

WaterSplatting: Fast Underwater 3D Scene Reconstruction Using Gaussian Splatting

Huapeng Li
University of Zurich
huapeng.li@uzh.ch

Wenxuan Song
ETH Zurich
wesong@student.ethz.ch

Tianao Xu
ETH Zurich
tianaxu@student.ethz.ch

Alexandre Elsig
ETH Zurich
elsiga@student.ethz.ch

Jonas Kulhanek
CTU in Prague, ETH Zurich
jonas.kulhanek@cvut.cz

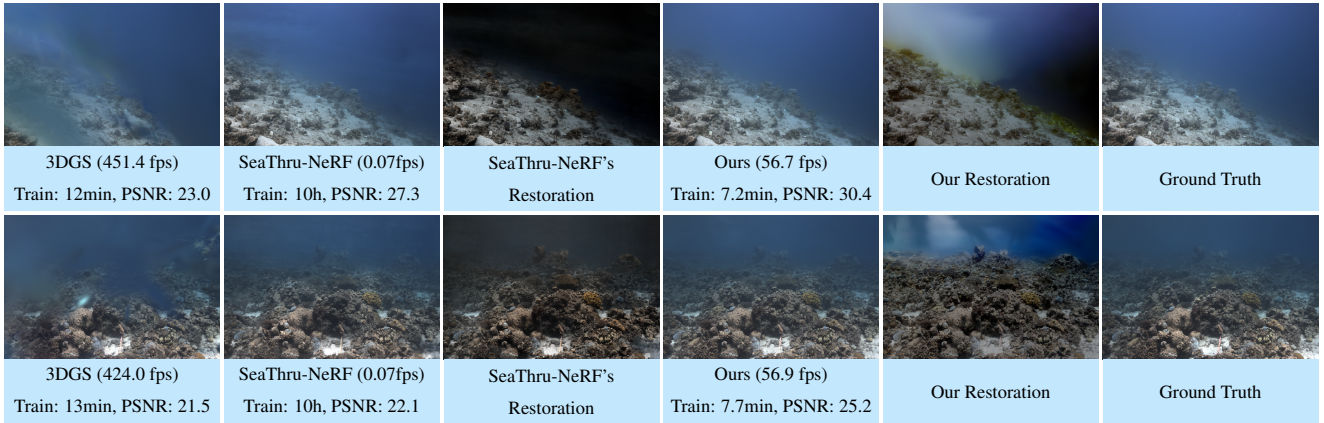


Figure 1. Our approach surpasses the performance of state-of-the-art NeRF-based underwater reconstruction methods [18] while offering real-time rendering speed [15].

Abstract

The underwater 3D scene reconstruction is a challenging, yet interesting problem with applications ranging from naval robots to VR experiences. The problem was successfully tackled by fully volumetric NeRF-based methods which can model both the geometry and the medium (water). Unfortunately, these methods are slow to train and do not offer real-time rendering. More recently, 3D Gaussian Splatting (3DGS) method offered a fast alternative to NeRFs. However, because it is an explicit method that renders only the geometry, it cannot render the medium and is therefore unsuited for underwater reconstruction. Therefore, we propose a novel approach that fuses volumetric rendering with 3DGS to handle underwater data effectively. Our method employs 3DGS for explicit geometry representation and a separate volumetric field (queried once per pixel) for capturing the scattering medium. This dual representation further allows the restoration of the scenes by removing the scattering medium. Our method outperforms state-of-the-art NeRF-based methods in rendering quality on the un-

derwater SeaThru-NeRF dataset. Furthermore, it does so while offering real-time rendering performance, addressing the efficiency limitations of existing methods.

1. Introduction

Neural Radiance Fields (NeRFs) [24] have recently gained significant popularity due to their ability to offer photorealistic 3D scene reconstruction quality. This has opened up new avenues in the field of 3D rendering and reconstruction. However, the landscape of 3D rendering techniques is rapidly evolving. More recently, point splatting methods have experienced a resurgence in the form of 3D Gaussian Splatting (3DGS) [15], which matches NeRFs in terms of rendering quality and offers real-time rendering speed, better editability, and control.

The reconstruction of scattering scenes, such as foggy and underwater environments, is an interesting research area with applications ranging from naval robots to VR experiences. Reconstructing geometry inside a water volume

is challenging due to the presence of the scattering medium with properties different from air. In a typical scene, the primary requirement is to represent the surface. Both NeRFs and Gaussian splatting methods are optimized to focus on representing the surfaces only, thereby gaining better efficiency. In the case of NeRFs, since they are fully volumetric, they should theoretically be able to represent the medium fully volumetrically. However, this is no longer the case as the proposal sampler used to speed up NeRFs prevents them from learning volumes well.

To address this issue, a NeRF approach, SeaThru-NeRF [18], was proposed, which uses two fields: one for the geometry and one for the volume in between. However, it is slow in both rendering and training. Therefore, we propose a novel approach to represent the geometry explicitly using 3DGS but to represent the volume in between using a volumetric representation. The renderer we propose not only surpasses the rendering quality of fully volumetric representations, as demonstrated by [18] but also achieves rendering and training speeds comparable to 3DGS.

To validate our method, we evaluate it on the established benchmark underwater dataset - SeaThru-NeRF [18]. The results of our evaluation demonstrate the effectiveness of our proposed method in achieving high-quality, efficient underwater reconstruction. In summary, we make the following contributions:

1. **Splatting with Medium:** We introduce a novel approach that combines the strengths of Gaussian Splatting (GS) and volume rendering. Our method employs GS for explicit geometry representation and a separate volumetric field for capturing the scattering medium. This dual representation allows for the synthesis of novel views in scattering media and the restoration of clear scenes without a medium.

2. **Loss Function Alignment:** We propose a novel loss function designed to align 3DGS with human perception of High Dynamic Range (HDR) and low-light scenes.

3. **Efficient Synthesis and Restoration:** We demonstrate that our method outperforms other models on synthesizing novel view on real-world underwater data and restoring clean scenes on synthesized back-scattering scenes with much shorter training and rendering time.

