Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

Full Length Article

Fast underwater scene reconstruction using multi-view stereo and physical imaging

Shuyi Hu^D, Qi Liu^D*

School of Future Technology, South China University of Technology, Guangzhou, 511442, China

ARTICLE INFO

Keywords: 3D scene reconstruction Novel view synthesis Multi-view stereo Underwater scene reconstruction

ABSTRACT

Underwater scene reconstruction poses a substantial challenge because of the intricate interplay between light and the medium, resulting in scattering and absorption effects that make both depth estimation and rendering more complex. While recent Neural Radiance Fields (NeRF) based methods for underwater scenes achieve high-quality results by modeling and separating the scattering medium, they still suffer from slow training and rendering speeds. To address these limitations, we propose a novel method that integrates Multi-View Stereo (MVS) with a physics-based underwater image formation model. Our approach consists of two branches: one for depth estimation using the traditional cost volume pipeline of MVS, and the other for rendering based on the physics-based image formation model. The depth branch improves scene geometry, while the medium branch determines the scattering parameters to achieve precise scene rendering. Unlike traditional MVSNet methods that rely on ground-truth depth, our method does not necessitate the use of depth truth, thus allowing for expedited training and rendering processes. By leveraging the medium subnet to estimate the medium parameters and combining this with a color MLP for rendering, we restore the true colors of underwater scenes and achieve higher-fidelity geometric representations. Experimental results show that our method enables high-quality synthesis of novel views in scattering media, clear views restoration by removing the medium, and outperforms existing methods in rendering quality and training efficiency.

1. Introduction

Underwater scene reconstruction is a crucial area of research with wide-ranging applications in marine science, underwater archaeology, ecology, and geoscience. Conventional methods of subaquatic investigation predominantly depend on divers and remotely operated vehicles (ROVs); however, these methodologies are frequently hindered by operational obstacles including restricted visibility, high expenses, and the requirement for specialized skills. Through the development of precise 3D depictions of submerged environments, researchers can enhance their ability to examine intricate underwater terrains, investigate marine biodiversity, and track alterations in underwater ecosystems throughout time.

Images taken underwater often experience a decrease in quality because of the distinct challenges presented by the aquatic environment. When light passes through water, it experiences attenuation that varies with distance and is sensitive to different wavelengths, as well as backscatter effects. Attenuation leads to the loss of certain colors, particularly red hue, while accentuating blue and green hues. Backscatter, in contrast, introduces a veil or haze over the image, diminishing clarity. The severity of these effects is contingent upon the proximity of the subject to the camera, in addition to its distance from the sea surface. These characteristics present a unique challenge when attempting to reconstruct the geometry of an underwater scene, due to the varying scattering properties of the medium compared to air.

In the field of Novel View Synthesis (NVS), Neural Radiance Fields (NeRF) (Mildenhall, Srinivasan et al., 2021) represented scenes as continuous volumetric fields encoded by multilayer perceptions (MLP), enabling photorealistic rendering through volume rendering technique. 3D Gaussian Splatting (3D-GS) (Kerbl, Kopanas, Leimkühler, & Drettakis, 2023) introduced differentiable rasterization for view synthesis, utilizing anisotropic 3D Gaussian primitives for explicit scene representation. Multi-View Stereo (MVS) (Yao, Luo, Li, Fang, & Quan, 2018) methods aggregated 2D information into 3D geometric perception representations by constructing cost volumes for depth estimation. However, these aforementioned methods are typically designed for clear media such as air, where light propagation is essentially unaffected by the medium. As a result, these methods are unable to achieve satisfactory outcomes when applied directly to underwater scenes, primarily due to inaccuracies in color and density estimations while conducting 3D reconstruction.

* Corresponding author. E-mail addresses: ft_shuyi@mail.scut.edu.cn (S. Hu), drliuqi@scut.edu.cn (Q. Liu).

https://doi.org/10.1016/j.neunet.2025.107568

Received 20 January 2025; Received in revised form 3 April 2025; Accepted 29 April 2025 Available online 14 May 2025 0893-6080/@ 2025 Elsevier Ltd. All rights are reserved including these for text and data mining. At training, or

0893-6080/© 2025 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.





Recently, a NeRF-based underwater scene reconstruction method, called SeaThru-NeRF (Levy et al., 2023), was proposed, achieving stateof-the-art quality by separating the medium. However, due to the inherent limitations of the NeRF method, its training and rendering speed were too slow. An in-depth analysis of the image formation model (Akkaynak & Treibitz, 2018) revealed that the depth prior played a crucial role. The MVSNet method, as described by Yao et al. (2018), was a widely recognized approach for estimating depth from multiple perspectives, utilizing a cost volume generated from differentiable homologous transformations. This indicates a favorable prospect for integrating these two methodologies.

We propose a novel method that combines MVS with a physically based underwater image formation model, which is divided into two branches: one branch for depth estimation based MVS, and the other for rendering based on the physical image formation model. Depth estimation is performed using a traditional cost volume construction pipeline (Yao et al., 2018). We design a medium subnet to estimate the parameters of the imaging model and combine a color MLP for efficient rendering. By estimating the medium parameters and scene depth, we can restore the scene's true colors while improving the accuracy of its geometric representation, enabling higher-fidelity and geometrically consistent rendering from novel viewpoints. Furthermore, our proposed approach does not rely on ground-truth depth required by MVSNet (Yao et al., 2018) pipeline and substantially improves both training and rendering speed when compared to SeaThru-NeRF. The source code is available in https://github.com/Azusa0811/uwMVS.

In summary, the contributions of this work are shown as follows:

- We, for the first time, propose a novel pipeline with multi-view stereo for underwater scene reconstruction, which can synthesize novel views in scattering media and restore clear views with the medium removed.
- We introduce a physics-based image formation model into MVS to infer the complete appearance of the scene from images without ground-truth depth, thereby enhancing the reconstruction quality of underwater environments and training efficiency.
- We propose a network that can independently retrieve the medium parameters of the imaging model for underwater scene reconstruction.

2. Related work

2.1. Novel view synthesis

NVS aims to generate new perspectives of a scene or object based on limited input views. Over the years, researchers have developed various methods to tackle this problem, spanning from conventional geometrycentered approaches to sophisticated deep learning techniques. Light field methods (Davis, Levoy, & Durand, 2012; Gortler, Grzeszczuk, Szeliski, & Cohen, 1996) captured dense arrays of light rays from multiple viewpoints, enabling realistic view synthesis but requiring extensive input data. Image-based rendering (Kalantari, Wang, & Ramamoorthi, 2016; Riegler & Koltun, 2020) interpolated views through warping and blending, working well for moderate viewpoint changes. However, both approaches face challenges with sparse inputs, occlusions, and large viewpoint deviations, limiting their applicability in complex scenes.

