2	0.09	26.25	0.89
3	0.14	23.42	0.83
4	0.16	22.42	0.80
5	0.27	18.09	0.65
6	0.25	18.71	0.69
7	0.41	15.66	0.57

Table 7: Quantitative results of large viewpoint differences.

against GT models. Our method ensures robust reconstruction under significant viewpoint variations through view-consistency detection and repair mechanisms, while metrics degrade in extreme scenarios with missing key viewpoints due to irreversible information loss.

VISTA in limited scenarios Our uncertainty maps are built by observing a set of adjacent perspectives/views, thus fully utilizing complementary visual cues. However, in some extreme conditions (A sparse 360° capture using only eight input views evenly spaced around the object, yielding an average angular separation of 45°, which severely challenges reconstruction and editing), we don’t have enough valid adjacent perspectives/views to get the visual cues. To investigate the performance of our method under such extreme conditions, we manually synthesized an extreme scenario where the camera rapidly changes poses, resulting in very few available adjacent viewpoints. In this case, the VISTA-GI can hardly detect the inconsistency between different views, requiring VISTA-CL to produce better results.

As shown in Figure 9 (a), thanks to the 2D diffusion model, our method utilizes its results to effectively inpaint the scene in such extreme conditions. Meanwhile, as shown in Figure 9 (b), the InFusion result is unrealistic due to neglecting consistency in inpainting. SPI-NeRF shows a reasonable result, but with blurry and indistinguishable inpainting areas. Compared to other methods, our approach benefits from Scene Conceptual Learning, resulting in clearer and more reasonable repairs in the target areas, and the textures and content maintain consistency with the original scene.

A.2 VISUALIZATION OF OUR METHOD

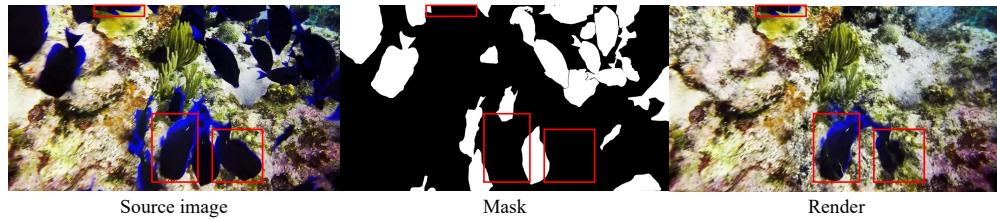
In this part, we visualize more results to demonstrate the effectiveness of our method and the potential failure scenarios that may arise.



Figure 10: Case of real-world pedestrian removal from the nerf-on-the-go dataset.

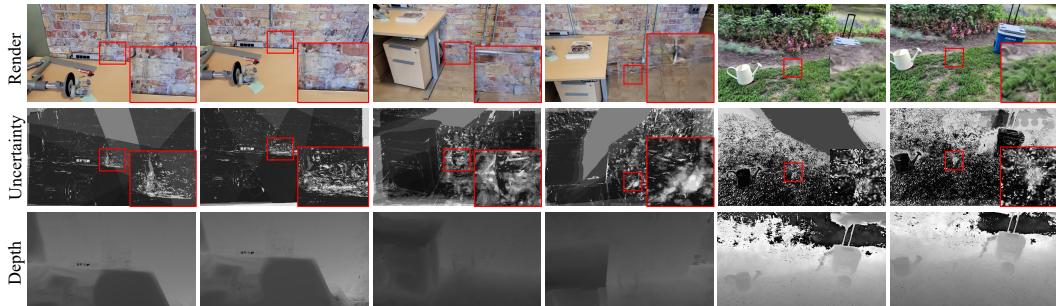
Real-world case. The underwater dataset we used is derived from real-world diving videos, and due to the effects of the underwater medium and floating debris, these scenes are challenging scenarios in the real world. We also tested our dataset on a scene related to pedestrian removal from the NeRF-on-the-Go dataset. This scene, called Tree, contains 212 images, with the main distractors from moving pedestrians. As shown in Figure 10, our method achieved high-quality results on this

864 dataset. Due to the abundance of viewpoints in the dataset, there is a lot of supplementary information
 865 between perspectives, allowing our method to effectively utilize other viewpoints to repair the
 866 blurring caused by moving distractors.
 867



868
 869
 870
 871
 872
 873
 874
 875
 876 Figure 11: Failure case from our underwater 3D inpainting dataset.
 877

878 **Fail case from our dataset.** We provide failure cases of our algorithm in Figure 11. Due to errors
 879 in the prior mask, some fish were not detected by the object detection model. Furthermore, since
 880 the fish did not move significantly during the shooting process, these areas did not produce incon-
 881 sistencies across multiple viewpoints during reconstruction, making it difficult for the VISTA-GI
 882 component to identify these areas through uncertainty. This also validates our algorithm design ap-
 883 proach: VISTA-CL introduces semantic information through masks, while VISTA-GI incorporates
 884 geometric information through uncertainty, complementing each other to remove distractors. How-
 885 ever, in this failed case, issues arose in both aspects, resulting in poor reconstruction quality of the
 886 final scene.
 887



888
 889
 890
 891
 892
 893
 894
 895
 896
 897 Figure 12: Visualization of the uncertainty map and depth of static scenes.
 898

901 **Uncertainty and depth maps of static scenes.** As shown in Figure 12, we further visualize the
 902 uncertainty and depth maps of the static scenes. The deeper the color, the closer the depth. It can
 903 be observed that our method identifies areas in the rendered image that are inconsistent with other
 904 viewpoints and generates reasonable content.
 905

Resolution	LPIPS ↓	PSNR ↑	SSIM ↑
64×64	0.51	16.27	0.68
128×128	0.42	18.89	0.69
256×256	0.26	21.33	0.71
512×512	0.11	26.04	0.84
1299×974	0.10	26.38	0.86

912 Table 8: Quantitative ablation results of different resolutions.
 913

914 A.3 IMPACTS OF DIFFERENT IMAGE RESOLUTION

915 In our experiment setup, we use the stable diffusion v1.5 as the inpainting model and train and test
 916 the model following its default setup: if the input image has a resolution higher than 512×512, we
 917

918 crop the image to a new size that is both the closest to the original image size and a multiple of 8; if
919 the input image is smaller than 512×512 , we rescale the image to 512×512 . To analyze the influence
920 of the strategy on different original resolutions, given an original scene with input images having a
921 size of 1299×974 , we downsample these images to four resolutions: 64×64 , 128×128 , 256×256 ,
922 and 512×512 . Then, for each resolution, we can build a 3D model and evaluate the rendering quality.
923 As shown in Table 8, we observe that: (1) reducing the resolution to 512×512 does not significantly
924 impact any of the metrics, demonstrating our method’s robustness to substantial resolution changes.
925 (2) further decreasing the resolution leads to gradual degradation in reference-based metrics, while
926 non-reference metrics remain relatively stable.

927

928

929

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971