

OceanSplat: Object-aware Gaussian Splatting with Trinocular View Consistency for Underwater Scene Reconstruction

Minseong Kweon¹, Jinsun Park^{2*}

¹University of Minnesota, MN, USA

²Pusan National University, South Korea
kweon021@umn.edu, jspark@pusan.ac.kr

Abstract

We introduce OceanSplat, a novel 3D Gaussian Splatting-based approach for high-fidelity underwater scene reconstruction. To overcome multi-view inconsistencies caused by scattering media, we design a trinocular setup for each camera pose by rendering from horizontally and vertically translated virtual viewpoints, enforcing view consistency to facilitate spatial optimization of 3D Gaussians. Furthermore, we derive synthetic epipolar depth priors from the virtual viewpoints, which serve as self-supervised depth regularizers to compensate for the limited geometric cues in degraded underwater scenes. We also propose a depth-aware alpha adjustment that modulates the opacity of 3D Gaussians during early training based on their depth along the viewing direction, deterring the formation of medium-induced primitives. Our approach promotes the disentanglement of 3D Gaussians from the scattering medium through effective geometric constraints, enabling accurate representation of scene structure and significantly reducing floating artifacts. Experiments on real-world underwater and simulated scenes demonstrate that OceanSplat substantially outperforms existing methods for both scene reconstruction and restoration in scattering media.

Project Page — <https://oceansplat.github.io>

Introduction

Underwater scene reconstruction is essential for several marine robotics tasks, including seafloor mapping (Kapoutsis et al. 2016; Liu et al. 2020; Joshi et al. 2022), ecological monitoring (Marre et al. 2019; Girdhar et al. 2023), and subsea infrastructure inspection (Fang et al. 2023; Huang et al. 2025a). While visual data can generally support robotic operations effectively, the underwater optical properties, such as wavelength-dependent attenuation (Akkaynak et al. 2017), scattering (McGlamery 1980; Jaffe 2002; Chen et al. 2021b), and low illumination (Marques and Albu 2020), significantly degrade perceptual cues (Yu et al. 2023; Zhang, Yuan, and Cai 2024) and hinder the deployment of vision-based autonomous systems (De and Pedersen 2021). Consequently, achieving geometric understanding and spatial representation from this low-quality imagery is essential

*Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

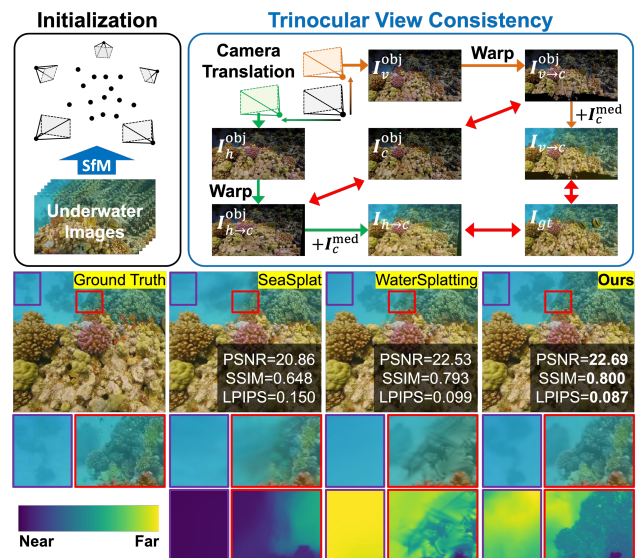


Figure 1: OceanSplat overcomes scattering and attenuation effects through trinocular view consistency, preserving object structure and enabling high-quality underwater 3D reconstruction.

for practical operation of autonomous platforms. Although advances in remotely operated vehicles (ROVs) and autonomous underwater vehicles (AUVs) have greatly enabled the acquisition of large-scale underwater imagery (Wynn et al. 2014), leveraging this data for scene representation remains highly challenging due to adverse visual conditions (Beall et al. 2010; Huang et al. 2025a). To address these issues in underwater perception, the use of Neural Radiance Fields (NeRF) (Mildenhall et al. 2020) and 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023), which have shown strong capability in 3D reconstruction, has emerged as a promising direction. However, these methods rely on assumptions tailored to in-air image acquisition, motivating recent research to integrate underwater image formation model (Akkaynak and Treibitz 2019).

NeRF-based methods (Sethuraman, Ramanagopal, and Skinner 2023; Zhang and Johnson-Roberson 2023; Levy et al. 2023; Ramazzina et al. 2023; Tang et al. 2024) extend their volumetric rendering framework by integrating light-

transmitting physics and scattering phenomena. However, these methods implicitly represent both object and backscatter components, which impedes accurate geometric representation and also results in slow rendering speed. More recently, 3DGS-based methods (Yang, Leonard, and Girdhar 2024; Li et al. 2024a; Wang et al. 2024) have been adapted for underwater environments, benefiting from their superior rendering performance and fast training. Despite notable progress in these studies, medium intensity is often absorbed into the 3D Gaussians, leading to numerous floating artifacts and degrading reconstruction quality.

Unlike previous approaches, our goal is to robustly recover scene geometry in underwater environments subject to scattering and attenuation, without depending on external geometric supervision. To this end, we propose **OceanSplat: Object-aware Gaussian Splatting** for geometrically consistent underwater scene reconstruction. Our method enforces trinocular view consistency across translated camera views along two orthogonal axes, resolving multi-view inconsistencies and effectively guiding 3D Gaussians to align with the scene structure. Additionally, we derive a synthetic epipolar depth prior through triangulation between the translated viewpoints and leverage it for self-supervised depth regularization. Furthermore, we apply a depth residual loss that constrains the z -component of each 3D Gaussian toward the alpha-blended depth to reduce floating artifacts. We also propose a depth-aware alpha adjustment that regulates the transparency of 3D Gaussians during the early stages of training based on their z -component and viewing direction, thereby discouraging the formation of medium-induced primitives. As shown in Figure 1, our method improves geometric accuracy for object representation under challenging scattering conditions by preventing the entanglement of 3D Gaussians with the scattering medium. Our key contributions are outlined in the following points:

- We propose OceanSplat, a novel underwater scene reconstruction method leveraging 3D Gaussian Splatting, which employs trinocular view consistency to align 3D Gaussians with scene geometry in scattering media.
- We introduce synthetic epipolar depth priors derived from translated virtual viewpoints via triangulation, enabling self-supervised depth regularization.
- We propose a depth-aware alpha adjustment that regulates the opacity of 3D Gaussians during the early training stage, suppressing medium-induced primitives.
- Our approach surpasses existing state-of-the-art methods in both qualitative and quantitative evaluations across real-world and simulated datasets.