The advent of deep learning revolutionized NVS, enabling datadriven approaches to infer novel views directly from images (Basak, Corcoran, McDonnell, & Schukat, 2022). Neural representations (Jiang, Ji, Han, & Zwicker, 2020; Lombardi et al., 2019; Shih, Su, Kopf, & Huang, 2020; Wizadwongsa, Phongthawee, Yenphraphai, & Suwajanakorn, 2021) were widely used for novel view synthesis. Enhance novel view synthesis by improving multiplane images (Wizadwongsa et al., 2021) and layered depth images (Shih et al., 2020) using neural networks and inpainting, achieving better view-dependent effects and motion parallax. Jiang et al. (2020) utilize signed distance functions (SDFs) for high-quality, topology-aware 3D reconstruction. Lombardi et al. (2019) learn a volumetric representation of dynamic scenes through an encoder-decoder network and differentiable ray-marching. NeRF (Mildenhall, Srinivasan et al., 2021) and its derivatives marked a significant breakthrough in NVS. NeRF represented scenes as continuous volumetric fields encoded by multilayer perceptrons, enabling photorealistic rendering through volume rendering techniques (Barron et al., 2021; Barron, Mildenhall, Verbin, Srinivasan, & Hedman, 2022). Efficiency improvements had been a key focus, with methods like Mip-NeRF (Barron et al., 2021) and Instant-NGP (Müller, Evans, Schied, & Keller, 2022) significantly accelerating training and rendering using hierarchical sampling and optimized neural representations. PlenOctrees (Yu et al., 2021) and Plenoxels (Fridovich-Keil et al., 2022) improved efficiency by pre-computing scene representations and dividing large-scale environments into manageable regions. Extensions for dynamic scenes, such as D-NeRF (Pumarola, Corona, Pons-Moll, & Moreno-Noguer, 2021) incorporated temporal information to handle motion, while NeRF-W (Martin-Brualla et al., 2021) and Block-NeRF (Tancik et al., 2022) enhanced robustness in complex real-world environments. Few-shot and generalizable approaches, including PixelNeRF (Yu, Ye, Tancik and Kanazawa, 2021) and Depthsupervised NeRF (Deng, Liu, Zhu, & Ramanan, 2022), reduced dependency on dense multi-view data, expanding NeRF's applicability to sparse-view scenarios. Notwithstanding these advancements, NeRFbased approaches still encounter difficulties with slow training and rendering times, thereby impeding their wider acceptance.

Recently, 3D-GS (Kerbl et al., 2023) introduced differentiable rasterization as a novel approach to view synthesis, leveraging a set of anisotropic 3D Gaussian primitives for explicit scene representation. Its fast training speed and high-quality real-time rendering capabilities have made 3D-GS a prominent focus in the field of NVS. Many improvements have also been made to the 3D-GS method, such as antialiasing (Yan, Low, Chen, & Lee, 2024; Yu, Chen, Huang, Sattler, & Geiger, 2024), effectively density control (Ye, Li, Liu, Qiao, & Dou, 2024) and Sparse views reconstruction (Charatan, Li, Tagliasacchi, & Sitzmann, 2024; Yang et al., 2024). However, the explicit representation of 3D-GS is not well-suited for rendering semitransparent media, such as underwater scenes with light scattering and absorption.

2.2. Multi-view stereo method

MVS aims to reconstruct a dense 3D representation of a scene from multiple views. Traditional methods (Fua & Leclerc, 1995; Galliani, Lasinger, & Schindler, 2015; Schonberger & Frahm, 2016; Schönberger, Zheng, Frahm, & Pollefeys, 2016) relied on handcrafted features and similarity metrics, such as voxel-based and point-based approaches, and had laid the groundwork for this field. Fua and Leclerc (1995) proposed a surface reconstruction method based on object geometric center by combining multi-view stereo with light shadow cues for the first time. Galliani et al. (2015) proposed a massively parallel MVS method based on surface normal diffusion, but there are shortcomings in the reconstruction of thin structures. Schonberger and Frahm (2016) revised the Structure-from-Motion (SFM) process, provided robust camera pose estimation for disordered image sets through feature point matching. This work became the benchmark framework for modern MVS systems. Schönberger et al. (2016) proposed a pixel-level view selection strategy based on a probabilistic model to reduce the deep fusion artifacts in low-texture areas by evaluating the luminosity consistency weights between viewing angles. Voxel-based methods (Kutulakos & Seitz, 2000; Seitz & Dyer, 1999) evaluated the occupancy of discrete 3D units but were limited by high memory usage, while point-based methods, like PMVS (Furukawa & Ponce, 2010), focused on expanding reliable feature matches but may struggle in textureless regions.

Learning-based methods have become the mainstream due to their flexibility. With the development of deep learning, cost volume has been widely used for depth estimation in MVS methods (Zhou, Zhao, Wang, Hao, & Lei, 2023). MVSNet (Yao et al., 2018) proposed for the

first time an end-to-end pipeline that aggregated 2D information into a 3D geometric-aware representation by building a cost volume. The subsequent works follow the pipeline of multi-view 3D reconstruction with depth estimation from cost volume and improve from various aspects, such as improving memory consumption with recurrent plane sweeping (Yan et al., 2020; Yao et al., 2019) or coarse-to-fine architectures (Cheng et al., 2020; Gu et al., 2020; Yu & Gao, 2020). However, the aforementioned MVS networks require ground truth depth as the geometric supervision for training. Thus, in the field of novel view synthesis, (Chen et al., 2021, 2025; Lin et al., 2022; Liu et al., 2025; Oechsle, Peng, & Geiger, 2021) attempted to combine MVS methods with NeRF methods or 3D-GS methods. Unisurf combines MVS and NeRF to achieve high-quality 3D reconstruction without the need for an input mask through a unified implicit surface and radiance field representation. MVSNeRF (Chen et al., 2021) integrated traditional MVS cost volumes with NeRF, using 3D CNNs to predict radiance fields directly from multi-view images while maintaining view consistency. ENeRF (Lin et al., 2022) enhanced NeRF by incorporating MVS-derived depth priors and depth consistency losses to guide the radiance field optimization. Moving beyond implicit representations, MVSGaussian (Liu et al., 2025) converted MVS geometry into 3D Gaussians for real-time rendering, preserving geometric details through Gaussian parameterization. MVSplat (Chen et al., 2025) proposed a hybrid point-radiance field representation that combines dense MVS point clouds with continuous neural fields via learnable splatting, achieving both geometric accuracy and rendering efficiency. These methods demonstrate progressive improvements in either reconstruction quality, geometric constraints, or rendering speed, collectively advancing the state of 3D scene reconstruction and novel view synthesis. Our model extrapolates the complete appearance of the scene directly from images without the need for ground truth depth, drawing inspiration from the integration of these various methodologies.

2.3. Image processing in scattering media

Scattering media, such as underwater environments, often face challenges like color shift, image warping, and contrast reduction. These issues arise primarily due to complex lighting conditions, including light scattering and attenuation. Current underwater image enhancement (UIE) techniques can generally be divided into two main categories: non-deep learning-based approaches and deep learningbased approaches. CNN-based models (Chang, Wang, Du, & Xu, 2025; Li et al., 2021; Zhang et al., 2024) have achieved end-to-end underwater image restoration, while Transformer-based architectures (Peng, Zhu, & Bian, 2023) have further improved restoration outcomes. In contrast, non-deep learning methods are typically based on physical models and rely on prior assumptions. For example, Drews, Nascimento, Moraes, Botelho, and Campos (2013) employed the dark channel prior to estimate transmission maps specific to underwater conditions, whereas (Akkaynak & Treibitz, 2018) refined the atmospheric scattering model for more accurate underwater image restoration. However, these non-deep learning methods rely on the accuracy of prior assumptions such as depth information, which has limited their further development. This question can be well solved by 3D vision. Various works (Akkaynak & Treibitz, 2019; Ramazzina et al., 2023; Sethuraman, Ramanagopal, & Skinner, 2023) in underwater image processing have achieved remarkable results using the imaging model. These methods combine imaging models with NeRF to inspire underwater scene reconstruction using physical imaging. The state-of-the-art NeRFbased underwater scene reconstruction method, SeaThru-NeRF (Levy et al., 2023), incorporated the image formation model (Akkaynak & Treibitz, 2018) into the NeRF rendering equations by estimating direct and backscatter components.