2. Related Work

2.1. NeRF

The field of 3D scene reconstruction has gained significant attention with the advent of NeRF methods [23, 24, 35]. NeRFs represent the 3D scene as a radiance field—comprising differential volume density and view-dependent color—rendered using a volume rendering integral from a list of samples sampled along the ray [29]. Originally, NeRFs utilized Multilayer Perceptrons (MLPs)

for representing the radiance field [3, 4, 24], but they were slow to train and render. To accelerate the training and rendering processes, alternative methods have been proposed using discrete grids [11, 46], hash grids [5, 26, 37], tensorial decomposition [9, 31], point clouds [42], or tetrahedral mesh [17]. NeRFs have been enhanced in various ways, including improved anti-aliasing [3, 5], handling of large 3D scenes [36], and complex camera trajectories [4, 40]. Moreover, NeRFs have been extended to a wide range of applications such as semantic segmentation [6, 16], few-view novel view synthesis [7, 8, 20, 41, 47], and generative 3D modeling [22, 27]. Despite these advancements, the slow rendering speed of NeRFs remains a critical limitation, hindering their widespread adoption on end-user devices.

2.2. 3D Gaussian Splatting

Recently, Gaussian Splatting (3DGS) [15] has seen a resurgence as a powerful method for real-time 3D rendering, matching the quality of Neural Radiance Fields (NeRFs) [24] but with significantly faster speeds even suitable for end-user devices [15]. This technique enhances control and editability since scenes are stored as editable sets of Gaussians, allowing for modifications, merging, and other manipulations. Additionally, the original 3DGS method has been refined to improve anti-aliasing [48] and adapt density control more effectively [45]. Owing to these advancements, 3DGS has been widely adopted in various applications such as large-scale reconstructions [21], 3D generation [10], simultaneous localization and mapping (SLAM) [14, 19], and open-set segmentation [28]. Despite its state-of-the-art rendering quality and impressive handling of complex scenes, 3DGS’s explicit representation nature limits its use in scenarios requiring the depiction of semi-transparent volumes, such as underwater reconstructions where light scattering and absorption are significant challenges [15].

2.3. Computer Vision in Scattering Media

There are many challenges in underwater computer vision. The complex lighting conditions including scattering and attenuation of light leading to distorted images and the failure of traditional algorithms trained on clear-air scenes [12]. SeaThru [2] introduces a method for removing water from underwater images. This method addresses color distortion in underwater images by revising the image formation model from [1], accurately estimating backscatter, and correcting colors along the depth axis. WaterNeRF[33] estimates medium parameters separately from rendering, while ScatterNeRF[30] extends NeRF’s volumetric rendering to model scattering in adverse weather. [51] introduces a neural reflectance field to jointly learn scene albedo, normals, and medium effects, achieving color consistency by modeling light attenuation and backscatter through logistic regres-

sion and differentiable volume rendering. [38] explores hybrid neural-explicit approaches for reconstructing dynamic underwater environments, though its reliance on volumetric rendering limits real-time applicability.

SeaThru-NeRF [18] implements the image formation model [1] that separates direct and backscatter components, into the NeRF rendering equations, which is highly specialized for underwater scenes. We implemented a similar model but on Gaussian Splatting, which yields higher performance and enables real-time rendering.

3. Method

We start by briefly reviewing 3DGS and the rendering model in scattering media in Sec. 3.1. Then, we illustrate our proposed rendering model combining 3DGS with medium encoding in Sec. 3.2. At last, we explain our proposed loss function to align 3DGS with human perception of HDR scenes in Sec. 3.3.

3.1. Preliminaries

3D Gaussian Splatting models a scene with explicit learnable primitives $\mathcal{G}_0, \mathcal{G}_1, \dots, \mathcal{G}_N$. Each Gaussian \mathcal{G}_i is defined by a central position μ_i and covariance matrix Σ_i [15]:

$$G_i(p) = e^{-\frac{1}{2}(p-\mu_i)^T(\Sigma_i)^{-1}(p-\mu_i)}. \quad (1)$$

3DGS primitive also has two additional parameterized properties: opacity o_i and spherical harmonics coefficients SH_i to represent directional appearance component (anisotropic color). In order to render pixel-wise color, primitives are transformed into camera space. This involves a viewing transformation W , the Jacobian of the affine approximation of the projective transformation J on Σ_i and a projection matrix P . Then we get the projected 2D means $\hat{\mu}_i$ and 2D covariance matrices $\hat{\Sigma}_i$ as:

$$\hat{\Sigma}_i = (JW\Sigma_iW^TJ^T)_{1:2,1:2}, \quad \hat{\mu}_i = (P\mu_i)_{1:2}, \quad (2)$$

and the depth of \mathcal{G}_i on z-coordinate:

$$s_i = (P\mu_i)_3. \quad (3)$$

The Gaussian kernel \hat{G}_i of 2D Gaussian is represented as:

$$\hat{G}_i(p) = e^{-\frac{1}{2}(p-\hat{\mu}_i)^T(\hat{\Sigma}_i)^{-1}(p-\hat{\mu}_i)}, \quad (4)$$

where p is the coordinate of the pixel. For rasterization, each Gaussian is truncated at 3 sigma, considering only those intersecting with the patch comprising 16×16 pixels within this range, based on the property that about 99.7% of the probability lies within 3 sigma of the mean. Thus, the pixel color is computed by alpha blending of the sorted intersected Gaussians \mathcal{G}_i whose α_i are higher than a threshold:

$$C = \sum_{i=1}^N c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad \alpha_i = \sigma(o_i) \cdot \hat{G}_i(p), \quad (5)$$

where c_i is the color given the view direction, $\sigma(\cdot)$ is the Sigmoid function and N is the number of Gaussians involved in alpha blending. During optimization, 3DGS periodically densify Gaussians with high average gradient on 2D coordinates $\hat{\mu}_i$ across frames via splitting large ones and duplicating small ones. In the meantime, 3DGS prunes primitives with low opacity for acceleration and periodically set α_i close to zero for all Gaussians to moderate the increase of floaters close to the input cameras.