Related Work

Underwater Scene Reconstruction Underwater scene reconstruction aims to represent precise scene structure and restore true color despite severe light attenuation and scattering. Initial approaches jointly estimated medium-dependent attenuation during reconstruction (Skinner, Iscar, and Johnson-Roberson 2017) and exploited global context with active labeling for sonar-based modeling (DeBortoli

et al. 2018). Subsequent works (Sethuraman, Ramanagopal, and Skinner 2023; Zhang and Johnson-Roberson 2023; Tang et al. 2024) embed physics-based underwater light transport and scattering into NeRF (Mildenhall et al. 2020) to jointly recover geometry and true color appearance. SeaThru-NeRF (Levy et al. 2023) leverages the *SeaThru* image formation model (Akkaynak and Treibitz 2019) within NeRF to successfully achieve novel view synthesis in real-world underwater scenes. ScatterNeRF (Ramazzina et al. 2023) uses inverse rendering, while DecoNeRF (Zhang et al. 2025) employs pseudo-labeling to disentangle scattering media from scene content and restore clear views in foggy scenes. Although these methods significantly improve reconstruction quality, they are often hindered by slow rendering speed and high memory consumption. Recently, 3DGS (Kerbl et al. 2023) has gained increasing attention for underwater scene reconstruction. SeaSplat (Yang, Leonard, and Girdhar 2024) incorporates underwater physics into 3D Gaussians to improve color restoration and structural fidelity, while WaterSplatting (Li et al. 2024a) combines implicit medium modeling with explicit object representation for efficient, high-quality reconstruction. UW-GS (Wang et al. 2024) further enhances geometry using pseudo-depth priors from foundation models (Yang et al. 2024). Despite these advances, existing 3DGS-based methods have difficulty recovering accurate scene structure in underwater environments.

In contrast, our method enforces a trinocular view consistency constraint that enhances the spatial coherence of 3D Gaussians, and leverages self-supervised depth regularization based on epipolar geometry to achieve geometrically accurate underwater scene reconstruction.

Multi-View Stereo Multi-view stereo (MVS) algorithms reconstruct detailed 3D geometry by leveraging images acquired from multiple calibrated viewpoints. Early approaches employed voxelization (Kutulakos and Seitz 2000), dense point clouds (Furukawa and Ponce 2009), and depth maps (Galliani, Lasinger, and Schindler 2015; Schonberger and Frahm 2016). In the deep learning era, MVS-Net (Yao et al. 2018) introduced cost volumes to form geometry-aware 3D representations, greatly enhancing accuracy. Building upon these advancements, MVS has been integrated into neural rendering (Mildenhall et al. 2020; Kerbl et al. 2023), where NeRF-based methods (Chen et al. 2021a; Wang et al. 2021; Fu et al. 2022; Wu et al. 2025) utilize these concepts for surface reconstruction, while 3DGS-based approaches (Chen, Li, and Lee 2023; Guédon and Lepetit 2024; Huang et al. 2024; Chung, Oh, and Lee 2024; Chen et al. 2024a,b; Liu et al. 2024; Li et al. 2024b; Huang et al. 2025b) harness multi-view relationships to enhance geometric fidelity. Among these, recent efforts (Safadoust et al. 2024; Han et al. 2024) have demonstrated promising results by constructing virtual stereo pairs to enforce binocular stereo consistency, effectively regularizing 3D Gaussians.

To better address geometric ambiguities inherent to scattering media, we extend binocular stereo vision to a trinocular configuration by introducing a vertically translated camera, facilitating spatial optimization of 3D Gaussians.

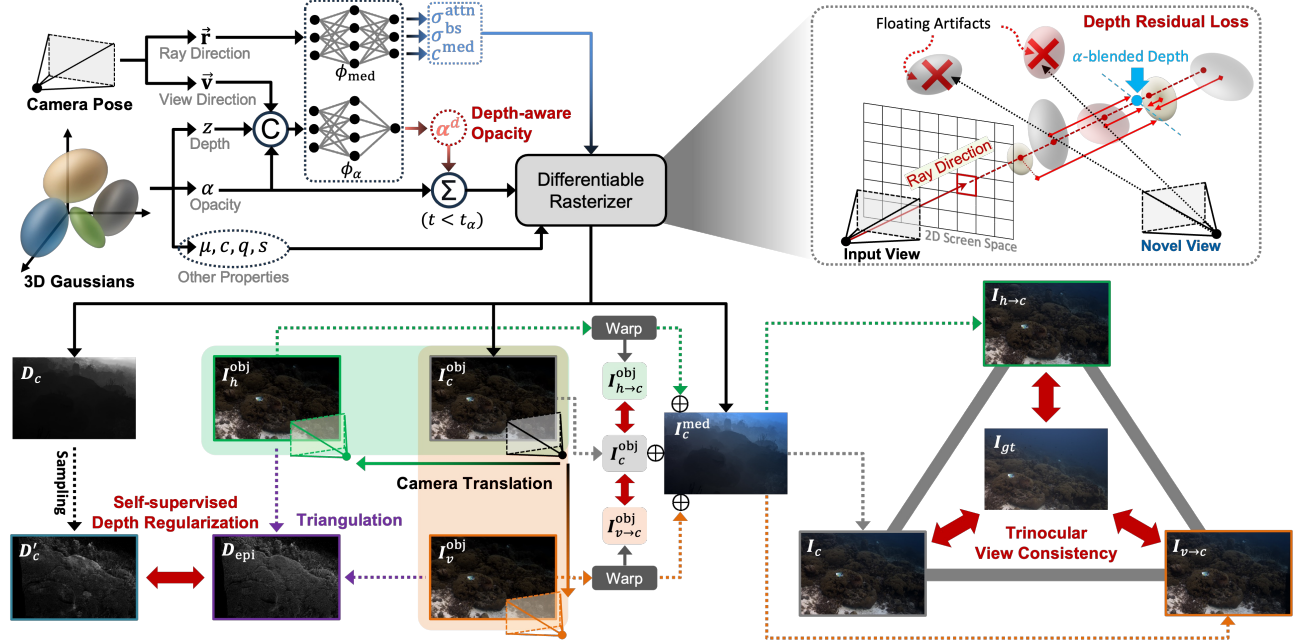


Figure 2: Overview of OceanSplat. We enforce trinocular view consistency by inverse warping rendered images from two translated camera poses to guide the spatial optimization of 3D Gaussians. From these poses, we derive a synthetic epipolar depth prior via triangulation, which provides self-supervised geometric constraints. Additionally, depth-aware alpha adjustment suppresses erroneous 3D Gaussians early and aligns rendered depth with the Gaussian z -component to prevent floating artifacts.

Preliminaries

3D Gaussian Splatting

3D Gaussian Splatting (Kerbl et al. 2023) models a scene with anisotropic 3D Gaussian primitives, each parameterized by a center position $\mu \in \mathbb{R}^3$, covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$, color c , opacity α , and view-dependent appearance modeled by spherical harmonics. Each 3D Gaussian defines its spatial influence at a point $\mathbf{X} \in \mathbb{R}^3$ centered at μ as:

$$G(\mathbf{X}) = e^{-\frac{1}{2}(\mathbf{X}-\mu)^\top \Sigma^{-1}(\mathbf{X}-\mu)}. \quad (1)$$

For rendering, 3D Gaussians are projected onto 2D screen-space and composited via alpha-blending of N primitives that are sorted by depth and overlap at each pixel:

$$C = \sum_{i=1}^N T_i \cdot \alpha_i \cdot c_i, \quad T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (2)$$

where C , N , and T_i denote the rendered color, the number of 3D Gaussians intersected by the ray, and the transmittance, respectively. Then, 3D Gaussians are optimized via an L1 loss combined with D-SSIM on the rendered images.