3. Method

In this section, we detail the proposed MVS method for underwater scenes, which leverages an underwater image formation model to enhance MVS performance by better accounting for underwater optical characteristics. Section 3.1 introduces the preliminary. The overall framework of the method is described in Section 3.2. In Section 3.3, we detail the use of the multi-view approach for depth estimation. The application of the underwater image formation model is explained in Section 3.4. Finally, Section 3.5 describes the specific implementation details.

3.1. Preliminary

Underwater Image Formation Model. The image quality captured by underwater imaging systems is significantly hindered by the intricate nature of the underwater environment, predominantly because of light attenuation and scattering in water. The primary focus of underwater imaging theory is on these phenomena to elucidate and enhance the deterioration of underwater images.

The decrease in the intensity of light as it passes through water adheres to an exponential relationship, resulting from two distinct physical phenomena: absorption and scattering. The process of absorption leads to the dissipation of light energy, whereas scattering causes a redirection in the path of light propagation. Light attenuation is a multifaceted phenomenon influenced by wavelength, exhibiting varying levels of attenuation across different wavelengths, demonstrating selectivity. Red, yellow, and light green are more greatly attenuated in the visible spectrum than blue and green light, which are attenuated to a lesser degree. Consequently, underwater images generally display a blue-green coloration. The reduction of light limits the operational range of underwater imaging systems. The suspended particles and impurities found in the water are the main contributors to light scattering. Scattering effects can be classified into forward scattering and backscatter. Forward scattering pertains to the scattering of light at small angles as it reflects off the target surface before reaching the camera, causing the image to appear blurred. On the contrary, backscatter refers to light that enters the camera directly from natural or artificial light sources, following its scattering by suspended particles, causing a decrease in image contrast.

We adopt the revised underwater image formation model (Akkaynak & Treibitz, 2018) as the general model under ambient illumination. The final image I is decomposed into the direct component and the backscatter component as follows:

Direct

$$I = \overbrace{J \cdot exp(-\beta^{D}(\mathbf{v}_{D}) \cdot z)}^{D} + \underbrace{B^{\infty} \cdot (1 - exp(-\beta^{B}(\mathbf{v}_{B}) \cdot z))}_{\text{Backscatter}}$$
(1)

where *J* is the clear scene captured at depth *z* without medium, and B^{∞} is the backscatter water color at infinity distance. The β^D and β^B are attenuation coefficients of the direct and backscatter components, respectively. The colors will be multiplied with attenuation coefficients to represent the effects of the medium on the color. The vectors v_D and v_B represent the dependencies of β^D and β^B on range, object reflectance, spectrum of ambient light, spectral response of the camera, and the physical scattering and beam attenuation coefficients of the water body.

3.2. Overview

Given multi-view images, we aim to synthesize the target image from a novel camera pose in an underwater scene. The overview of our proposed method is depicted in Fig. 1. The initial step involves utilizing a Feature Pyramid Network (FPN) (Lin et al., 2017) to extract multi-scale features from the input multi-view source views. These



Fig. 1. Overview. We first use an FPN to extract image features from the source view and warp them into warped features $\{F_i^w\}_{i=1}^N$. The distorted characteristics are combined into a cost volume, which is then processed using a 3D CNN to generate depth. Subsequently, a pooling network is utilized to consolidate features for each 3D point at the predicted depth and then integrate them with the source image features. Following this, a color MLP and a medium-sized subnet are utilized to analyze these features, resulting in a refined image with the medium removed. Ultimately, the medium subnet is utilized once more to conduct additional processing on the image, resulting in the extraction of the backscatter and attenuation images. These are subsequently combined to produce the ultimate reconstructed outcome.

characteristics are subsequently transferred onto the designated camera frustum through differentiable homography to create a cost volume. This volume is then utilized by a 3D CNN for regularization to generate the depth map. The depth prediction branch of this system is constructed utilizing a cascading structure that propagates the depth map sequentially and systematically. Utilizing these depth maps, we aggregate multi-view and spatial information for each 3D point at the predicted depth by encoding features. Following this, we utilize a color MLP and the medium subnet to analyze the encoded features, which leads to the creation of a clear image with the medium removed. The intermediate subnet encodes the perspective of the view, integrates the depth map, and subsequently decodes the intermediate parameters. Finally, the medium subnet is used again to enhance the clear image by obtaining backscatter and attenuation images. These images are then combined to produce the final reconstructed result.

3.3. Depth estimation using MVS

The depth map produced by MVS is an essential component of our workflow, facilitating seamless integration with the underwater image formation model. Our approach is grounded in the insights of learning-based MVS methods (Yao et al., 2018).

Feature Map. We employ a Feature Pyramid Network (FPN) to extract multi-scale image features $\{F_i \mid i = 1, ..., N\}$ from input source multi-view images $\{I_i \mid i = 1, ..., N\}$. The structure of our implemented FPN is illustrated in Fig. 2. Specifically, each input image $I_i \in \mathbb{R}^{H \times W \times 3}$ is processed through the network to generate a low-resolution feature map $F_{i,1} \in \mathbb{R}^{H/4 \times W/4 \times 32}$. Subsequently, these features are fused using lateral connections and upsampling, producing two additional feature maps at higher resolutions: $F_{i,2} \in \mathbb{R}^{H/2 \times W/2 \times 16}$ and $F_{i,3} \in \mathbb{R}^{H \times W \times 8}$. This multi-resolution feature representation is then refined via smoothing convolutions to enhance the quality of the features before being passed to subsequent stages of the pipeline. The multi-scale feature extractor integrates high-resolution features with information from lower resolutions through a learned upsampling mechanism. This design ensures



Fig. 2. The structure of FPN. We employ an FPN to extract multi-scale image features.

that each stage in the multi-stage depth prediction process uses meaningful feature representations from previous stages, making it easier to extract high-frequency features effectively (Cheng et al., 2020).

Cost Volume. We first partition a set of sampling planes from an initial scene range depth $\{L_i \mid i = 1, ..., N\}$, and then warp image features of source views F_i onto D sweeping planes to build cost volume. This step requires the use of the differentiable homography,



Fig. 3. The structure of 3DCNN. We employ a 3D CNN for cost volume regularization.

as described below:

$$H_{i}(z) = K_{i}R_{i}\left(I + \frac{\left(R_{i}^{-1}t_{i} - R_{i}^{-1}t_{i}\right)a^{T}R_{i}}{z}\right)R_{i}^{-1}K_{i}^{-1}$$
(2)

where $[K_i, R_i, t_i]$ and $[K_t, R_t, t_t]$ denote camera intrinsic, rotation and translation of input source view and target view, respectively. *I* is the identity matrix and *a* is the principal axis of the target view camera. We use the matrix $H_i(z)$ warping the pixels at (u, v) from the source view to the target view at sampled depth z. Finally, we can obtain the warped feature maps of the target view, which can be defined as:

$$F_i^w(u, v, z) = F_i \cdot \left(\mathbf{H}_i(z) \cdot [u, v, 1]^T \right)$$
(3)

To build the cost volume *C*, we compute the variance of these warped multi-view feature volumes $\{F_i^w \mid i = 1, ..., N\}$, which is widely used in MVS (Cheng et al., 2020; Gu et al., 2020; Yao et al., 2018) for geometry reconstruction. For each voxel in *C*, located at the coordinates (u, v, z), we calculate its cost feature vector as:

$$C = \frac{\sum_{i=1}^{N} \left(F_i^w - \overline{F_i^w} \right)^2}{N} \tag{4}$$

where $\overline{F_{i}^{w}}$ is the average volume among all feature volumes. This cost volume is created utilizing variance to represent deviations in image appearance among various input perspectives. These discrepancies are a result of variations in scene geometry and view-dependent shading effects.