For **scene rendering in scattering media** we use the revised underwater image formation model from [1] where the final image I is separated into a direct and backscatter component

$$I = \underbrace{O \cdot e^{-\beta^D(\mathbf{v}_D) \cdot z}}_{\text{Direct Image component}} + \underbrace{B^\infty \cdot (1 - e^{-\beta^B(\mathbf{v}_B) \cdot z})}_{\text{Backscatter Image component}}, \quad (6)$$

where O is the clear scene captured at depth z in no medium, B^∞ is the backscatter color of the water at infinite distance. The colors get multiplied with attenuations, where the β^D and β^B are attenuation coefficients for the direct and backscatter components of the image which represent the effect the medium has on the color. The vector \mathbf{v}^D represents the dependencies for the direct component, which includes factors such as the depth z , reflectance, ambient light, water scattering properties, and the attenuation coefficient of the water. The vector \mathbf{v}^B represents the dependencies for the backscatter component, which includes ambient light, water scattering properties, the backscatter coefficient, and the attenuation coefficient of the water.

3.2. Splatting with Medium

We illustrate the pipeline of our method in Fig. 2. The input to our model is a set of images with scattering medium and corresponding camera poses. We initialize a set of 3D Gaussians via SfM [15] and optimize them with medium properties encoded by a neural network. Under the occlusion of both primitives and medium, our model acquires the transmittance along the ray and is capable of synthesizing medium component and object component in the novel view. Below we derive the whole model in detail.

Considering the expected color of a pixel integrated along the camera ray $r(s) = o + d(s)$ from the camera to infinitely far $C(r) = \int_0^\infty T(s)\sigma(s)c(s)ds$ [13] because of unbounded rendering of 3DGS [15], we release the constraints on light traveling in clear air to through a scattering medium [18] by adding the medium term:

$$C(r) = \int_0^\infty T(s)(\sigma^{\text{obj}}(s)c^{\text{obj}}(s) + \sigma^{\text{med}}(s)c^{\text{med}}(s))ds \quad (7)$$

$$T(s) = \exp(-\int_0^\infty (\sigma^{\text{obj}}(s) + \sigma^{\text{med}}(s))ds), \quad (8)$$

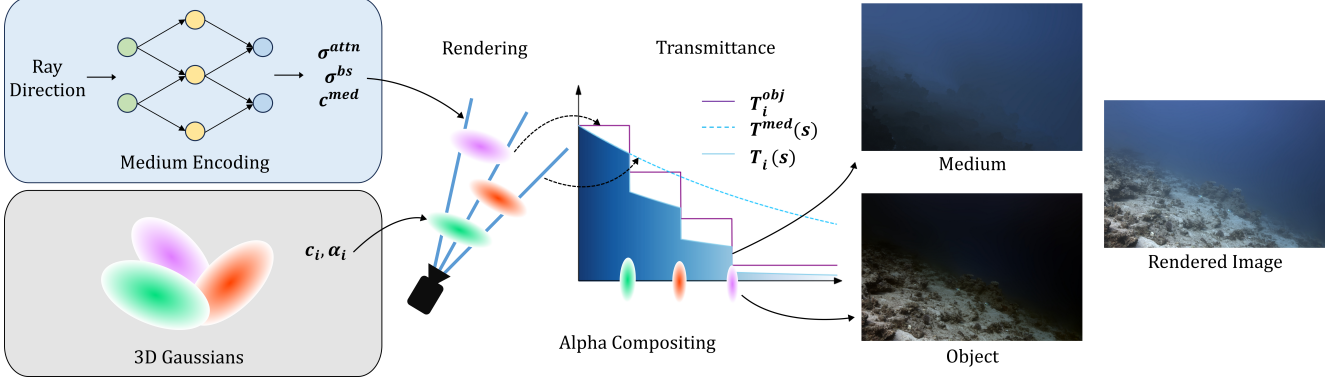


Figure 2. **Splatting with Medium**: We start rendering by casting a ray per pixel and collect the patch-intersected Gaussians along the ray and their color given ray direction. Then, we walk through the list of sorted Gaussians per pixel, query their opacity and depth, based on which we acquire the transmittance of both Gaussians and medium, rendering the Gaussians and the segments between adjacent two to obtain the **Medium** component and the **Object** component.

where $\sigma^{obj}/\sigma^{med}$ and c^{obj}/c^{med} are density and color of objects and medium respectively.

Following [18], we take σ^{med} and c^{med} to be constant per ray and separate per color channel. In order to apply discretized representation in 3DGS, the transmittance $T_i(s)$ in front of the i -th Gaussian \mathcal{G}_i (and behind $(i-1)$ -th Gaussian \mathcal{G}_{i-1}) with depth $s \in [s_{i-1}, s_i]$ can be decomposed as

$$T_i(s) = T_i^{obj} T^{med}(s), \quad T_i^{obj} = \prod_{j=1}^{i-1} (1 - \alpha_j) \quad (9)$$

where T_i^{obj} is the accumulated transmittance contributed by previous primitives' occlusion [15] and

$$T^{med}(s) = \exp\left(-\int_0^s \sigma^{med}(s) ds\right) = \exp(-\sigma^{med} s) \quad (10)$$

is the accumulated transmittance under the effect of medium from the camera to depth s . Then, the color is composed with discretized Gaussians and integrable medium

$$C(r) = \sum_{i=1}^N C_i^{obj}(r) + \sum_{i=1}^N C_i^{med}(r). \quad (11)$$

The contributed color of the \mathcal{G}_i to final output is

$$\begin{aligned} C_i^{obj}(r) &= T_i^{obj} T^{med}(s_i) \alpha_i c_i \\ &= T_i^{obj} \alpha_i c_i \exp(-\sigma^{med} s_i), \end{aligned} \quad (12)$$

where α_i is the opacity given the relative position between the pixel p and μ_i in Eq. (5) and $c_i = c_i^{obj}$ is the color given the ray direction. The contributed color of the medium between the $(i-1)$ -th and \mathcal{G}_i is

$$\begin{aligned} C_i^{med}(r) &= \int_{s_{i-1}}^{s_i} T_i^{obj} T^{med}(s) \sigma^{med} c^{med} ds \\ &= T_i^{obj} c^{med} [\exp(-\sigma^{med} s_{i-1}) - \exp(-\sigma^{med} s_i)]. \end{aligned} \quad (13)$$

To precisely estimate the properties of medium, we also include the background medium term from the last Gaussian \mathcal{G}_N to infinitely far

$$\begin{aligned} C_\infty^{med}(r) &= \int_{s_N}^{\infty} T_i^{obj} T^{med}(s) \sigma^{med} c^{med} ds \\ &= T_i^{obj} c^{med} \exp(-\sigma^{med} s_N) \end{aligned} \quad (14)$$

into the accumulated color.