Underwater Image Formation Model

To render a 3D underwater scene, the image formation is guided by the revised underwater image formation model (Akkaynak and Treibitz 2018), which decomposes the observed image into a direct component attenuated by the medium with respect to depth z , and a backscatter component accumulated along the viewing direction, as follows:

$$C = C^{obj} \cdot e^{-\sigma^{attn} \cdot z} + C^\infty \cdot (1 - e^{-\sigma^{bs} \cdot z}), \quad (3)$$

where C^{obj} , C^∞ , σ^{attn} , and σ^{bs} denote the unattenuated object color, the asymptotic backscatter color, the attenuation coefficient, and the backscattering coefficient, respectively.

Methodology

Figure 2 shows the overall architecture of our method. We begin by initializing 3D Gaussians from a collection of RGB images $I \in \mathbb{R}^{H \times W \times 3}$, along with their corresponding camera intrinsics $K \in \mathbb{R}^{3 \times 3}$ and extrinsic parameters $P \in \mathbb{R}^{4 \times 4}$, which represent the homogeneous transformation matrix, recovered using Structure-from-Motion (SfM) (Schonberger and Frahm 2016). We model medium properties along the ray \vec{r} using an MLP ϕ_{med} , following (Li et al. 2024a), which directly predicts attenuation σ^{attn} , backscattering σ^{bs} , and medium color c^{med} as follows:

$$\sigma^{attn}, \sigma^{bs}, c^{med} = \phi_{med}(\vec{r}). \quad (4)$$

Subsequently, by integrating Eq. (2) with Eq. (3), the contributed colors of the objects and medium along the ray in underwater scenes are formulated as follows:

$$C^{obj} = \sum_{i=1}^N T_i^{obj} \cdot \alpha_i \cdot c_i \cdot e^{-\sigma^{attn} z_i}, \quad T_i^{obj} = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (5)$$

$$C^\infty = T_N^{obj} \cdot c_{med} \cdot e^{-\sigma^{bs} z_N}, \quad (6)$$

$$C^{med} = \sum_{i=1}^N [T_i^{obj} \cdot c_{med} \cdot (e^{-\sigma^{bs} z_{i-1}} - e^{-\sigma^{bs} z_i})] + C^\infty, \quad (7)$$

where C^{obj} , C^{med} , C^∞ , and T^{obj} denote the pixel-wise colors of the object, medium, background medium, and represent the transmittance of 3D Gaussians, respectively. Here,

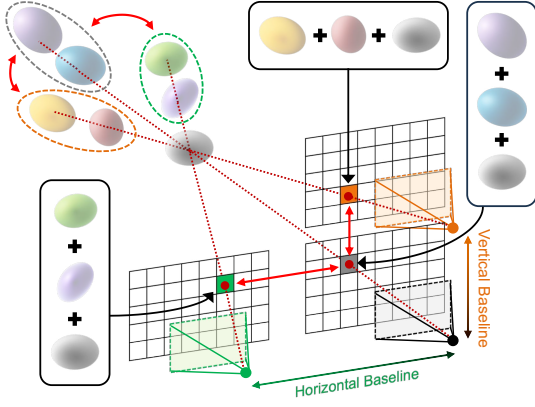


Figure 3: Pixel-wise color contributions across translated camera views in 3D Gaussian Splatting.

the object refers to physical structures submerged in water, while the medium represents the water itself. Based on this rendering framework, we describe our proposed methods for robust 3DGS-based underwater scene reconstruction.

Trinocular View Consistency

We propose a trinocular stereo framework for 3DGS to improve object representation in scattering media. While prior works (Safadoust et al. 2024; Han et al. 2024) use only binocular stereo consistency, we show that a vertically translated viewpoint provides orthogonal constraints that complement horizontal stereo for 3D Gaussians. As shown in Figure 3, view-dependent sampling in alpha-blending often causes inconsistencies under minor viewpoint changes. Furthermore, previous stereo matching studies (Okutomi and Kanade 1993; Imran et al. 2020; Shamsafar and Zell 2021) have shown that multi-baseline stereo surpasses single-baseline methods in accuracy and robustness. Inspired by these observations, we enforce trinocular view consistency with orthogonal baselines of varying lengths to better regularize the spatial position of 3D Gaussians.

From a given camera pose P_c , we obtain the rasterized object and medium components, denoted by I_c^{obj} and I_c^{med} , respectively. We define two virtual viewpoints by translating P_c along the horizontal h and vertical v axes as:

$$P_h = \begin{bmatrix} \mathbb{I} & \mathbf{t}_h \\ \mathbf{0}^\top & 1 \end{bmatrix} P_c, \quad P_v = \begin{bmatrix} \mathbb{I} & \mathbf{t}_v \\ \mathbf{0}^\top & 1 \end{bmatrix} P_c, \quad (8)$$

where \mathbb{I} is the 3×3 identity matrix, $\mathbf{t}_h = (b_h, 0, 0)^\top$ and $\mathbf{t}_v = (0, b_v, 0)^\top$ are translation vectors defining the horizontally and vertically translated camera poses P_h and P_v , respectively, with $b_h, b_v \in \mathbb{R}$ denoting the baseline distances. Using these poses, we render two images I_h^{obj} and I_v^{obj} from viewpoints P_h and P_v . We then apply disparity-based inverse warping to the rendered images. Let $D_c \in \mathbb{R}^{H \times W \times 1}$ be the rendered depth map from P_c . Based on stereo geometry, the horizontal and vertical disparity maps are computed as:

$$d_h(x, y) = \frac{f_h \cdot b_h}{D_c(x, y)}, \quad d_v(x, y) = \frac{f_v \cdot b_v}{D_c(x, y)}, \quad (9)$$

where $f_h, f_v \in \mathbb{R}$ denote the focal lengths in the horizontal and vertical directions. The disparity maps $d_h, d_v \in \mathbb{R}^{H \times W \times 1}$ are used to align the images rendered from the translated virtual viewpoints to the original view, and (x, y) denotes the spatial pixel coordinates. Then, the rendered images from the virtual viewpoints, I_h^{obj} and I_v^{obj} , are inverse warped using disparity maps d_h and d_v , respectively, in the direction opposite to the virtual camera shift, aligning them with the central image rendered from P_c . The resulting inverse warped images $I_{h \rightarrow c}^{\text{obj}}, I_{v \rightarrow c}^{\text{obj}} \in \mathbb{R}^{H \times W \times 3}$ are defined as:

$$\begin{cases} I_{h \rightarrow c}^{\text{obj}}(x, y) = I_h^{\text{obj}}(x - d_h(x, y), y), \\ I_{v \rightarrow c}^{\text{obj}}(x, y) = I_v^{\text{obj}}(x, y - d_v(x, y)), \end{cases} \quad (10)$$

where $d_h, d_v \in \mathbb{R}^{H \times W}$ denote the horizontal and vertical disparities. To reconstruct the complete scene, the shifted object images are composited with the medium component I_c^{med} as:

$$\begin{cases} I_c = I_c^{\text{obj}} + I_c^{\text{med}}, \\ I_{h \rightarrow c} = I_{h \rightarrow c}^{\text{obj}} + I_c^{\text{med}}, \\ I_{v \rightarrow c} = I_{v \rightarrow c}^{\text{obj}} + I_c^{\text{med}}. \end{cases} \quad (11)$$