Depth Prediction from Probability Volumes. During this step, a 3D CNN is utilized to analyze the cost volume, producing a depth probability volume \hat{C} , as illustrated in Fig. 3. Afterwards, a depth distribution can be calculated based on this probability volume. The final depth prediction is computed by weighting each depth hypothesis using the depth probability distribution. Specifically, we obtain its probability P_i at a specific depth plane L_i by performing the softmax operation on the depth probability volume, namely:

$$P_i = Softmax(\hat{C}) \tag{5}$$

The depth value and its confidence at pixel (u, v) in the target view are defined as the weighted average $\hat{L}(u, v)$ and the standard deviation $\hat{\sigma}$ are calculated by

$$\hat{L}(u,v) = \sum_{i=1}^{D} P_i(u,v) \cdot L_i(u,v)$$
(6)

$$\hat{\sigma}(u,v) = \sqrt{\sum_{i=1}^{D} P_i(u,v) \cdot \left(L_i(u,v) - \hat{L}(u,v)\right)^2}$$
(7)

By conducting a calculation using the depth prediction and its variance to establish a confidence interval, we can measure the uncertainty of the prediction and identify the probable range of depth within which the object is positioned. That is:

$$\hat{U}(u,v) = [\hat{L}(u,v) - \lambda\hat{\sigma}(u,v), \hat{L}(u,v) + \lambda\hat{\sigma}(u,v)]$$
(8)

where λ is a scalar parameter that sets the size of the confidence interval. This depth range will become the initial depth range for fine processing.

Coarse-to-fine Prediction. The depth prediction pipeline is developed in a cascade structure, which allows for the propagation of the depth map in a coarse-to-fine manner. The objective is to accurately capture the geometric characteristics of the scene using multiple images. It then aims to generate a detailed depth prediction map for a specific viewpoint. Subsequently, we construct a high-resolution cost volume utilizing the depth map that was previously estimated. Through the processing of this refined volume, we are able to produce a more detailed depth map and a 3D feature volume. All depth maps are stored in the buffer for utilization in the subsequent phase.

3.4. Underwater image formation

Traditional MVS methods typically use cost volume only for geometric reconstruction, but recent works (Chen et al., 2021; Lin et al., 2022; Liu et al., 2025) show that it can also be leveraged to infer the complete appearance of the scene. Inspired by ENeRF (Lin et al., 2022), we propose to apply an underwater imaging model to MVS and reconstruct the entire underwater scene.

In Section 3.3, we obtain the image features $\{F_i \mid i = 1, ..., N\}$ from the input source views, as well as the depth maps \hat{L}_{tar} for the target view and $\{\hat{L}_{src,i} \mid i = 1, ..., N\}$ for the source views. In this section, we further process these to obtain the final reconstruction.

To effectively extract corresponding pixel-aligned features $\{f_i \mid i = 1, \ldots, N\}$ of target view from multi-view source image features, we first compute the 3D coordinates based on the depth map \hat{L}_{tar} of the target view and transform them from the world coordinate system to the image coordinate systems of the source cameras. By utilizing the converted 2D coordinates, we extract features at the respective locations from the feature maps of the source images using bilinear interpolation. For every 3D point, we compute the unit direction vectors in relation to both the target camera and the source cameras. This process involves standardizing the vectors and calculating differences in ray directions to account for variations in light directions. The extracted characteristics of the sampled image are combined with the calculated

ray difference vectors to create a cohesive feature representation. Finally, we used the pooling algorithm to aggregate all the pixel-aligned features to obtain the image features f_{img} of the target view, namely:

$$f_{\rm img} = Pooling(f_1, \dots, f_N) \tag{9}$$

Similarly, we translate the coordinates of each 3D point in the target view into a two-dimensional sampling grid, and then sample the warped multi-view feature volumes $\{F_i^w \mid i = 1, ..., N\}$ to obtain the grid features f_{grid} at the corresponding coordinates. Through the acquisition of feature representations at distinct locations, we augment the model's understanding of scene geometry. The color of a 3D point is determined by extending the color from the originating viewpoint along the path of the destination viewpoint. Specifically, the color from the source view is modulated by a blend weight w_i , which is computed using a color MLP φ_{color} via:

$$w_i = \varphi_{color}(f_{img}, f_{grid}, f_i) \tag{10}$$

Our model is derived from the image formation model described in Section 3.1, the color of each pixel in both the source views and the target view adheres to the following imaging equation:

$$\hat{c} = c^{\text{clr}} \cdot \exp(-\sigma^{atten} \cdot \hat{L}) + c^{\text{bs}} \cdot (1 - \exp(-\sigma^{bs} \cdot \hat{L}))$$
(11)

where \hat{c} represents the pixel color in water medium and $c^{\rm clr}$ is the clear image after accounting for the removal of the water medium. There are three medium parameters in the equation, σ^{atten} and σ^{bs} are the attenuation and scattering coefficients, respectively, and $c^{\rm bs}$ is the scattering color of the water medium.

To forecast medium parameters, a distinct subnet called the 'medium module' is employed. It computes these parameters based on the orientation of individual pixels within an image, aligned with the world coordinate system. This module consists of an encoder and a decoder. The medium encoder employs spherical harmonic encoding (Ramamoorthi & Hanrahan, 2001) to project the input direction map into a higher-dimensional space, thereby enhancing its ability to capture subtle directional variations and extract characteristic information across various frequencies.

$$\hat{d} = \operatorname{SHencoding}(d)$$
 (12)

where *d* is the input direction map and \hat{d} is the result after encoding. Subsequently, we put \hat{d} into medium decoder to get a base output s_{base} . We apply different activation functions to the base output to get the final media parameters.

$$s_{base} = \text{MediumMLP}(\hat{d}),$$
 (13)

$$\sigma^{atten}, \sigma^{bs} = \text{Softplus}(s_{base}), \tag{14}$$

$$c^{\rm DS} = {\rm Sigmoid}(s_{base}) \tag{15}$$

Based on Eq. (11), we first recover the clear image $\{c_{src,i}^{clr} \mid 1, ..., N\}$ of source views and then perform a weighted averaging operation to obtain the clear image c_{tar}^{clr} of the target view.

$$c_{tar}^{\text{chr}} = \sum_{i=1}^{N} w_i \cdot c_{src,i}^{\text{chr}}$$
(16)

We continue to use Eq. (11) to reconstruct the image in the water medium. Specifically, the attenuation coefficient is applied to diminish the clarity of the image in order to obtain the attenuated image, which is subsequently combined with the predicted scattering image to produce the ultimate reconstruction outcome.

$$\hat{I} = I_{atten} + I_{bs},\tag{17}$$

$$I_{atten} = c_{tar}^{clr} \cdot \exp(-\sigma^{atten} \cdot \hat{L}_{tar}),$$
(18)

$$I_{bs} = c^{bs} \cdot (1 - \exp(-\sigma^{atten} \cdot \hat{L}_{tar}))$$
⁽¹⁹⁾

where \hat{I} represents the final reconstruction result, while I_{atten} and I_{bs} denote the attenuated image and the scattering image, respectively.