As discussed in [1], the effective σ^{med} experienced a camera with wide-band color channels by differs in $C_i^{obj}(r)$ and $C_i^{med}(r)$, following [18], we use two sets of parameters, object attenuation σ^{attn} and medium back-scatter σ^{bs} for $C_i^{obj}(r)$ and $C_i^{med}(r)$ respectively. By setting $s_0 = 0$, our final equations of rendered color are:

$$C(r) = \sum_{i=1}^N C_i^{obj}(r) + \sum_{i=1}^N C_i^{med}(r) + C_\infty^{med}(r), \quad (15)$$

$$C_i^{obj}(r) = T_i^{obj} \alpha_i c_i \exp(-\sigma^{attn} s_i), \quad (16)$$

$$C_i^{med}(r) = T_i^{obj} c^{med} [\exp(-\sigma^{bs} s_{i-1}) - \exp(-\sigma^{bs} s_i)], \quad (17)$$

$$C_\infty^{med}(r) = T_N^{obj} c^{med} \exp(-\sigma^{bs} s_N). \quad (18)$$

3.3. Loss Function Alignment

In vanilla 3DGS, the loss function is a combination of \mathcal{L}_1 loss and D-SSIM loss. In low-light situations, [25] proposed a regularized \mathcal{L}_2 loss

$$\mathcal{L}_{\text{Reg-}\mathcal{L}_2} = ((sg(\hat{y}) + \epsilon)^{-1} \odot (\hat{y} - y))^2, \quad (19)$$

to boost the weight of the dark regions in optimization to align with how humans perceive dynamic range, which is applied to underwater scene reconstruction by [18]. For

the case of our 3DGS-based model, we propose a regularized loss function \mathcal{L}_{Reg} : we apply pixel-wise weight $W = \{w_{i,j}\}$ on both rendered estimate \hat{y} and target image y , where $w_{i,j} = (sg(\hat{y}_{i,j}) + \epsilon)^{-1}$ with pixel coordinate (i, j) and $sg(\cdot)$ denotes stopping gradient of its argument, which backpropagates zero derivative.

By introducing the weight W to the loss function of 3DGS composing of \mathcal{L}_1 and D-SSIM loss, we have the regularized \mathcal{L}_1 loss

$$\mathcal{L}_{\text{Reg-}\mathcal{L}_1} = |W \odot (\hat{y} - y)|, \quad (20)$$

and the regularized D-SSIM loss

$$\mathcal{L}_{\text{Reg-DSSIM}} = \mathcal{L}_{\text{DSSIM}}(W \odot y, W \odot \hat{y}). \quad (21)$$

$\mathcal{L}_{\text{Reg-DSSIM}}$ prioritizes structural similarity and perceptual fidelity, particularly in dark regions where human perception is more sensitive. Integrating regularization into the $\mathcal{L}_{\text{Reg-DSSIM}}$ formulation becomes particularly critical for 3DGS optimization due to the discrete nature of its primitives, necessitating structural regularization to maintain perceptual consistency across independently optimized elements - a distinct requirement from NeRF-based approaches where parameter-shared volumetric representations can achieve adequate low-light adaptation through pixel-level regularized losses alone. To model the smoothness of volumetric medium, we employ $\mathcal{L}_{\text{Reg-}\mathcal{L}_2}$ as our pixel-level loss. Our final proposed loss function is

$$\mathcal{L}_{\text{Reg}} = (1 - \lambda)\mathcal{L}_2 + \lambda\mathcal{L}_{\text{Reg-DSSIM}}. \quad (22)$$

4. Experiments

Implementation Details: Our implementation is based on the reimplemented version of 3DGS released by NeRF-Studio [44]. Following [45, 49], we accumulate the norms of the individual pixel gradients of μ_i for primitive densification. For the medium encoding, we use a spherical harmonic encoding [39] and a MLP with 2 linear layers with 128 hidden units and Sigmoid activation, followed by Sigmoid activation for c^{med} and Softplus activation for σ^{attn} and σ^{bs} . Upon each densification and pruning step of the 3DGS, the moving averages in the Adam optimizer of the medium encoding are reset to ensure the independence of subsequent iterations. We reset opacity to 0.5 every 500 training steps and prune gaussians with opacity below 0.5 every 100 training steps. The culling and reset thresholds [15] are set to 0.5 to encourage Gaussians to model opaque objects.

SeaThru-NeRF Dataset: SeaThru-NeRF Dataset released by [18] contains real-world scenes acquired from four different scenes in sea: IUI3 Red Sea, Curaçao, Japanese Gardens Red Sea, and Panama. There are 29, 20, 20 and 18 images respectively, among which 25, 17, 17 and 15 images are used for training and the rest 4, 3, 3 and 3 are used for

validation. The dataset encompasses a variety of water and imaging conditions. The images were captured in RAW format using a Nikon D850 SLR camera housed in a Nauticam underwater casing with a dome port, which helped prevent refractions that could disrupt the pinhole model. These images were then downsampled to an approximate resolution of 900×1400 . Prior to processing, the input linear images underwent white balancing with a 0.5% clipping per channel to eliminate extreme noise pixels. Lastly, COLMAP [32] was employed to determine the camera poses and correct lens distortions inherent.

