

Initial Experiments on the Use of Radiance Fields for Underwater 3D Reconstruction

Original

Initial Experiments on the Use of Radiance Fields for Underwater 3D Reconstruction / Tanduo, Beatrice; Matrone, Francesca; Murtiyoso, Arnadi. - In: INTERNATIONAL ARCHIVES OF THE PHOTOGRAMMETRY, REMOTE SENSING AND SPATIAL INFORMATION SCIENCES. - ISSN 2194-9034. - ELETTRONICO. - XLVIII-M-9-2025:(2025), pp. 1475-1481. [10.5194/isprs-archives-xlvi-m-9-2025-1475-2025]

Availability:

This version is available at: 11583/3004124 since: 2025-10-16T14:16:52Z

Publisher:

Copernicus Publications

Published

DOI:10.5194/isprs-archives-xlvi-m-9-2025-1475-2025

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Initial Experiments on the Use of Radiance Fields for Underwater 3D Reconstruction

Beatrice Tanduo¹, Francesca Matrone², Arnadi Murtiyoso³

¹ Laboratory of Geomatics for Cultural Heritage (LabG4CH), Department of Architecture and Design (DAD), Politecnico di Torino, Viale Pier Andrea Mattioli, 39, Torino, Italy - beatrice.tanduo@polito.it

² Geomatics Lab, Department of Environment, Land and Infrastructure Engineering (DIATI), Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129, Torino, Italy - francesca.matrone@polito.it

³ Université de Strasbourg, INSA Strasbourg, CNRS, ICube Laboratory UMR 7357, Photogrammetry and Geomatics Group, 67000 Strasbourg, France - arnadi.murtiyoso@insa-strasbourg.fr

Keywords: Underwater Photogrammetry, NeRF, 3D Gaussian Splatting, 3D Modelling, Submerged Archaeology

Abstract

Underwater photogrammetry presents unique challenges, including light attenuation, refraction, and turbidity, that affect the accuracy and quality of 3D reconstructions. This study investigates the performance of novel neural rendering techniques, Neural Radiance Fields (NeRF), SeaThru-NeRF, and 3D Gaussian Splatting (3DGS), in comparison to conventional Structure-from-Motion (SfM) workflows. Using a dataset acquired during the SIFET benchmark campaign on a submerged Roman archaeological site, we processed image data via Nerfacto, SeaThru, and Jawset Postshot (3DGS) and compared outputs against a reference model produced in Agisoft Metashape. Evaluation criteria included processing time, geometric accuracy (via M3C2 analysis), point cloud density and roughness, and point cloud completeness. Results show that radiance fields-based methods significantly reduce processing time while providing competitive visual results. SeaThru-NeRF demonstrated the highest geometric accuracy, benefiting from underwater-specific corrections, while 3DGS offered photorealistic rendering. These findings highlight the potential of neural methods for underwater cultural heritage documentation, though further improvements are needed in data fidelity and robustness under challenging underwater conditions.

1. Introduction

1.1 Background

Despite the challenges, photogrammetry remains a valuable method in underwater archaeology due to its ability to generate diverse outcomes, particularly, but not limited to, texture-based outputs. When combined with sonar surveys, the question of data fusion parallels the integration of land-based photogrammetry and lidar (Menna et al., 2018).

Data acquisition aided by Remotely Operated Vehicles (ROV) has also been tested as a solution to minimise human intervention (Drap et al., 2015), with ongoing research focusing on the question of automation. Image data collected from these systems often requires a specific protocol for geometric distortion calibration due to the two-media nature of the environment (air through the lens cap and water).

Recent advancements in 3D rendering have seen a move towards employing neural-based solutions. Neural Radiance Fields (NeRF) (Mildenhall et al., 2021) presented a novel approach to perform this task, in which a Multi Layer Perceptron (MLP) is used to represent a continuous density function for light rays. This new 3D representation in the form of radiance fields is different from conventional point clouds but offers advantages, which includes a faster computation time and a richer output.