3.5. Loss function

Our loss function follows the 3D-GS method with some improvements. In the initial 3D-GS method (Kerbl et al., 2023), the loss function is generally composed of an \mathcal{L}_1 Loss and a D-SSIM Loss as:

$$\mathcal{L} = (1 - \lambda)\mathcal{L}_1 + \lambda\mathcal{L}_{\text{D-SSIM}}$$
(20)

we generally set $\lambda = 0.2$. Inspired by Mildenhall, Hedman, Martin-Brualla, Srinivasan, Barron (2021), applying a loss function that significantly penalizes errors in dark areas, in accordance with the way human perception condenses dynamic range, can enhance the quality of reconstructing under-illuminated scenes. To be more specific, we adopt the reconstruction loss as follows:

$$\mathcal{L}_{\text{recon}}\left(\hat{C}, C^*\right) = \left(\frac{\hat{C} - C^*}{\operatorname{sg}(\hat{C}) + \epsilon}\right)^2 \tag{21}$$

where $sg(\cdot)$ stands for stop gradient and $\epsilon = 10^{-3}$, \hat{C} and C^* denote the reconstructed pixel color and the ground truth pixel color. We use the reconstruction loss instead of the \mathcal{L}_1 Loss to build the proposed final loss function as:

$$\mathcal{L} = (1 - \lambda)\mathcal{L}_{recon} + \lambda\mathcal{L}_{D-SSIM}$$
(22)

where $\lambda = 0.2$ is set in all our tests.

4. Experiments

4.1. Experimental setting

SeaThru-NeRF Dataset. SeaThru-NeRF (Levy et al., 2023) released a real forward-facing (Mildenhall, Srinivasan et al., 2021) dataset consisting of four multi-view underwater scenes captured in different sea regions: IUI3 Red Sea, Curaçao, Japanese Gardens Red Sea, and Panama. These four scenes include 29, 20, 20, and 18 images, respectively, with 25, 17, 17, and 15 images designated for training, while the remaining 4, 3, 3, and 3 images are reserved for validation. All images were captured in RAW format using a Nikon D850 SLR camera housed in a Nauticam underwater casing with a dome port, effectively minimizing refraction effects that could interfere with the pinhole camera model. The RAW images were then downsampled to an approximate resolution of 900 \times 1400. Prior to further processing, the linear input images were white-balanced with a 0.5% channel-wise clipping to eliminate extreme noise pixels. Finally, camera poses were estimated using COLMAP (Schonberger & Frahm, 2016).

Implementation Details. Our method is implemented using Py-Torch and trained on a single RTX 3090 GPU. We use the Adam (Kingma & Ba, 2014) optimizer and train the model for 3k iterations per scene. In practice, we utilize 16 and 8 depth planes for constructing the cost volumes at the coarse and fine levels, respectively. During training, 2, 3, or 4 source views are randomly chosen as inputs with probabilities of 0.1, 0.8, and 0.1, respectively.

For the Medium Subnet, the Direction Encoder utilizes spherical harmonic encoding at level 4. The Medium Decoder is an MLP consisting of two linear layers: the first layer contains 128 hidden units, while the second layer contains 64 hidden units, both with ReLU activation. The color MLP comprises two linear layers with 24 hidden units and ReLU activation.

Baseline Methods and Evaluation Metrics. To demonstrate improvements specific to underwater scenes, we first compare our method with other MVS-based approaches, such as ENeRF (Lin et al., 2022) and MVSGaussian (Liu et al., 2025), as well as the widely used 3D reconstruction method, 3D-GS. Most importantly, we conduct an in-depth comparison with SeaThru-NeRF (Levy et al., 2023), the state-of-the-art NeRF-based method for underwater scene reconstruction, to emphasize the advancements introduced by our method. For evaluation, we adopt the criteria established in prior works, including ENeRF (Lin et al., 2022), MVSNeRF (Chen et al., 2021), and MVSGaussian (Liu et al.,

Table 1

Quantitative evaluation results on the SeaThru-NeRF dataset. \uparrow indicates that larger values are better, while \downarrow signifies the opposite. Bolded values highlight the best results, and underlined values represent the second-best.

Scene	Cueaçao		IUI3 Red Sea		Panama		Japanese Gardens Red Sea		Avg.				
methods	PSNR ↑	SSIM \uparrow	LPIPS \downarrow	PSNR ↑	SSIM ↑	LPIPS \downarrow	PSNR ↑	SSIM ↑	LPIPS \downarrow	PSNR ↑	SSIM ↑	LPIPS \downarrow	time
SeaThru-NeRF	32.01	0.964	0.198	28.22	0.937	0.271	29.59	0.940	0.219	24.09	0.879	0.264	10 h
3D-GS	28.93	0.907	0.195	22.59	0.728	0.351	29.11	0.894	0.193	21.50	0.837	0.2188	29 min
ENeRF	21.49	0.724	0.620	18.76	0.713	0.469	19.85	0.671	0.608	18.22	0.634	0.565	15 h
MVSGaussian	30.27	0.911	0.232	27.42	0.905	0.244	27.38	0.877	0.237	25.59	0.916	0.187	1 h
Ours	32.29	0.942	0.187	31.82	0.953	0.241	29.95	0.924	0.206	26.44	0.940	0.180	15 min



Fig. 4. Restorations and depth maps. We contrast our approach with SeaThru-NeRF through the exhibition of renderings devoid of the medium. Underneath each image, we present the respective depth maps. Our restoration technique effectively preserves a greater level of color detail. In comparison to SeaThru-NeRF, the depth renderings exhibit a higher level of smoothness and coherence.

2025). For the Real Forward-facing (Mildenhall, Srinivasan et al., 2021) dataset, where the peripheral regions of images are typically not visible in the input views, we focus our evaluation on the central 80% area of the images. We use the widely-used PSNR, SSIM (Wang, Bovik, Sheikh, & Simoncelli, 2004), and LPIPS (Zhang, Isola, Efros, Shechtman, & Wang, 2018) metrics to compare the quality of synthesized views. All models are trained on the same GPU using the same data set to ensure fairness.

4.2. Results

Quantitative Results. To evaluate the effectiveness of our method, we compare its rendering quality against several baseline methods using the standard SeaThru-NeRF dataset. Table 1 presents a detailed comparison of PSNR, SSIM, LPIPS, and average training time across four distinct scenes for novel view synthesis. The results clearly demonstrate the advantages of our method, which surpasses other methods in most scenarios and shows notable enhancements in training efficiency. Our method demonstrates superior performance, resulting in an average PSNR enhancement of 1.65 dB compared to the leading NeRF-based underwater scene reconstruction method, SeaThru-NeRF. In combination with the scenes in Fig. 5, we can see that our improvement in PSNR is reflected mainly in two deeper scenes. Of utmost importance, our approach significantly decreases the time required for

Table 2

Statistical analysis of performance metrics. Δ shows the mean difference (Ours -
SeaThru-NeRF). Statistical significance is assessed using Wilcoxon signed-rank test, with
* indicating $p < 0.1$. Effect size (δ) is calculated using Cliff's delta, where $ \delta > 0.33$
represents medium effect.