Afterward, to emphasize dark regions as perceived in human vision, we follow prior studies (Mildenhall et al. 2022; Levy et al. 2023; Li et al. 2024a), and employ a regularized L1 loss L_{R-L1} defined as:

$$L_{R-L1}(I_1, I_2) = \frac{1}{HW} \sum_{x,y} \left| \frac{I_1(x, y) - I_2(x, y)}{\text{sg}(I_c(x, y)) + \varepsilon} \right|, \quad (12)$$

where $\text{sg}(\cdot)$ and ε denote the stop-gradient operator and a small constant for numerical stability, respectively. This loss enforces object consistency across rendered images and photometric alignment with the ground truth I_{gt} as follows:

$$L_{\text{obj-stereo}} = L_{R-L1}(I_{h \rightarrow c}^{\text{obj}}, I_c^{\text{obj}}) + L_{R-L1}(I_{v \rightarrow c}^{\text{obj}}, I_c^{\text{obj}}), \quad (13)$$

$$L_{\text{full-stereo}} = L_{R-L1}(I_{h \rightarrow c}, I_{gt}) + L_{R-L1}(I_{v \rightarrow c}, I_{gt}). \quad (14)$$

In addition, we define the smoothness loss for horizontal and vertical disparities as follows:

$$L_{\text{smooth}} = \frac{1}{HW} \sum_{d \in \{d_h, d_v\}} \sum_{x,y} \sum_{k \in \{h,v\}} |\nabla_k d| \cdot e^{-\gamma |\nabla_k I_{gt}|}, \quad (15)$$

where ∇_k denotes the directional gradient along k -axis and γ is a weighting factor, respectively. Finally, the overall trinocular view consistency loss is defined as:

$$L_{\text{tri}} = L_{\text{obj-stereo}} + L_{\text{full-stereo}} + L_{\text{smooth}}. \quad (16)$$

Synthetic Epipolar Depth Prior

To enhance geometric alignment, we introduce a synthetic epipolar depth prior D_{epi} , derived via triangulation between the translated viewpoints. For reliable correspondences in triangulation, we use 3D Gaussians sampled from the intersection of the view frusta of all trinocular camera views, with opacity above τ_α . Each selected Gaussian is projected onto the image planes of P_h and P_v via projection matrices $\mathbf{M}_h, \mathbf{M}_v \in \mathbb{R}^{3 \times 4}$, producing screen-space points $\mathbf{x}_h^i, \mathbf{x}_v^i \in \mathbb{R}^2$.

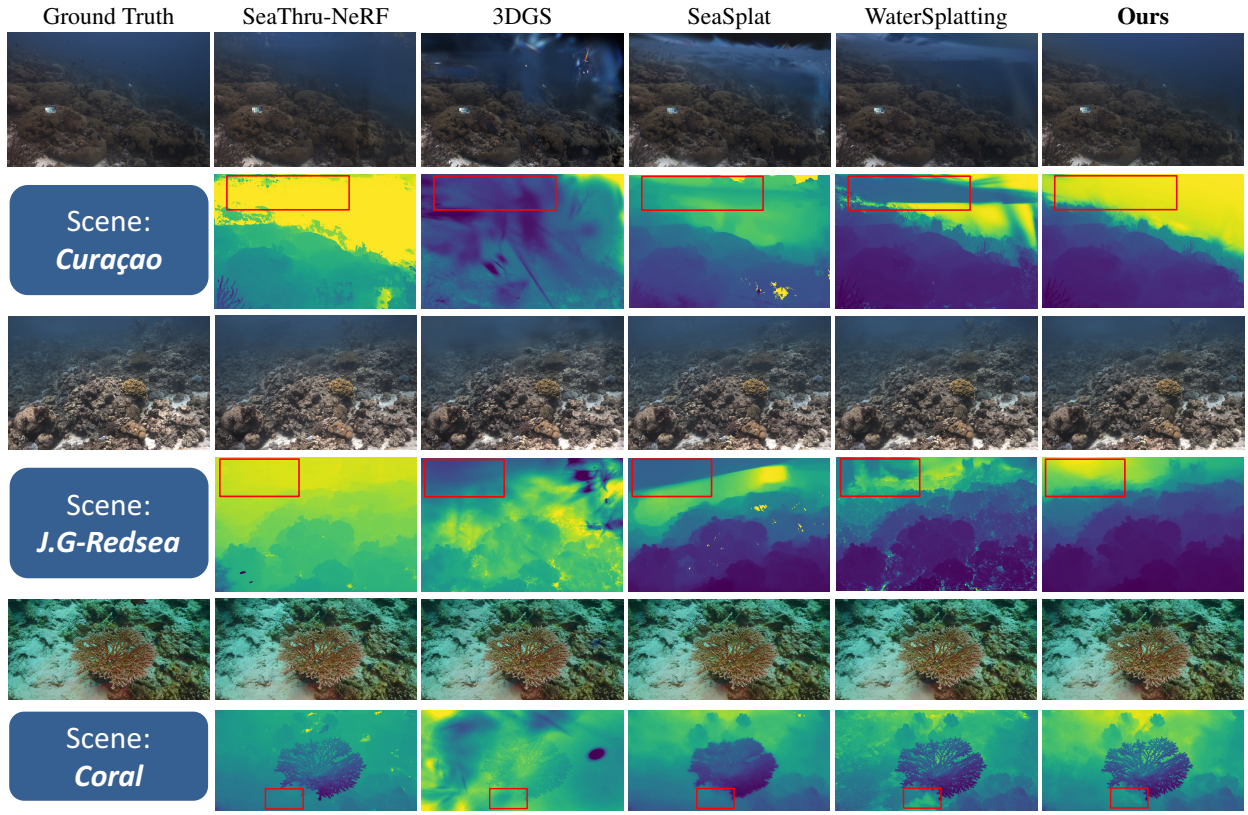


Figure 4: Qualitative evaluation of novel view synthesis on diverse real-world underwater 3D scenes.