Another paradigm, 3D Gaussian splatting (3DGS) uses a similar image rendering process to NeRF, but discards the implicit nature of the MLP-based representation in favour of a more explicit 3D Gaussians. Using this explicit approach, 3DGS is able to perform a faster rendering of the 3D scene. It is, however, less suitable than NeRF or conventional Multi-View

Stereo (MVS) photogrammetry for overall 3D reconstruction since it relies on the anisotropic scaling of its elements to render a scene, instead of the presence of dense points.

These novel rendering paradigms offer promising alternatives to traditional photogrammetric methods, especially in complex environments like submerged archaeological sites. Their ability to reconstruct scenes quickly and with high visual fidelity, even under challenging lighting and visibility conditions, makes them particularly suited for documenting and monitoring underwater cultural heritage.

In addition to supporting accurate 3D recording, neural-based approaches like NeRF and 3DGS can facilitate the creation of realistic visualisations for interpretation and conservation planning, significantly enhancing the preservation, and dissemination of underwater heritage assets, supporting both scientific research and public engagement.

These tools are especially relevant where access is restricted or repeated data acquisition is impractical, as they can maximise the value of limited imagery. Additionally, their integration into heritage workflows could open new perspectives in the development of digital twins for submerged cultural sites.

1.2 Related works

These approaches have been recently applied also to underwater data for colour correction (Zhang and Johnson-Roberson, 2023) or to increase rendering quality (Li et al. 2024). SeaSplat (Yang et al., 2024), for example, integrates 3D Gaussian Splatting with a physically-based underwater image formation model to enable real-time rendering of underwater scenes. Applied to real datasets like SeaThru-NeRF, it enhances novel view synthesis and restores true scene colours by compensating for underwater

light distortions. However, further studies and analyses on data accuracy should still be carried out and deepened.

3D Gaussian splatting built upon the concept of radiance fields but proposes an explicit representation of the 3D rendering process (Kerbl et al., 2023). Instead of representing the radiance fields as a density function along ray lines like NeRF, 3DGS uses anisotropically scaled kernels with attributes such as colour, opacity, and size; the eponymous 3D Gaussians. For rendering purposes, stacked 3D Gaussians are simply combined, thus "splatted" into the user's screen. In this way, 3DGS avoids the costly neural rendering process of NeRF.

In terms of cultural heritage application, both NeRF and 3DGS has seen some uses, mostly to create 3D point clouds in comparison to conventional methods such as lidar and Multi-View Stereo (MVS) (Croce et al., 2023; Clini et al., 2024). However, its use in underwater environments remains scarce.

For this reason, in this contribution, we propose the comparison of the results obtained from SeaThru, the 3DGS and Nerfacto implemented in the software Jawset Postshot with the ground truth processed with Agisoft Metashape, in terms of time, metric accuracy, density and data completeness.

2. Materials and Methods

2.1 Case study and dataset

The dataset used for this study is the one acquired in 2019 within the context of the SIFET (Italian Society of Photogrammetry and Topography) benchmark. The surveyed site is an Italian underwater area of 18x10 m, at a depth of 5 to 7 m, containing three Roman columns (Figure 1), three squared blocks, and four irregular marble blocks. The photogrammetric survey was conducted using a Nikon D700 with a 20 mm fixed lens in an underwater housing with a hemispherical dome. The acquisition followed a scheme of 9 nadiral stripes (Figure 2) and two radial captures around the blocks to better record their vertical development for a total of 323 images (221 nadiral and 102 oblique).




Figure 1. Overview of the site and GCPs.




Figure 2. Example of the acquired images and stripes.

In addition, 15 Ground Control Points (GCPs), measured with direct trilateration, were placed and acquired (Figure 3). Due to the difficulty of maintaining a static position during the diving, it was necessary to set the ISO value to 1600 during the acquisition phase. This approach has led to higher noise in the captured images than traditional terrestrial acquisitions. To remove the blue dominant, typical of underwater images, a white balance was performed using a medium grey colour control board.




Figure 3. Trilateration operations for the GCPs measuring.