Simulated Dataset: To further evaluate the performance of the proposed method, we take a standard NeRF dataset - the Garden scene from the Mip-NeRF 360 dataset [4] - and added fog to it to simulate the presence of medium. We used 3DGS to extract the depth maps. These maps were then utilized to create scenarios simulating both underwater and foggy conditions. In line with method Eq. (6), we utilized the following parameters to simulate easy foggy scenario: $\beta^D = [0.6, 0.6, 0.6]$, $\beta^B = [0.6, 0.6, 0.6]$, and $B^\infty = [0.5, 0.5, 0.5]$. The parameters for the hard foggy case are: $\beta^D = [0.8, 0.8, 0.8]$, $\beta^B = [0.6, 0.6, 0.6]$, and $B^\infty = [0.5, 0.5, 0.5]$.

Baseline methods: All methods were trained on the same set of white-balanced images. For rendering scenes with the medium, we compare several NeRF techniques: SeaThru-NeRF [18], the reimplement of SeaThru-NeRF released on NeRF-Studio (SeaThru-NeRF-NS) [34], ZipNeRF [5] and 3DGS [15]. For each baseline method, we use the PSNR, SSIM, and LPIPS [50] metrics to compare rendering quality. We present the alpha blending of depth as the depth map and the rendering without medium to demonstrate the ability to decouple the medium and the object for SeaThru-NeRF and our method. We also calculate the FPS and total training time using the same A100 GPU to illustrate the speed difference between baseline methods and our method. All of the results were averaged over three runs.

4.1. Results

First, we evaluated the performance of our method using the standard benchmark dataset, the SeaThru-NeRF Dataset. Table 1 compares PSNR, SSIM, LPIPS, average FPS and average training time with the baseline methods across the validation sets of four scenes. Our method demonstrates its superiority in the majority of cases and efficiency on both rendering and training. 'Panama' is a special case where there is little medium and the medium properties stays the same across different ray directions. Therefore, ZipNeRF can reconstruct it well. However, ZipNeRF training takes orders of magnitude more time than our method and does not offer real-time rendering.

Fig. 3 shows that the mainstream 3DGS and NeRF approaches fail to reconstruct well scenes in the presence of

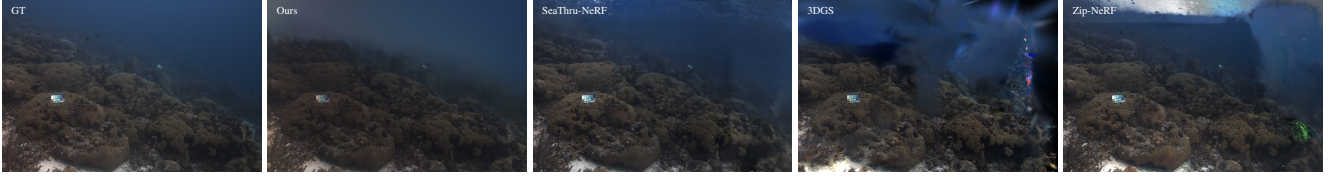


Figure 3. Underwater scene rendering in the 'Curacao' scene. From left to right: white-balanced ground-truth image, our result, SeaThru-NeRF's result, 3DGS' result, and Zip-NeRF's result. Both traditional 3DGS and NeRF with a proposal sampler cannot handle semi-transparent medium well.