Let the homogeneous coordinates be $\tilde{\mathbf{x}}_k^i = [\mathbf{x}_k^{iT}, 1]^\top \in \mathbb{R}^3$ and $\tilde{\mathbf{X}}_i = [\mathbf{X}_i^\top, 1]^\top \in \mathbb{R}^4$. We apply epipolar geometry to the projected correspondences, which yields a homogeneous linear system $\mathbf{A}_i \tilde{\mathbf{X}}_i = \mathbf{0}$, where $\mathbf{A}_i \in \mathbb{R}^{4 \times 4}$ stacks the constraints from both views. To solve it stably, we rewrite it in the least-squares form:

$$\hat{\mathbf{X}}_i = \arg \min_{\mathbf{X}} \|\mathbf{A}'_i \mathbf{X} + \mathbf{b}_i\|_2^2, \quad (17)$$

where $\mathbf{A}'_i \in \mathbb{R}^{4 \times 3}$ and $\mathbf{b}_i \in \mathbb{R}^4$ are obtained by separating the last column of \mathbf{A}_i . The triangulated point $\hat{\mathbf{X}}_i$ is transformed to the central camera frame via extrinsics $\mathbf{R}_c \in \mathbb{R}^{3 \times 3}$, $\mathbf{t}_c \in \mathbb{R}^3$, and the depth prior is defined as the z -component:

$$D_{\text{epi}} = [\mathbf{R}_c \hat{\mathbf{X}}_i + \mathbf{t}_c]_z. \quad (18)$$

After that, we apply an edge-aware Log-L1 loss (Turkulainen et al. 2025) using D_{epi} as follows:

$$L_{\text{epi}} = \frac{1}{HW} \sum_{x,y} \sum_{k \in \{h,v\}} \log(1 + |D'_c - D_{\text{epi}}|) \cdot e^{-|\nabla_k I_c|}, \quad (19)$$

where D'_c denotes the depth sampled at the same spatial locations as D_{epi} . Through this approach, we enforce depth regularization in a self-supervised manner.

Depth Residual Loss

Overly dispersed 3D Gaussians along a camera ray may cause those far from the actual surface to appear as floating artifacts in novel views. To alleviate this issue, we apply

a depth residual loss that aligns the rendered depth with the z -component of individual 3D Gaussians. This residual loss is computed as follows:

$$L_{\text{res}} = \frac{1}{N'} \sum_{i=1}^{N'} |D_c(\mathbf{x}_i) - z_i|, \quad (20)$$

where N' denotes the number of 3D Gaussians located within the camera view frustum, $D_c(\mathbf{x}_i)$ is the rendered depth via alpha-blending at pixel location \mathbf{x}_i , and z_i is the camera-space z -component of the corresponding 3D Gaussian.

Depth-aware Alpha Adjustment

In scattering media, misplaced 3D Gaussians acquire medium-colored contributions that appear as floating artifacts in novel views. To tackle this challenge, we introduce a depth-aware alpha adjustment. During the adjustment stage $t < t_\alpha$, the adjusted opacity α'_i is used for rasterization and is formulated as a weighted sum of the original opacity α_i and the depth-aware opacity α_i^d as follows:

$$\alpha'_i = (1 - w) \alpha_i + w \alpha_i^d, \quad \alpha_i^d = \phi_\alpha(\alpha_i, z_i, \vec{\mathbf{v}}_i), \quad (21)$$

where ϕ_α is the MLP and $\vec{\mathbf{v}}_i$ is the viewing direction vector. Note that, after transition step t_α , the weight w is decayed to zero to eliminate inference-time overhead. This approach adaptively adjusts the transparency of each 3D Gaussian using depth and viewing direction cues, suppresses their opacity in scattering-heavy directions, and promotes the pruning of suboptimal 3D Gaussians.

Method	SeaThru-NeRF												In-the-Wild					
	Curaçao			Panama			JapaneseGardens-Redsea			IU13-Redsea			Coral			Composite		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
SeaThru-NeRF	30.86	0.869	0.216	28.41	0.833	0.219	22.34	0.763	0.263	27.27	0.791	0.293	25.28	0.732	0.315	24.15	0.722	0.315
SeaThru-NeRF-NS	28.58	0.907	0.145	30.60	0.935	0.086	23.11	0.866	0.148	29.00	0.885	0.133	27.60	0.882	0.119	20.14	0.732	0.338
3DGS	28.92	0.881	0.202	29.35	0.899	0.137	21.03	0.859	0.205	22.48	0.832	0.258	29.11	0.906	0.112	26.94	0.877	0.164
SeaSplat	29.77	0.898	0.172	28.65	0.904	0.114	23.07	0.875	0.155	27.23	0.868	0.183	28.41	0.897	0.125	26.22	0.866	0.156
WaterSplatting	32.32	0.954	0.119	31.71	0.946	0.078	24.77	0.897	0.119	29.84	0.890	0.201	28.19	0.877	0.160	25.47	0.849	0.177
Ours	34.56	0.961	0.113	32.74	0.957	0.072	25.35	0.908	0.114	30.17	0.913	0.183	29.15	0.909	0.105	26.39	0.883	0.128

Table 1: Quantitative evaluation of novel view synthesis performance on real-world underwater scenes compared with reproducible prior methods. *1st*, *2nd*, and *3rd* indicate performance ranking.

Method	Underwater			Foggy		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
Novel View Synthesis						
SeaThru-NeRF-NS	18.94	0.669	0.359	23.04	0.724	0.280
SeaSplat	15.62	0.750	0.247	27.52	0.831	0.137
WaterSplatting	28.12	0.858	0.094	28.45	0.860	0.097
Ours	28.80	0.871	0.085	29.12	0.873	0.089
Scene Restoration						
SeaThru-NeRF-NS	12.54	0.544	0.400	14.87	0.605	0.303
SeaSplat	17.15	0.719	0.179	19.84	0.744	0.153
WaterSplatting	18.39	0.748	0.141	19.40	0.770	0.129
Ours	17.16	0.768	0.138	20.17	0.791	0.114

Table 2: Quantitative evaluation of novel view synthesis and scene restoration performance on simulated scenes.

Training Objective

We adopt the photometric loss as follows:

$$L_{\text{photo}} = (1 - \lambda_s)L_{R-L1}(I_c, I_{gt}) + \lambda_s L_{R-SSIM}(I_c, I_{gt}), \quad (22)$$

where L_{R-SSIM} and λ_s denote the Regularized SSIM loss (Li et al. 2024a) and its corresponding weight, respectively, with $\lambda_s = 0.2$. In addition to the baseline, we incorporate our proposed components into the final training objective:

$$L_{\text{total}} = L_{\text{photo}} + \lambda_{\text{tri}}L_{\text{tri}} + \lambda_{\text{epi}}L_{\text{epi}} + \lambda_{\text{res}}L_{\text{res}}, \quad (23)$$

where λ_{tri} , λ_{epi} , and λ_{res} are the weights of the proposed loss terms. We empirically set $\lambda_{\text{tri}} = 0.1$ and $\lambda_{\text{res}} = 0.01$, while λ_{epi} is annealed from 0.4 to 0.2.

Experiment

Datasets and Metrics

We evaluate underwater scene reconstruction on four scenes from SeaThru-NeRF dataset (Levy et al. 2023) and two scenes from In-the-Wild dataset (Tang et al. 2024). In addition, to evaluate scene restoration performance, we generate simulated underwater and foggy scenes using the *Fern* scene from the LLFF dataset (Mildenhall et al. 2019). We simulate underwater scenes following previous methods (Levy et al. 2023) using normalized depth obtained from a depth foundation model (Yang et al. 2024) with $\beta_D = [1.3, 1.2, 0.9]$, $\beta_B = [0.95, 0.85, 0.7]$, and $\beta_\infty = [0.07, 0.2, 0.39]$. For foggy scenes, we use $\beta_D = [0.5, 0.5, 0.5]$ and $\beta_B = \beta_\infty = 1.2$. To validate our method, we compare scene quality using PSNR, SSIM (Wang et al. 2004), and LPIPS (Zhang et al. 2018).