2.2 Research design

The datasets were processed using three radiance fields-based methods: SeaThru, 3DGS and Nerfacto. The latter pipeline (Tancik et al. 2022), a NeRF technique tuned for quicker convergence and superior scene representation, was used to handle the same dataset; however, due to the poor results obtained, it will be presented only in the Discussion section. On the other side, the SeaThru-NeRF technique (Levy et al., 2023) was used to process the dataset in order to solve issues with underwater imagery, such as color distortion and light attenuation. This technique enhances the quality of the recreated scene by incorporating corrections unique to underwater environment.

The same dataset was also processed using the software Jawset Postshot (<https://www.jawset.com/>, accessed 17 March 2025) exploiting the 3DGS method. This method produced very fast results with very photorealistic rendering when presented in the form of Gaussian splats. Indeed, a stable result was achieved after a few minutes, which is equivalent to around 8,000 iteration steps. A conversion to a more conventional point cloud produced a less dense result when compared to traditional MVS

and noticeably more noise. This noise is, nevertheless, easily reduced using standard point cloud noise cleaning algorithms.

The generated results were then compared to a point cloud generated using Agisoft Metashape as a reference. After an appropriate data registration and cleaning, several analyses were performed.

First, a look at the processing time. Indeed, one of the main advantages of radiance fields is its claim to a faster processing time when compared to traditional MVS. In this regard, a comparison of processing time was performed to assess whether the results can be obtained faster than MVS all while keeping their quality.

To assess the geometric quality, a cloud-to-cloud distance analysis was performed using the M3C2 plugin in the software CloudCompare. This method enables the computation of a signed distance, thus providing a normal distribution of the errors. From the normal distribution, statistical parameters such as average error and standard deviation are derived.

From CloudCompare, an additional analysis on the density and presence of noise in each dataset is also presented. This analysis enabled us to ascertain whether the radiance fields-based methods can provide a homogeneous result, as with conventional photogrammetry. In addition, the completeness of the point cloud is also assessed with regard to the Metashape point cloud reference. This analysis is equally important to determine whether this method is applicable for underwater heritage documentation purposes.

Finally, a visual analysis of the quality of the RGB colour of the resulting point clouds is also conducted. It is worth noting that color-based analysis is a challenge in underwater 3D reconstruction. SeaThru also proposes a colorimetric correction to specifically address this issue.

3. Results

The acquired data has been processed following three different workflows to assess the accuracy and effectiveness of novel NeRF and 3DGS-based approaches compared to traditional Structure-from-Motion (SfM) photogrammetry (Figure 4).

3.1 Processing time and density

Starting from a basic comparison of processing time (Table 1), the results show a clear improvement in time. Both neural radiance outputs required up to 5-7 times less processing time than the SfM point cloud (using NVIDIA GeForce RTX 4090 GPU). These times, however, must also be correlated with the density of the point cloud, which is clearly lower than the reference point cloud.

As far as the number of points is concerned (Table 1), a measurement unit of 1 m² has been selected for the analysis. Without any filtering operations, the SeaThru output ensured a higher density with respect to the 3D Gaussian Splatting solution, which has about 3/4 of fewer points. Both solutions have significantly fewer points with regards to the SfM point cloud; about 1/3 less for the SeaThru and more than 1/10 for PostShot.




Figure 4. Visual comparison of the SfM point cloud, processed with Agisoft Metashape (a); the NeRF result - SeaThru (b) and the 3DGS - PostShot (c).

	SfM	NeRF - SeaThru	3DGS - PostShot
Number of points	23 215 341	7 940 218	2 235 093
Processing time	110 min	20 min	15 min*
Density	150-200 000 points/m ²	50-60 000 points/m ²	10-15 000 points/m ²

*processed with a NVIDIA RTX 3000 Ada Generation Laptop GPU.

Table 1. Comparison of the number of points and the processing time.

3.2 Cloud-to-cloud distances – M3C2

A M3C2 analysis (which computes robust distances) has been carried out in CloudCompare in order to assess the precision and the accuracy of the geometric reconstruction.