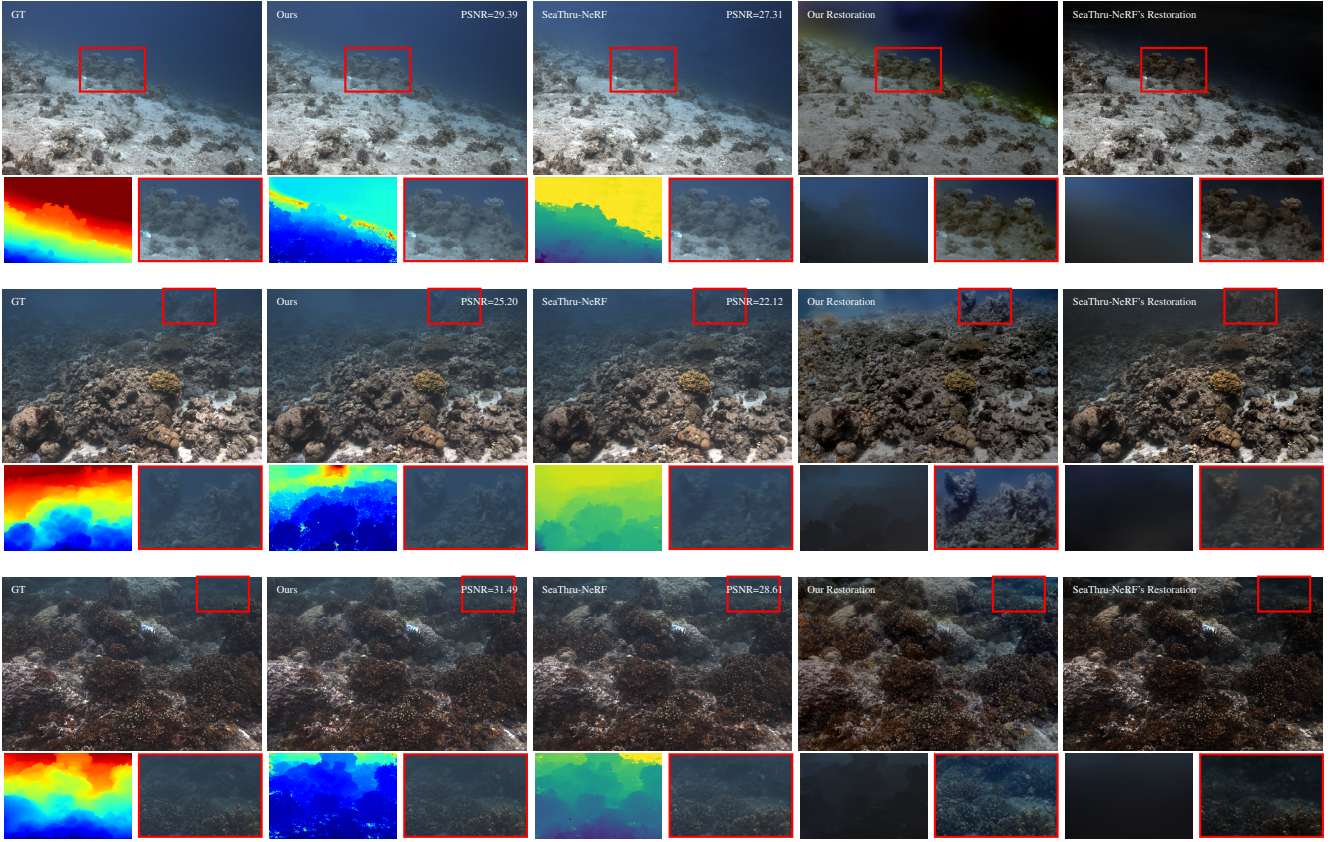


Figure 4. Underwater scene rendering in the 'IUI3 Red Sea' scene, 'Japanese Gardens Red Sea' scene and 'Panama' scene. We compare our method with SeaThru-NeRF by showing both the full image and the rendering without the medium. Furthermore, under each image, we show the depth maps (for GT the depth map from pre-trained model [43]), and highlighted region from the image. For Restoration, we further show the rendered medium without rendering objects. Our method achieves better rendering quality and preserves finer distant geometric details while reducing the amount of floaters.

Table 1. **Quantitative evaluation on the SeaThru-NeRF dataset.** We show PSNR \uparrow , SSIM \uparrow , LPIPS \downarrow , Avg. FPS \uparrow , and Avg. Training Time \downarrow .

Dataset/Metric	IUI3 Red Sea			Curaçao			J.G. Red Sea			Panama			Avg. FPS	Avg. Time
Method	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS		
SeaThru-NeRF	27.310	0.861	0.221	29.305	0.898	0.214	22.116	0.802	0.232	28.608	0.868	0.202	0.07	10h
SeaThru-NeRF-NS	27.651	0.885	0.133	30.609	0.928	0.159	23.542	0.889	0.116	31.857	0.948	0.071	0.9	2h
ZipNeRF	29.350	0.899	0.106	29.925	0.938	0.124	23.446	0.883	0.136	32.335	0.956	0.064	0.9	6h
3D Gauss.	22.980	0.843	0.2458	28.313	0.873	0.221	21.493	0.854	0.216	29.200	0.893	0.152	412.1	17.4min
Ours	29.386	0.910	0.180	32.672	0.957	0.110	25.201	0.904	0.113	31.490	0.948	0.075	41.8	9.4min



Figure 5. Simulated scene rendering with the easy foggy scene (upper) and hard foggy scene (lower). We compare our method with SeaThru-NeRF by showing both the full image and the rendering without the attenuation (restoration). Furthermore, under each image, we show the depth maps (for GT the depth map from pre-trained model [43]), and highlighted region from the image. For restoration, we further show the rendered medium without rendering objects. Our results exhibit better restoring quality and reasonable depth map compared to SeaThru-NeRF-NS’ results.

back-scattering media. 3DGS prunes the Gaussians with low opacity, leaving dense and muddy cloud-like primitives to fit the medium, which causes artifacts in the novel views. Zip-NeRF struggles to model the geometrical surface, leading to an unreal scene reconstructed with little media left.

Fig. 4 demonstrates that in both the ‘IUI3 Red Sea’, the ‘Japanese Gardens Red Sea’ and ‘Panama’ scenes, our method delivers superior quality and more effectively separates medium and object than the SeaThru-NeRF, especially in deeper and more complex scenes (as highlighted in the red square). Additionally, our depth map reveals much finer details compared to SeaThru-NeRF, which struggles to produce a reasonable depth map at greater distances, as indicated by the red color in the upper right corner of the depth map. We also achieve higher PSNR values in both scenes. The same advantages are observed in simulated scenes, where our method renders better details (indicated by the red square) than the SeaThru-NeRF in both easy and hard foggy scenes as depicted in Fig. 5. Our rendering without medium and depth maps significantly outperform those from the SeaThru-NeRF, especially in scenes that are farther from the camera. While our method’s predictions may appear blurry and the object map unclear in the upper right corner for the hard foggy scene, the results from SeaThru-NeRF are considerably worse. The restoration quality comparison presented in Table 2 further quantitatively demonstrates the superiority of our method in simulated scenes. Overall, our method surpasses SeaThru-NeRF in both underwater and simulated scenes.

Table 2. **Restoration Performance.** (PSNR↑/SSIM↑/LPIPS↓)

Dataset/Metric Method	Foggy-Easy			Foggy-Hard		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
SeaThru-NeRF-NS	13.11	0.32	0.58	10.76	0.29	0.63
Ours	15.70	0.37	0.56	14.06	0.45	0.54

4.2. Ablation Study

We isolate the different contributions to show their importance. We conduct a quantitative analysis on different combination of loss functions, between pixel-wise component $\{\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_{\text{Reg-}\mathcal{L}_1}, \mathcal{L}_{\text{Reg-}\mathcal{L}_2}\}$ and frame-wise $\{\mathcal{L}_{\text{DSSIM}}, \mathcal{L}_{\text{Reg-DSSIM}}\}$, as well as removing the medium effect and removing both the medium and our proposed loss function \mathcal{L}_{Reg} as a reimplement of 3DGS. These comparisons are made across validation sets for the SeaThru-NeRF dataset in Table 3.

Fig. 6 clearly demonstrates that our frame-level $\mathcal{L}_{\text{Reg-DSSIM}}$ improves the reconstruction quality and structural similarity of distant objects in low light. When used alone, $\mathcal{L}_{\text{Reg-}\mathcal{L}_1}$ can also provide better details in far distance. Used with $\mathcal{L}_{\text{Reg-DSSIM}}$, pixel-level $\mathcal{L}_{\text{Reg-}\mathcal{L}_1}$ and $\mathcal{L}_{\text{Reg-}\mathcal{L}_2}$ can further enhance reconstruction quality. However, $\mathcal{L}_{\text{Reg-}\mathcal{L}_2}$ alone [25] is not sufficient to train 3DGS-based models, and $\mathcal{L}_{\text{Reg-DSSIM}}$ is required. Our proposed \mathcal{L}_{Reg} shows superiority over other configurations in leading the 3DGS-based model to better fit HDR scenes and removing the medium component (basically 3DGS) significantly harms the performance of our method, which indicates the necessity of our approach.