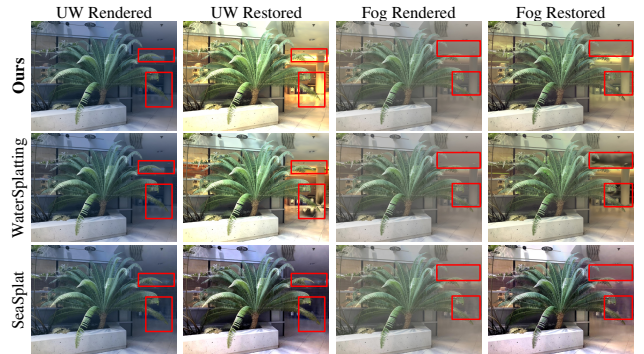


Figure 5: Qualitative comparison of scene reconstruction and restoration in underwater and foggy environments.

Implementation Details

Our method builds on Nerfstudio (Tancik et al. 2023), and employs a TCNN-based MLP (Müller 2021). We use 7K/3K training steps for densification/finetuning on SeaThru-NeRF (Levy et al. 2023) and simulated datasets, and 10K/5K steps on In-the-Wild (Tang et al. 2024). At each iteration, b_v is sampled from $[-0.4, 0.4]$, and b_h is set to $1.5b_v$. We enable L_{tri} and L_{epi} during iterations 4K–8K and 6K–12K for 10K and 15K training steps, respectively. The τ_α for synthetic epipolar depth is set to 0.8. The input resolution is increased from $\times 1/4$ to $\times 1/2$ at $1/5$ of total training steps, and from $\times 1/2$ to full resolution at $2/5$ to boost the training process. For depth-aware alpha adjustment, w is set to 0.5 for $t < t_\alpha$ and is decayed to 0 at t_α . We set the transition step t_α to 4K and 10K for 10K and 15K training steps, respectively. All experiments use a single RTX A6000.

Performance Comparison

Real-world Underwater Scenes We evaluate OceanSplat with previous methods, including the official and Nerfstudio versions of SeaThru-NeRF (Levy et al. 2023), as well as 3DGS (Kerbl et al. 2023), SeaSplat (Yang, Leonard, and Girdhar 2024), and WaterSplatting (Li et al. 2024a). All results were reproduced using the publicly released code. As shown in Table 1, OceanSplat achieves superior performance in novel view synthesis across most real-world underwater scenes. On the SeaThru-NeRF dataset (Levy et al. 2023), OceanSplat outperforms the previous state-of-the-art, WaterSplatting (Li et al. 2024a), and SeaThru-NeRF-

Method	Training ↓	FPS ↑	Memory (GB) ↓
SeaThru-NeRF	18 h 25 m	0.10	37.7
SeaThru-NeRF-NS	3 h 21 m	0.29	15.1
3DGS	22 m	150.28	3.7
SeaSplat	51 m	90.05	6.0
WaterSplatting	10 m	88.19	5.8
Ours	19 m	85.67	7.6

Table 3: Comparison of training time, rendering speed (FPS), and memory consumption with prior methods.

NS (Levy et al. 2023) by average PSNR margins of 1.05 dB and 2.88 dB, respectively. On the In-the-Wild dataset (Tang et al. 2024), OceanSplat reduces the average LPIPS by 0.021 compared to 3DGS and surpasses SeaSplat (Yang, Leonard, and Girdhar 2024) with average gains of 0.46 dB in PSNR and 0.014 in SSIM, along with a 0.024 reduction in LPIPS.

Figure 4 compares RGB and depth renderings across three underwater scenes. In the *Curaçao* scene, prior methods suffer from degraded reconstructions due to floating artifacts, whereas our method models the 3D scene with significantly fewer floating artifacts and accurate object representation. In *J.G-Redsea* scene, prior methods yield plausible RGB images but their depth maps reveal erroneous 3D Gaussians in water regions. In contrast, our method maintains consistent geometry throughout. In the *Coral* scene, most prior methods (Levy et al. 2023; Kerbl et al. 2023; Yang, Leonard, and Girdhar 2024) fail to capture fine-grained coral structures, and WaterSplatting (Li et al. 2024a) shows geometry loss and depth holes, while ours reconstructs a faithful 3D scene.

We additionally compare training time, inference FPS, and VRAM usage with prior methods. As shown in Table 3, our method trains faster and consumes less memory than SeaThru-NeRF (Levy et al. 2023), while remaining competitive among 3DGS-based approaches. Compared to WaterSplatting (Li et al. 2024a), our method requires more rasterizations per iteration and least-squares solving for the synthetic epipolar depth prior, resulting in longer training time per-iteration. Nevertheless, it trains over $\times 2$ faster than SeaSplat (Yang, Leonard, and Girdhar 2024) while retaining comparable inference speed to other 3DGS-based methods.

Simulated Scattering Scenes We compare the scene reconstruction and restoration performance of SeaThru-NeRF-NS (Levy et al. 2023), SeaSplat (Yang, Leonard, and Girdhar 2024), WaterSplatting (Li et al. 2024a), and our proposed OceanSplat using a simulated dataset. As shown in Table 2, for the novel view synthesis task, OceanSplat achieves an average PSNR improvement of 7.39 dB over SeaSplat (Yang, Leonard, and Girdhar 2024) and 0.67 dB over WaterSplatting (Li et al. 2024a). For scene restoration, our method also demonstrates superior performance, yielding average SSIM gains of 0.048 over SeaSplat (Yang, Leonard, and Girdhar 2024) and 0.021 over WaterSplatting (Li et al. 2024a). Figure 5 shows that, in both simulated scenes, previous methods assign 3D Gaussians to the scattering medium (e.g., ceilings or walls), resulting in restoration holes, whereas our method accurately represents scene structures in both underwater and foggy scenes.

Configuration	PSNR ↑	SSIM ↑	LPIPS ↓
Full Model	34.56	0.961	0.113
w/o L_{res}	34.30	0.960	0.115
w/o L_{epi}	33.82	0.959	0.120
w/o L_{tri}	33.20	0.957	0.115
w/o α^d	33.90	0.960	0.116

Table 4: Ablation study demonstrating the effectiveness of each component in our method.

Configuration	PSNR ↑	SSIM ↑	LPIPS ↓	Training ↓
Baseline (P_c)	32.20	0.954	0.120	9m
Binocular ($P_c + P_h$)	33.62	0.957	0.119	13m
Trinocular ($P_c + P_h + P_{h'}$)	33.72	0.957	0.118	16m
Trinocular ($P_c + P_h + P_v$)	33.91	0.958	0.115	16m

Table 5: Ablation study on stereo configurations, comparing the impact of horizontally and vertically translated camera poses on performance and training time.

Ablation Study

In Table 4, we evaluate the effectiveness of each component on the *Curaçao* scene. Removing L_{tri} yields the largest PSNR drop, confirming that trinocular view consistency guides spatial optimization of 3D Gaussians, enabling structurally coherent novel view synthesis. Excluding L_{epi} significantly increases LPIPS, validating our self-supervised depth regularization in underwater scenes with a lack of geometric cues. Omitting L_{res} and depth-aware opacity α^d degrades overall quality, highlighting their role in mitigating medium entanglement through reduced depth ambiguity.