The SfM point cloud was kept as a reference, while the SeaThru and PostShot point clouds were used for the comparison. In the first case, a mean of 0.003 m with a standard deviation of 0.089 m was obtained (Figure 5); while in the second case a mean of -

0.007 m with a standard deviation of 0.081 m was observed (Figure 6). It clearly appears that the results of the two approaches are fully comparable, with higher values in the partially submerged areas next to the seabed.

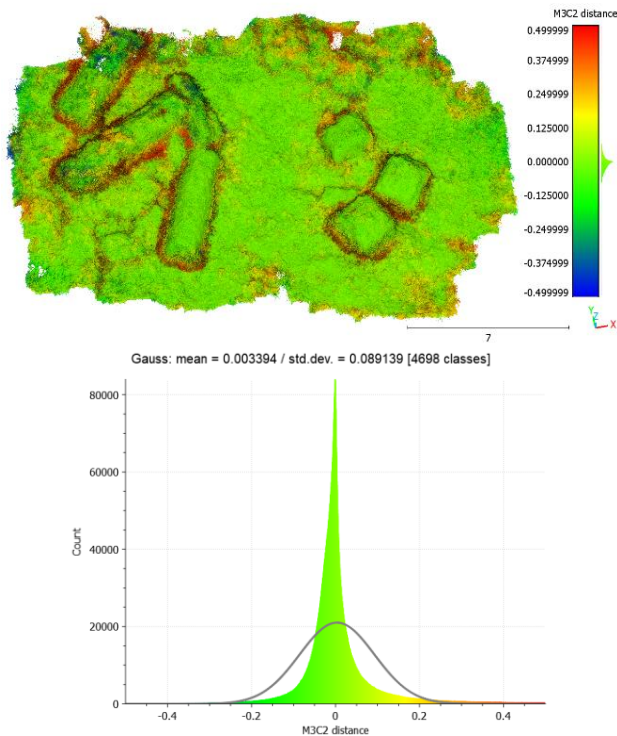


Figure 5. M3C2 analysis and Gaussian curve obtained from the SfM and NeRF – SeaThru point clouds comparison.

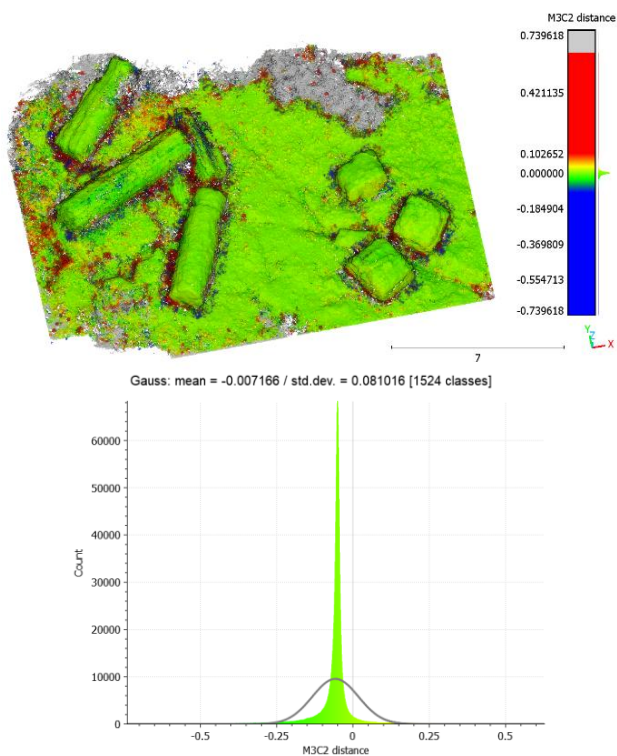


Figure 6. M3C2 analysis and Gaussian curve obtained from the SfM and 3DGS - PostShot point clouds comparison.

3.3 Roughness

The roughness of a point cloud typically refers to quantifying how much a point (or region) deviates from a locally fitted surface, to evaluate and highlight local surface irregularities. It has been computed in CloudCompare with a radius of 0.1 m.