Figure 6. **Ablation Study: loss function alignment.** Our proposed $\mathcal{L}_{\text{Reg-DSSIM}}$ improves the reconstruction quality of distant details in dark areas, and the benefit is obvious even when used alone. Regularized pixel-level losses $\mathcal{L}_{\text{Reg-L}_2}$ and $\mathcal{L}_{\text{Reg-L}_1}$ further improve the quality of the reconstruction.

Table 3. **Ablation Study** Avg. over SeaThru-NeRF Scenes

Configuration	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
$\mathcal{L}_1 + \mathcal{L}_{\text{DSSIM}}$	29.219	0.915	0.161
$\mathcal{L}_2 + \mathcal{L}_{\text{DSSIM}}$	28.652	0.911	0.172
$\mathcal{L}_1 + \mathcal{L}_{\text{Reg-DSSIM}}$	29.574	0.929	0.121
$\mathcal{L}_2 + \mathcal{L}_{\text{Reg-DSSIM}}$	29.169	0.924	0.132
$\mathcal{L}_{\text{Reg-L}_1} + \mathcal{L}_{\text{DSSIM}}$	28.890	0.916	0.153
$\mathcal{L}_{\text{Reg-L}_2} + \mathcal{L}_{\text{DSSIM}}$	29.064	0.914	0.165
$\mathcal{L}_{\text{Reg-L}_1} + \mathcal{L}_{\text{Reg-DSSIM}}$	29.603	0.928	0.123
$\mathcal{L}_{\text{Reg-L}_2} + \mathcal{L}_{\text{Reg-DSSIM}}$ (Ours)	29.687	0.930	0.119
$\mathcal{L}_1 + \mathcal{L}_{\text{DSSIM}}$ w/o Medium	29.227	0.915	0.164
$\mathcal{L}_{\text{Reg-L}_1} + \mathcal{L}_{\text{Reg-DSSIM}}$ w/o Med.	29.351	0.923	0.134
$\mathcal{L}_{\text{Reg-L}_2} + \mathcal{L}_{\text{Reg-DSSIM}}$ w/o Med.	29.353	0.925	0.130



Figure 7. **Limitation: simulating distant medium with Gaussians.** Our method (left) models distant medium with Gaussians, while SeaThru-NeRF [18] (right) also struggles with the background.

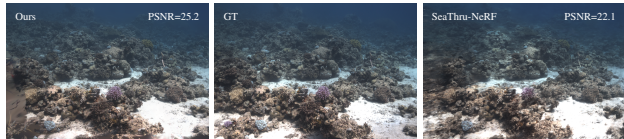


Figure 8. **Limitation: insufficient supervision.** Our method (left) has low-detail visuals in regions not sufficiently covered by training views. SeaThru-NeRF [18] (right) suffers from blurring in these regions.

5. Limitations

Although our method achieves good reconstruction quality, there are some limitations to consider. Firstly, our method, similar to NeRF-based approaches [18], has some difficul-

ties with distinguishing the background-like object and the medium far in the distance, as illustrated on the top of Fig. 3 and Fig. 7. However, in the foreground, our method prunes medium-role primitives well while SeaThru-NeRF cannot prevent the geometrical field from fitting the medium, resulting in wave-like artifacts. Secondly, same as other NVS methods [15, 24], our method relies on the camera poses being available which might prove difficult to obtain in underwater 3D scenes. Thirdly, our 3DGS-based method has artifacts in regions lacking observation [15], which is also suffered by NeRF-based models, as illustrated in the left side of Fig. 8 and the top part of Fig. 4, while the NeRF-based SeaThru-NeRF approach (right image) will introduce some blurring, distortion and interpolation. Lastly, the restored color of the scene from the scene is not ensured to be precise (especially for the background-like object), as under the effect of medium, the color of object and the attenuation attribute are entangled during training, which is shown by Fig. 5.

6. Conclusions

In our work, we focused on the problem of underwater reconstruction, previously tackled by fully volumetric representations that are slow to train and render. Therefore, we proposed to fuse the explicit point-splatting method (3DGS) with volume rendering to achieve both fast training and real-time rendering speed. Our method interleaves alpha compositing of splatted Gaussians with integrated ray segments passing through the scattering medium. We have demonstrated that our method achieves state-of-the-art results while enabling real-time rendering. Furthermore, the explicit scene representation enables disentanglement of geometry and the scattering medium. In future work, we would like to extend our method for larger scenes with both water and fog.

Acknowledgements. Jonas Kulhanek was supported by the Czech Science Foundation (GAČR) EXPRO (grant no. 23-07973X) and by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90254).