Table 5 compares rendering performance and training time across different stereo configurations, including the baseline (Li et al. 2024a). Binocular stereo setup ($P_c + P_h$) improves over baseline, while adding another horizontal stereo view ($P_c + P_h + P_{h'}$) provides minimal gain despite longer training. In contrast, introducing vertical disparity ($P_c + P_h + P_v$) achieves a 0.29 dB increase in PSNR over the binocular setup. This shows that using horizontally and vertically translated camera poses enhances performance without compromising training time, as shown in Figure 3.

Conclusion

We introduce OceanSplat, a novel framework for enhancing geometric consistency in underwater scene reconstruction. Primarily, we impose trinocular view consistency across horizontally and vertically translated virtual viewpoints, effectively aligning 3D Gaussians with scene structure. In addition, we introduce synthetic epipolar depth priors derived via triangulation of correspondences from the translated viewpoints, enabling self-supervised depth regularization. Lastly, we propose a depth-aware alpha adjustment that uses depth and viewing direction to suppress medium-induced floating artifacts. Through these contributions, OceanSplat achieves high-fidelity underwater scene reconstruction. Future work will focus on addressing non-rigid object representations and eliminating the dependence on SfM, thereby broadening the applicability to dynamic underwater scenes.

Acknowledgments

This work was supported in part by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (RS-2023-00217689, 50%; RS-2024-00358935, 40%), and in part by Institute of Information & communications Technology Planning & Evaluation (IITP) under the Artificial Intelligence Convergence Innovation Human Resources Development grant funded by the Korea government (MSIT) (IITP-2026-RS-2023-00254177, 10%). Minseong Kweon acknowledges travel support from the Minnesota Robotics Institute (MnRI).

References

- Akkaynak, D.; and Treibitz, T. 2018. A revised underwater image formation model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6723–6732.
- Akkaynak, D.; and Treibitz, T. 2019. Sea-thru: A method for removing water from underwater images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1682–1691.
- Akkaynak, D.; Treibitz, T.; Shlesinger, T.; Loya, Y.; Tamir, R.; and Iluz, D. 2017. What is the space of attenuation coefficients in underwater computer vision? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4931–4940.
- Beall, C.; Lawrence, B. J.; Ila, V.; and Dellaert, F. 2010. 3D reconstruction of underwater structures. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4418–4423. IEEE.
- Chen, A.; Xu, Z.; Zhao, F.; Zhang, X.; Xiang, F.; Yu, J.; and Su, H. 2021a. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14124–14133.
- Chen, D.; Li, H.; Ye, W.; Wang, Y.; Xie, W.; Zhai, S.; Wang, N.; Liu, H.; Bao, H.; and Zhang, G. 2024a. Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*.
- Chen, H.; Li, C.; and Lee, G. H. 2023. Neusg: Neural implicit surface reconstruction with 3d gaussian splatting guidance. *arXiv preprint arXiv:2312.00846*.
- Chen, X.; Zhang, P.; Quan, L.; Yi, C.; and Lu, C. 2021b. Underwater image enhancement based on deep learning and image formation model. *arXiv preprint arXiv:2101.00991*.
- Chen, Y.; Xu, H.; Zheng, C.; Zhuang, B.; Pollefeys, M.; Geiger, A.; Cham, T.-J.; and Cai, J. 2024b. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. In *Proceedings of the European Conference on Computer Vision*, 370–386. Springer.
- Chung, J.; Oh, J.; and Lee, K. M. 2024. Depth-regularized optimization for 3d gaussian splatting in few-shot images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 811–820.
- De, K.; and Pedersen, M. 2021. Impact of colour on robustness of deep neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 21–30.
- DeBortoli, R.; Nicolai, A.; Li, F.; and Hollinger, G. A. 2018. Real-time underwater 3D reconstruction using global context and active labeling. In *IEEE International Conference on Robotics and Automation*, 6204–6211. IEEE.
- Fang, X.; Li, H.; Zhang, S.; Zhang, J.; Wang, C.; Wang, X.; Ma, Z.; and Jia, H. 2023. Integration of ROV and vision-based underwater inspection for *Limnoperna fortunei* in water conveyance structure. *Engineering Applications of Artificial Intelligence*, 124: 106575.
- Fu, Q.; Xu, Q.; Ong, Y. S.; and Tao, W. 2022. Geo-neus: Geometry-consistent neural implicit surfaces learning for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 35: 3403–3416.
- Furukawa, Y.; and Ponce, J. 2009. Accurate, dense, and robust multiview stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(8): 1362–1376.
- Galliani, S.; Lasinger, K.; and Schindler, K. 2015. Massively parallel multiview stereopsis by surface normal diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 873–881.
- Girdhar, Y.; McGuire, N.; Cai, L.; Jamieson, S.; McCammon, S.; Claus, B.; Soucie, J. E. S.; Todd, J. E.; and Mooney, T. A. 2023. CUREE: A Curious Underwater Robot for Ecosystem Exploration. In *IEEE International Conference on Robotics and Automation*, 11411–11417. IEEE.
- Guédon, A.; and Lepetit, V. 2024. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5354–5363.
- Han, L.; Zhou, J.; Liu, Y.-S.; and Han, Z. 2024. Binocular-guided 3d gaussian splatting with view consistency for sparse view synthesis. *arXiv preprint arXiv:2410.18822*.
- Huang, B.; Yu, Z.; Chen, A.; Geiger, A.; and Gao, S. 2024. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, 1–11.
- Huang, G.; Wang, H.; Seymour, B.; Kovacs, E.; Ellerbrock, J.; Blackham, D.; and Anantrasirichai, N. 2025a. Visual enhancement and 3D representation for underwater scenes: a review. *arXiv preprint arXiv:2505.01869*.
- Huang, H.; Wu, Y.; Deng, C.; Gao, G.; Gu, M.; and Liu, Y.-S. 2025b. FatesGS: Fast and Accurate Sparse-View Surface Reconstruction Using Gaussian Splatting with Depth-Feature Consistency. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Imran, S.; Khan, M. U. K.; Mukaram, S.; and Kyung, C.-M. 2020. Unsupervised Monocular Depth Estimation with Multi-Baseline Stereo. In *Proceedings of the British Machine Vision Conference*.
- Jaffe, J. S. 2002. Computer modeling and the design of optimal underwater imaging systems. *IEEE Journal of Oceanic Engineering*, 15(2): 101–111.
- Joshi, B.; Xanthidis, M.; Roznere, M.; Burgdorfer, N. J.; Mordohai, P.; Li, A. Q.; and Rekleitis, I. 2022. Underwater exploration and mapping. In *2022 IEEE/OES Autonomous Underwater Vehicles Symposium*, 1–7. IEEE.
- Kapoutsis, A. C.; Chatzichristofis, S. A.; Doitsidis, L.; De Sousa, J. B.; Pinto, J.; Braga, J.; and Kosmatopoulos, E. B. 2016. Real-time adaptive multi-robot exploration with application to underwater map construction. *Autonomous Robots*, 40(6): 987–1015.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4): 139–1.
- Kutulakos, K. N.; and Seitz, S. M. 2000. A theory of shape by space carving. *International Journal of Computer Vision*, 38(3): 199–218.
- Levy, D.; Peleg, A.; Pearl, N.; Rosenbaum, D.; Akkaynak, D.; Korman, S.; and Treibitz, T. 2023. Seathru-nerf: Neural radiance fields in scattering media. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 56–65.
- Li, H.; Song, W.; Xu, T.; Elsig, A.; and Kulhanek, J. 2024a. Water-splatting: Fast underwater 3d scene reconstruction using gaussian splatting. *arXiv preprint arXiv:2408.08206*.