Metric	SeaThru	Ours	Δ	p-val	δ	Sig. $(p < 0.1)$	
PSNR SSIM	28.44 0.9379	30.09 0.9351	+1.65 -0.0028	0.2274 0.8302	0.1243 -0.3373	-	
LPIPS	0.2402	0.2115	-0.0287	0.0955	-0.3136	*	

training. While our approach may not excel in all metrics, its consistent performance secured a second-place finish, closely behind the leader, yet it is noteworthy for its superior overall quality. Furthermore, our approach surpasses other MVS methods, demonstrating that the integration of MVS with a physical imaging model effectively enhances the ability to reconstruct underwater scenes.

Fig. 5 shows the novel view renderings in the water medium. To further emphasize the superiority of our approach, we zoom in on the areas highlighted by the red squares. As shown in Fig. 6, our approach maintains sharp geometric features in remote regions, demonstrating its capacity to accurately represent intricate scene geometry, particularly in deep and complex environments. Furthermore, our approach successfully isolates objects from the scattering medium. In particular, we strive to eliminate the presence of water in the image, thereby



Fig. 5. Novel view synthesis in water medium. The columns, from left to right, show the underwater scene renderings for the 'Cueaçao', 'IUI3 Red Sea', 'Japanese Gardens Red Sea', and 'Panama' scenes. The rendering quality of distant details, highlighted within the red squares, is compared and presented in Fig. 6.

presenting restored results that are clearer and based on the estimated medium parameters derived from our image formation model. Fig. 4 demonstrates the potential of our method for restoring the color of underwater scenes. Although SeaThru-NeRF's restored images may appear grayscale, our method preserves a greater amount of color detail. Furthermore, our depth maps exhibit increased smoothness and coherence in comparison to those generated by SeaThru-NeRF, thus suggesting a higher level of accuracy in depth estimation. This is particularly advantageous for our physics-based model, as the integration of MVS's depth estimation capabilities into the physical imaging model is essential for the success of our approach.

Statistical tests. We employ statistical testing methods to assess the robustness of our proposed approach compared to SeaThru-NeRF. Specifically, we conduct an in-depth analysis of the quantitative evaluation results from the SeaThru-NeRF dataset, collecting 13 pairs of test samples across four different scenes. For each evaluation metric (PSNR, SSIM, and LPIPS), we utilize the Wilcoxon signed-rank test to determine the statistical significance of performance differences. Given the small sample size (n = 13) and its potential limitations in statistical testing, we compute Cliff's Delta (δ) as an effect size measure and adopt a significance level of α = 0.1. This robust non-parametric metric quantifies the magnitude of the effect by assessing the dominance probability between the two data distributions.

As shown in Table 2, our method achieves an average PSNR improvement of +1.65 dB in test scenarios. Although this enhancement did not reach conventional statistical significance (p = 0.2274) due to limited sample size, the boxplot analysis in Fig. 7 reveals notable distributional advantages: our method demonstrates a higher median value and a more compact interquartile range (IQR), indicating reduced performance variability and greater stability across diverse scenarios. While showing marginally lower SSIM values compared to SeaThruNeRF, the difference is not statistically significant and our method exhibits fewer outliers. More importantly, our approach demonstrates superior performance in LPIPS metrics with a more concentrated distribution and narrower fluctuation range, achieving statistical significance at $\alpha = 0.1$ (p = 0.0955) with an effect size approaching medium



Fig. 6. The rendering of distant details. We compare our method with several baseline methods. Our method outperforms them in rendering quality and better preserves distant geometric details.



Fig. 7. Boxplots of statistical analysis. Overall, our method exhibits a more compact and concentrated distribution compared to SeaThru-NeRF, reducing performance variability and ensuring stable results across diverse scenarios.

magnitude ($|\delta| = 0.3136$). This evidence collectively confirms the practical advantages of our method in visual quality perception.

stability proves especially valuable for real-world deployments where environmental conditions may vary substantially.

The robustness of underwater imaging methods holds particular importance for practical applications. Complex underwater environments present substantial variations in lighting conditions and turbidity levels across different scenarios. In such contexts, consistent performance reliability and predictable lower-bound performance become crucial for ensuring satisfactory user experience. Compared with SeaThru-NeRF, our method offers enhanced stability and more dependable performance guarantees, particularly evidenced by its tighter distribution characteristics across all evaluation metrics. This operational

4.3. Ablation study

Ablation of key components. We conduct ablation experiments on the 'IUI3 Red Sea' scene to validate the contributions of the three key components of our framework: cascade structure, reconstruction loss, and medium subnet. As shown in Table 3, both the cascade structure and reconstruction loss lead to improvements in reconstruction quality,

Table 3

Ablation of key components on the 'IUI3 Red Sea' scene.

	<u>.</u>		
Cascade	Recon loss	Medium subnet	PSNR↑
	1	1	31.25
1		1	31.73
1	1		29.53
1	1	1	31.82

Table 4

Ablation of direction encoder on the 'Japanese Gardens Red Sea' scene.

Configuration	PSNR↑	SSIM↑	LPIPS↓
w/ SHEncdoing	26.44	0.940	0.180
w/o SHEncdoing	26.23	0.932	0.201

as measured by PSNR. However, the greatest advantage is derived from our proposed Medium Subnet, which provides the most substantial performance enhancement. This underscores the critical role of the Medium Subnet in enhancing the effectiveness of our approach, especially when integrated with physics-based image formation models.

Ablation of Direction Encoder We conduct ablation experiments on the 'Japanese Gardens Red Sea' scene to validate the contributions of the direction encoder in medium subnet. The results in Table 4 show that using high level spherical harmonic coding directions can improve the quality of reconstruction.

4.4. Discussion

The experimental results can explain why we chose to integrate MVS with the physical model rather than other possible approaches. SeathruNeRF has demonstrated that incorporating imaging models can effectively enhance underwater 3D reconstruction. However, due to the inherent limitations of NeRF, both training and inference require significant time, hindering its practical application. Therefore, we aim to develop a method that enables faster reconstruction to meet realworld usage demands. Building on imaging models, we analyzed their structure and found that depth is a crucial component. Accelerating depth estimation during training could potentially enhance reconstruction speed. To address this challenge, we seek a faster depth training method. MVSNet (Yao et al., 2018), a MVS method that achieves 3D reconstruction through multi-view depth estimation, offers a promising direction. However, it primarily focuses on geometric reconstruction and relies on ground-truth depth as geometric supervision. Inspired by recent works that integrate MVS with NeRF or 3DGS, we adopt only the cost volume pipeline of MVSNet for efficient depth estimation. By combining this with an underwater imaging model, we aim to infer the complete appearance of the scene from images without requiring ground-truth depth. Experimental results validate the effectiveness of our approach, demonstrating a significant improvement in training speed while also enhancing reconstruction metrics.