References

- [1] Derya Akkaynak and Tali Treibitz. A revised underwater image formation model. In *CVPR*, 2018. 2, 3, 4
- [2] Derya Akkaynak and Tali Treibitz. Sea-thru: A method for removing water from underwater images. In *CVPR*, 2019. 2
- [3] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields, 2021. 2
- [4] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022. 2, 5
- [5] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. *ICCV*, 2023. 2, 5
- [6] Jiazhong Cen, Zanwei Zhou, Jiemin Fang, Chen Yang, Wei Shen, Lingxi Xie, Dongsheng Jiang, Xiaopeng Zhang, and Qi Tian. Segment anything in 3d with nerfs. In *NeurIPS*, 2023. 2
- [7] Eric R. Chan, Koki Nagano, Matthew A. Chan, Alexander W. Bergman, Jeong Joon Park, Axel Levy, Miika Aittala, Shalini De Mello, Tero Karras, and Gordon Wetzstein. GeNVS: Generative novel view synthesis with 3D-aware diffusion models. In *arXiv*, 2023. 2
- [8] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvnerf: Fast generalizable radiance field reconstruction from multi-view stereo. *arXiv preprint arXiv:2103.15595*, 2021. 2
- [9] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision (ECCV)*, 2022. 2
- [10] Jaeyoung Chung, Suyoung Lee, Hyeongjin Nam, Jaerin Lee, and Kyoung Mu Lee. Lucidreamer: Domain-free generation of 3d gaussian splatting scenes, 2023. 2
- [11] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 2
- [12] Salma P. González-Sabbagh and Antonio Robles-Kelly. A survey on underwater computer vision. *ACM Comput. Surv.*, 55(13s), 2023. 2
- [13] James T. Kajiya and Brian P Von Herzen. Ray tracing volume densities. In *Proceedings of the 11th Annual Conference on Computer Graphics and Interactive Techniques*, page 165–174, New York, NY, USA, 1984. Association for Computing Machinery. 3
- [14] Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathan Luiten. Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. *arXiv*, 2023. 2
- [15] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. *ACM TOG*, 2023. 1, 2, 3, 4, 5, 8
- [16] Chung Min* Kim, Mingxuan* Wu, Justin* Kerr, Matthew Tancik, Ken Goldberg, and Angjoo Kanazawa. Garfield: Group anything with radiance fields. In *arXiv*, 2024. 2
- [17] Jonas Kulhanek and Torsten Sattler. Tetra-NeRF: Representing neural radiance fields using tetrahedra. *arXiv preprint arXiv:2304.09987*, 2023. 2
- [18] Deborah Levy, Amit Peleg, Naama Pearl, Dan Rosenbaum, Derya Akkaynak, Simon Korman, and Tali Treibitz. Seathru-nerf: Neural radiance fields in scattering media. In *CVPR*, pages 56–65, 2023. 1, 2, 3, 4, 5, 8
- [19] Mingrui Li, Shuhong Liu, Heng Zhou, Guohao Zhu, Na Cheng, Tianchen Deng, and Hongyu Wang. Sgs-slam: Semantic gaussian splatting for neural dense slam, 2024. 2
- [20] Chen-Hsuan Lin, Chaoyang Wang, and Simon Lucey. Sdf-srn: Learning signed distance 3d object reconstruction from static images. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [21] Jiaqi Lin, Zhihao Li, Xiao Tang, Jianzhuang Liu, Shiyong Liu, Jiayue Liu, Yangdi Lu, Xiaofei Wu, Songcen Xu, Youliang Yan, and Wenming Yang. Vastgaussian: Vast 3d gaussians for large scene reconstruction. In *CVPR*, 2024. 2
- [22] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object, 2023. 2
- [23] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Trans. Graph.*, 38(4):65:1–65:14, 2019. 2
- [24] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis, 2020. 1, 2, 8
- [25] Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul Srinivasan, and Jonathan T. Barron. Nerf in the dark: High dynamic range view synthesis from noisy raw images, 2021. 4, 7
- [26] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022. 2
- [27] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. 2
- [28] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. *arXiv preprint arXiv:2312.16084*, 2023. 2
- [29] Chen Quei-An. *nerf_pl*: a pytorch-lightning implementation of nerf, 2020. 2
- [30] Andrea Ramazzina, Mario Bijelic, Stefanie Walz, Alessandro Sanvito, Dominik Scheuble, and Felix Heide. Scatternerf: Seeing through fog with physically-based inverse neural rendering, 2023. 2
- [31] Sara Fridovich-Keil and Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *CVPR*, 2023. 2
- [32] Johannes L Schonberger and Jan-Michael Frahm. Structure from-motion revisited. *CVPR*, 2016. 5
- [33] Advait Venkatraman Sethuraman, Manikandasri Ram Srinivasan Ramanagopal, and Katherine A. Skinner.

- WaterNerf: Neural radiance fields for underwater scenes, 2023. 2
- [34] Paul Setinek, Lukas Mosser, and Lluís Guasch. A nerfstudio implementation of SeaThru-NeRF, 2023. MIT License. 5
- [35] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, 2022. 2
- [36] Matthew Tancik, Vincent Casser, Xichen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P. Srinivasan, Jonathan T. Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8248–8258, 2022. 2
- [37] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Justin Kerr, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, David McAllister, and Angjoo Kanazawa. Nerfstudio: A modular framework for neural radiance field development. In *ACM SIGGRAPH 2023 Conference Proceedings*, 2023. 2
- [38] Yunkai Tang, Chengxuan Zhu, Renjie Wan, Chao Xu, and Boxin Shi. Neural underwater scene representation. In *CVPR*, 2024. 3
- [39] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T. Barron, and Pratul P. Srinivasan. Ref-NeRF: Structured view-dependent appearance for neural radiance fields. *CVPR*, 2022. 5
- [40] Peng Wang, Yuan Liu, Zhaoxi Chen, Lingjie Liu, Ziwei Liu, Taku Komura, Christian Theobalt, and Wenping Wang. F2-nerf: Fast neural radiance field training with free camera trajectories. *CVPR*, 2023. 2
- [41] Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P. Srinivasan, Dor Verbin, Jonathan T. Barron, Ben Poole, and Aleksander Holynski. Reconfusion: 3d reconstruction with diffusion priors. *arXiv*, 2023. 2
- [42] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5438–5448, 2022. 2
- [43] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv:2406.09414*, 2024. 6, 7
- [44] Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. gsplat: An open-source library for Gaussian splatting. *arXiv preprint arXiv:2409.06765*, 2024. 5
- [45] Zongxin Ye, Wenyu Li, Sidun Liu, Peng Qiao, and Yong Dou. Absgs: Recovering fine details for 3d gaussian splatting, 2024. 2, 5
- [46] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. 2
- [47] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021. 2
- [48] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. *arXiv:2311.16493*, 2023. 2
- [49] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient and compact surface reconstruction in unbounded scenes, 2024. 5
- [50] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, Los Alamitos, CA, USA, 2018. IEEE Computer Society. 5
- [51] Tianyi Zhang and Matthew Johnson-Roberson. Beyond nerf underwater: Learning neural reflectance fields for true color correction of marine imagery, 2023. 2