- Li, J.; Zhang, J.; Bai, X.; Zheng, J.; Ning, X.; Zhou, J.; and Gu, L. 2024b. Dngaussian: Optimizing sparse-view 3d gaussian radiance fields with global-local depth normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20775–20785.
- Liu, B.; Liu, Z.; Men, S.; Li, Y.; Ding, Z.; He, J.; and Zhao, Z. 2020. Underwater hyperspectral imaging technology and its applications for detecting and mapping the seafloor: A review. *Sensors*, 20(17): 4962.
- Liu, T.; Wang, G.; Hu, S.; Shen, L.; Ye, X.; Zang, Y.; Cao, Z.; Li, W.; and Liu, Z. 2024. Mvsgaussian: Fast generalizable gaussian splatting reconstruction from multi-view stereo. In *Proceedings of the European Conference on Computer Vision*, 37–53. Springer.
- Marques, T. P.; and Albu, A. B. 2020. L2uwe: A framework for the efficient enhancement of low-light underwater images using local contrast and multi-scale fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition workshops*, 538–539.
- Marre, G.; Holon, F.; Luque, S.; Boissery, P.; and Deter, J. 2019. Monitoring marine habitats with photogrammetry: a cost-effective, accurate, precise and high-resolution reconstruction method. *Frontiers in Marine Science*, 6: 276.
- McGlamery, B. 1980. A computer model for underwater camera systems. In *Ocean Optics VI*, volume 208, 221–231. SPIE.
- Mildenhall, B.; Hedman, P.; Martin-Brualla, R.; Srinivasan, P. P.; and Barron, J. T. 2022. Nerf in the dark: High dynamic range view synthesis from noisy raw images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16190–16199.
- Mildenhall, B.; Srinivasan, P. P.; Ortiz-Cayon, R.; Kalantari, N. K.; Ramamoorthi, R.; Ng, R.; and Kar, A. 2019. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics*, 38(4): 1–14.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *Proceedings of the European Conference on Computer Vision*.
- Müller, T. 2021. tiny-cuda-nn.
- Okutomi, M.; and Kanade, T. 1993. A multiple-baseline stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(4): 353–363.
- Ramazzina, A.; Bijelic, M.; Walz, S.; Sanvito, A.; Scheuble, D.; and Heide, F. 2023. Scatternerf: Seeing through fog with physically-based inverse neural rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17957–17968.
- Safadoust, S.; Tosi, F.; Güney, F.; and Poggi, M. 2024. Self-Evolving Depth-Supervised 3D Gaussian Splatting from Rendered Stereo Pairs. *arXiv preprint arXiv:2409.07456*.
- Schonberger, J. L.; and Frahm, J.-M. 2016. Structure-from-motion revisited. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4104–4113.
- Sethuraman, A. V.; Ramanagopal, M. S.; and Skinner, K. A. 2023. Waternerf: Neural radiance fields for underwater scenes. In *OCEANS 2023-MTS/IEEE US Gulf Coast*, 1–7. IEEE.
- Shamsafar, F.; and Zell, A. 2021. TriStereoNet: A Trinocular Framework for Multi-baseline Disparity Estimation. *arXiv preprint arXiv:2111.12502*.
- Skinner, K. A.; Iscar, E.; and Johnson-Roberson, M. 2017. Automatic color correction for 3D reconstruction of underwater scenes. In *IEEE International Conference on Robotics and Automation*, 5140–5147. IEEE.
- Tancik, M.; Weber, E.; Ng, E.; Li, R.; Yi, B.; Wang, T.; Kristoffersen, A.; Austin, J.; Salahi, K.; Ahuja, A.; et al. 2023. Nerfstudio: A modular framework for neural radiance field development. In *ACM SIGGRAPH 2023 Conference Papers*, 1–12.
- Tang, Y.; Zhu, C.; Wan, R.; Xu, C.; and Shi, B. 2024. Neural underwater scene representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11780–11789.
- Turkulainen, M.; Ren, X.; Melekhov, I.; Seiskari, O.; Rahtu, E.; and Kannala, J. 2025. DN-Splatter: Depth and Normal Priors for Gaussian Splatting and Meshing. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*.
- Wang, H.; Anantrasirichai, N.; Zhang, F.; and Bull, D. 2024. UW-GS: Distractor-aware 3d gaussian splatting for enhanced underwater scene reconstruction. *arXiv preprint arXiv:2410.01517*.
- Wang, P.; Liu, L.; Liu, Y.; Theobalt, C.; Komura, T.; and Wang, W. 2021. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612.
- Wu, Y.; Huang, H.; Zhang, W.; Deng, C.; Gao, G.; Gu, M.; and Liu, Y.-S. 2025. Sparis: Neural Implicit Surface Reconstruction of Indoor Scenes from Sparse Views. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Wynn, R. B.; Huvenne, V. A.; Le Bas, T. P.; Murton, B. J.; Connolly, D. P.; Bett, B. J.; Ruhl, H. A.; Morris, K. J.; Peakall, J.; Parsons, D. R.; et al. 2014. Autonomous Underwater Vehicles (AUVs): Their past, present and future contributions to the advancement of marine geoscience. *Marine Geology*, 352: 451–468.
- Yang, D.; Leonard, J. J.; and Girdhar, Y. 2024. Seasplat: Representing underwater scenes with 3d gaussian splatting and a physically grounded image formation model. *arXiv preprint arXiv:2409.17345*.
- Yang, L.; Kang, B.; Huang, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10371–10381.
- Yao, Y.; Luo, Z.; Li, S.; Fang, T.; and Quan, L. 2018. Mvsnets: Depth inference for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision*, 767–783.
- Yu, M.; Shen, L.; Wang, Z.; and Hua, X. 2023. Task-friendly underwater image enhancement for machine vision applications. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–14.
- Zhang, M.; Zhang, J.; Fang, F.; and Zhang, G. 2025. Decoupling Scattering: Pseudo-Label Guided NeRF for Scenes with Scattering Media. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 586–595.
- Zhang, T.; and Johnson-Roberson, M. 2023. Beyond nerf underwater: Learning neural reflectance fields for true color correction of marine imagery. *IEEE Robotics and Automation Letters*, 8(10): 6467–6474.
- Zhang, Y.; Yuan, J.; and Cai, Z. 2024. DCGF: Diffusion-Color Guided Framework for Underwater Image Enhancement. *IEEE Transactions on Geoscience and Remote Sensing*.