Interpretability is a crucial aspect of deep learning research, as it enhances trust in model predictions and facilitates further improvements (ŞAHiN, Arslan, & Özdemir, 2025). While our primary focus is on improving reconstruction efficiency, we acknowledge the importance of understanding how our model makes decisions. Our approach benefits from the integration of MVSNet's cost volume pipeline, which provides explicit depth estimation before scene appearance inference. This structure improves transparency by separating geometric and photometric reasoning, making the decision-making process more interpretable compared to end-to-end NeRF models. The loss function we designed for the underwater environment can also improve Perceptual Similarity Metrics (Güven, Şahin, & Talu, 2024). Additionally, incorporating an imaging model further strengthens physical consistency, allowing us to infer scene appearance in a way that aligns with real-world imaging principles.

5. Conclusion

We propose a method for underwater scene reconstruction by integrating MVS with a physical image formation model. By embedding the depth estimation capabilities of MVS directly into the imaging model, our approach achieves two key objectives: generating highquality novel view renderings within scattering media and accurately restoring the true colors of underwater scenes. This dual capability not only enhances the visual fidelity of rendered scenes but also provides a more accurate representation of the underwater environment, overcoming challenges associated with scattering effects and color distortion. Beyond the advancements in rendering quality, a standout feature of our methodology is the significant improvement in training efficiency than the top NeRF-based method. This efficiency enables faster processing times without compromising the accuracy or quality of the results, making our approach both practical and effective for real-world applications. In summary, the combination of superior rendering quality, efficient training, and adaptability to scattering media positions our method as a highly effective solution for underwater scene reconstruction.

Although our method represents an improvement, there are still the following directions for future improvements in underwater scene reconstruction. First, we still face challenges in handling objects with similar backgrounds in distant locations. For example, a distant blue water feature is mistakenly identified as being closer in our depth map (refer to Fig. 4). This issue is attributed in part to the feature extraction phase, especially through the FPN, which is influenced by the medium, affecting feature extraction accuracy. Future work will therefore focus on developing an improved FPN for more accurate feature extraction in underwater environments. Second, acquiring high-quality underwater scene data sets with removed water poses a challenge, leading to constraints on the quality of training data and affecting the accuracy of color recovery. In upcoming research endeavors, utilizing datasets consisting of underwater images alongside their corresponding waterremoval sets is expected to significantly improve the precision of color recovery.

CRediT authorship contribution statement

Shuyi Hu: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Data curation, Conceptualization. Qi Liu: Writing – review & editing, Supervision, Resources, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 62202174, inpart by the GJYC program of Guangzhou under Grant 2024D01J0081, in part by The Taihu Lake Innovation Fund for the School of Future Technology of South China University of Technology under Grant 2024B105611004, and in part by the Guangdong Provincial Key Laboratory of Human Digital Twin 2022B1212010004.

Data availability

Data will be made available on request.

S. Hu and Q. Liu

References

- Akkaynak, D., & Treibitz, T. (2018). A revised underwater image formation model. In 2018 IEEE/CVF conference on computer vision and pattern recognition (pp. 6723–6732).
- Akkaynak, D., & 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 (pp. 1682–1691).
- Barron, J. T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., & Srinivasan, P. P. (2021). Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer* vision (pp. 5855–5864).
- Barron, J. T., Mildenhall, B., Verbin, D., Srinivasan, P. P., & Hedman, P. (2022). Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 5470–5479).
- Basak, S., Corcoran, P., McDonnell, R., & Schukat, M. (2022). 3D face-model reconstruction from a single image: A feature aggregation approach using hierarchical transformer with weak supervision. *Neural Networks*, 156, 108–122.
- Chang, L., Wang, Y., Du, B., & Xu, C. (2025). Rectangling and enhancing underwater stitched image via content-aware warping and perception balancing. *Neural Networks*, 181, Article 106809.
- Charatan, D., Li, S. L., Tagliasacchi, A., & Sitzmann, V. (2024). Pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 19457–19467).
- Chen, A., Xu, Z., Zhao, F., Zhang, X., Xiang, F., Yu, J., et al. (2021). Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *ICCV* (pp. 14124–14133).
- Chen, Y., Xu, H., Zheng, C., Zhuang, B., Pollefeys, M., Geiger, A., et al. (2025). Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. In *European* conference on computer vision (pp. 370–386). Springer.
- Cheng, S., Xu, Z., Zhu, S., Li, Z., Li, L. E., Ramamoorthi, R., et al. (2020). Deep stereo using adaptive thin volume representation with uncertainty awareness. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 2524–2534).
- Davis, A., Levoy, M., & Durand, F. (2012). Unstructured light fields. vol. 31, In Computer graphics forum (pp. 305–314). Wiley Online Library.
- Deng, K., Liu, A., Zhu, J.-Y., & Ramanan, D. (2022). Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF conference on computer* vision and pattern recognition (pp. 12882–12891).
- Drews, P., Nascimento, E., Moraes, F., Botelho, S., & Campos, M. (2013). Transmission estimation in underwater single images. In Proceedings of the IEEE international conference on computer vision workshops (pp. 825–830).
- Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., & Kanazawa, A. (2022). Plenoxels: Radiance fields without neural networks. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 5501–5510).
- Fua, P., & Leclerc, Y. G. (1995). Object-centered surface reconstruction: Combining multi-image stereo and shading. *International Journal of Computer Vision*, 16(1), 35–56.
- Furukawa, Y., & Ponce, J. (2010). Accurate, dense, and robust multiview stereopsis. IEEE Transactions on Pattern Analysis and Machine Intelligence, 32(8), 1362–1376.
- Galliani, S., Lasinger, K., & Schindler, K. (2015). Massively parallel multiview stereopsis by surface normal diffusion. In Proceedings of the IEEE international conference on computer vision (pp. 873–881).
- Gortler, S. J., Grzeszczuk, R., Szeliski, R., & Cohen, M. F. (1996). The lumigraph. In Proceedings of the 23rd annual conference on computer graphics and interactive techniques (pp. 43–54). Association for Computing Machinery.
- Gu, X., Fan, Z., Zhu, S., Dai, Z., Tan, F., & Tan, P. (2020). Cascade cost volume for highresolution multi-view stereo and stereo matching. In *Proceedings of the IEEE/CVF* conference on computer vision and pattern recognition (pp. 2495–2504).
- Güven, S., Şahin, E., & Talu, M. (2024). Image-to-image translation with CNN based perceptual similarity metrics. *Computer Science*, http://dx.doi.org/10.53070/bbd. 1429596.
- Jiang, Y., Ji, D., Han, Z., & Zwicker, M. (2020). Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. In *Proceedings of the IEEE/CVF* conference on computer vision and pattern recognition (pp. 1251–1261).
- Kalantari, N. K., Wang, T.-C., & Ramamoorthi, R. (2016). Learning-based view synthesis for light field cameras. ACM Transactions on Graphics, 35(6), 1–10.
- Kerbl, B., Kopanas, G., Leimkühler, T., & Drettakis, G. (2023). 3D Gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics, 42(4).
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
- Kutulakos, K. N., & Seitz, S. M. (2000). A theory of shape by space carving. International Journal of Computer Vision, 38, 199–218.
- Levy, D., Peleg, A., Pearl, N., Rosenbaum, D., Akkaynak, D., Korman, S., et al. (2023). SeaThru-nerf: Neural radiance fields in scattering media. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 56–65).
- Li, C., Anwar, S., Hou, J., Cong, R., Guo, C., & Ren, W. (2021). Underwater image enhancement via medium transmission-guided multi-color space embedding. *IEEE Transactions on Image Processing*, 30, 4985–5000.

- Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., & Belongie, S. (2017). Feature pyramid networks for object detection. In CVPR (pp. 2117–2125).
- Lin, H., Peng, S., Xu, Z., Yan, Y., Shuai, Q., Bao, H., et al. (2022). Efficient neural radiance fields for interactive free-viewpoint video. In SIGGRAPH Asia 2022 conference papers (pp. 1–9).
- Liu, T., Wang, G., Hu, S., Shen, L., Ye, X., Zang, Y., et al. (2025). MvsGaussian: Fast generalizable Gaussian splatting reconstruction from multi-view stereo. In *European conference on computer vision* (pp. 37–53). Springer.
- Lombardi, S., Simon, T., Saragih, J., Schwartz, G., Lehrmann, A., & Sheikh, Y. (2019). Neural volumes: learning dynamic renderable volumes from images. ACM *Transactions on Graphics*, 38(4).
- Martin-Brualla, R., Radwan, N., Sajjadi, M. S., Barron, J. T., Dosovitskiy, A., & Duckworth, D. (2021). Nerf in the wild: Neural radiance fields for unconstrained photo collections. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 7210–7219).
- Mildenhall, B., Hedman, P., Martin-Brualla, R., Srinivasan, P. P., & Barron, J. T. (2021). Nerf in the dark: High dynamic range view synthesis from noisy raw images. In 2022 IEEE/CVF conference on computer vision and pattern recognition CVPR, (pp. 16169–16178).
- Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., & Ng, R. (2021). Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1), 99–106.
- Müller, T., Evans, A., Schied, C., & Keller, A. (2022). Instant neural graphics primitives with a multiresolution hash encoding. ACM Transactions on Graphics, 41(4), 1–15.
- Oechsle, M., Peng, S., & Geiger, A. (2021). Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In Proceedings of the IEEE/CVF international conference on computer vision (pp. 5589–5599).
- Peng, L., Zhu, C., & Bian, L. (2023). U-shape transformer for underwater image enhancement. *IEEE Transactions on Image Processing*, 32, 3066–3079.
- Pumarola, A., Corona, E., Pons-Moll, G., & Moreno-Noguer, F. (2021). D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 10318–10327).
- Ramamoorthi, R., & Hanrahan, P. (2001). An efficient representation for irradiance environment maps. In Proceedings of the 28th annual conference on computer graphics and interactive techniques (pp. 497–500).
- Ramazzina, A., Bijelic, M., Walz, S., Sanvito, A., Scheuble, D., & Heide, F. (2023). Scatternerf: Seeing through fog with physically-based inverse neural rendering. In Proceedings of the IEEE/CVF international conference on computer vision (pp. 17957–17968).
- Riegler, G., & Koltun, V. (2020). Free view synthesis. In Computer vision–ECCV 2020: 16th European conference, glasgow, UK, August 23–28, 2020, proceedings, part XIX 16 (pp. 623–640). Springer.
- ŞAHiN, E., Arslan, N. N., & Özdemir, D. (2025). Unlocking the black box: an in-depth review on interpretability, explainability, and reliability in deep learning. *Neural Computing and Applications*, 37(2), 859–965. http://dx.doi.org/10.1007/s00521-024-10437-2.
- Schonberger, J. L., & Frahm, J.-M. (2016). Structure-from-motion revisited. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 4104–4113).
- Schönberger, J. L., Zheng, E., Frahm, J.-M., & Pollefeys, M. (2016). Pixelwise view selection for unstructured multi-view stereo. In *Computer vision–ECCV 2016: 14th European conference, amsterdam, the netherlands, October 11-14, 2016, proceedings, part III 14* (pp. 501–518). Springer.
- Seitz, S. M., & Dyer, C. R. (1999). Photorealistic scene reconstruction by voxel coloring. International Journal of Computer Vision, 35, 151–173.
- Sethuraman, A. V., Ramanagopal, M. S., & Skinner, K. A. (2023). Waternerf: Neural radiance fields for underwater scenes. In OCEANS 2023-MTS/IEEE US gulf coast (pp. 1–7). IEEE.
- Shih, M.-L., Su, S.-Y., Kopf, J., & Huang, J.-B. (2020). 3D photography using contextaware layered depth inpainting. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 8028–8038).
- Tancik, M., Casser, V., Yan, X., Pradhan, S., Mildenhall, B., Srinivasan, P. P., et al. (2022). Block-nerf: Scalable large scene neural view synthesis. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 8248–8258).
- Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4), 600–612.
- Wizadwongsa, S., Phongthawee, P., Yenphraphai, J., & Suwajanakorn, S. (2021). Nex: Real-time view synthesis with neural basis expansion. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 8534–8543).
- Yan, Z., Low, W. F., Chen, Y., & Lee, G. H. (2024). Multi-scale 3d gaussian splatting for anti-aliased rendering. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 20923–20931).
- Yan, J., Wei, Z., Yi, H., Ding, M., Zhang, R., Chen, Y., et al. (2020). Dense hybrid recurrent multi-view stereo net with dynamic consistency checking. In *European conference on computer vision* (pp. 674–689). Springer.
- Yang, C., Li, S., Fang, J., Liang, R., Xie, L., Zhang, X., et al. (2024). Gaussianobject: High-quality 3d object reconstruction from four views with gaussian splatting. ACM Transactions on Graphics, 43(6), 1–13.
- Yao, Y., Luo, Z., Li, S., Fang, T., & Quan, L. (2018). Mvsnet depth inference for unstructured multi-view stereo. In ECCV (pp. 767–783).

S. Hu and Q. Liu

- Yao, Y., Luo, Z., Li, S., Shen, T., Fang, T., & Quan, L. (2019). Recurrent mvsnet for high-resolution multi-view stereo depth inference. In CVPR (pp. 5525–5534).
- Ye, Z., Li, W., Liu, S., Qiao, P., & Dou, Y. (2024). Absgs: Recovering fine details in 3d gaussian splatting. In Proceedings of the 32nd ACM international conference on multimedia (pp. 1053–1061).
- Yu, Z., Chen, A., Huang, B., Sattler, T., & Geiger, A. (2024). Mip-splatting: Alias-free 3d gaussian splatting. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 19447–19456).
- Yu, Z., & Gao, S. (2020). Fast-mysnet: Sparse-to-dense multi-view stereo with learned propagation and gauss-newton refinement. In *Proceedings of the IEEE/CVF conference* on computer vision and pattern recognition (pp. 1949–1958).
- Yu, A., Li, R., Tancik, M., Li, H., Ng, R., & Kanazawa, A. (2021). Plenoctrees for realtime rendering of neural radiance fields. In Proceedings of the IEEE/CVF international conference on computer vision (pp. 5752–5761).
- Yu, A., Ye, V., Tancik, M., & Kanazawa, A. (2021). Pixelnerf: Neural radiance fields from one or few images. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 4578–4587).
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., & Wang, O. (2018). The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR* (pp. 586–595).
- Zhang, D., Wu, C., Zhou, J., Zhang, W., Lin, Z., Polat, K., et al. (2024). Robust underwater image enhancement with cascaded multi-level sub-networks and triple attention mechanism. *Neural Networks*, 169, 685–697.
- Zhou, H., Zhao, H., Wang, Q., Hao, G., & Lei, L. (2023). Miper-MVS: Multi-scale iterative probability estimation with refinement for efficient multi-view stereo. *Neural Networks*, 162, 502–515.