The results show a mean of 0.015 m and a standard deviation of 0.011 m for the NeRF – SeaThru point cloud (Figure 7) and a mean of 0.014 m and a standard deviation of 0.013 m for the 3DGS – PostShot point cloud (Figure 8).

Also in this case the results are comparable and do not highlight particular differences. The behavior is homogeneous both for the reconstructed parts of the archaeological finds and for the seabed.

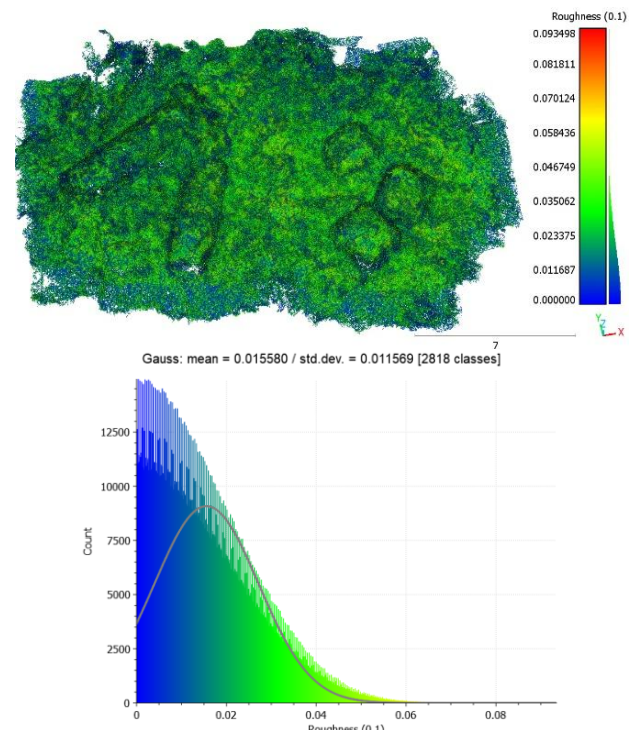



Figure 7. Roughness analysis for NeRF – SeaThru point cloud (radius 0.1 m).



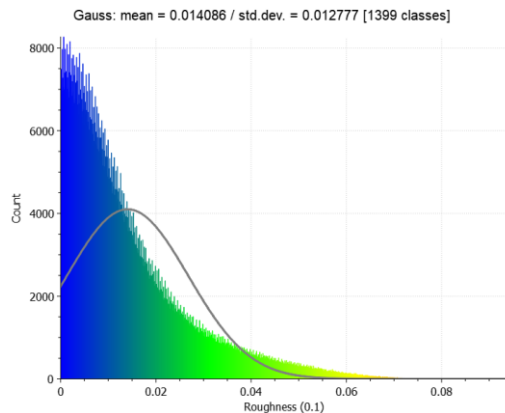


Figure 8. Roughness analysis for NeRF – SeaThru point cloud (radius 0.1 m).

3.4 Data completeness

The data completeness was investigated by subsampling the point clouds with a regular spacing of 0.005 m, then a cloud-to-cloud distance was computed with a maximum range of 0.002 m. This procedure allowed to highlight the areas of the point cloud with higher discrepancies.

In particular, red points represent those parts where no common points were identified between the SfM point cloud and the neural radiance solutions.

The results show for the NeRF point cloud a commonality of ~90% of the points, with red areas distributed both on the columns and on the border, where the multiplicity is lower (Figure 9). Part of the column on the top-right has not been reconstructed, but the reason has not been identified since, except for a slight variation in the radiometric data, the geometric content is similar to the next areas.

The 3DGS point cloud (Figure 10) shares ~85% of common points, and a missing part in a vegetated area on the top right is evident. This lack is due to the presence of *Posidonia Oceanica*, a moving seagrass that made the 3D reconstruction difficult.

Both the point clouds appear quite complete, but missing points are found on the side of the columns, in the partially submerged areas in contact with the seabed. This issue could be related to darker parts or shadows, as well as difficulties in capturing those points during the acquisition phases.




Figure 9. Data completeness result for the NeRF-SeaThru point cloud. Red points represent the missing parts.




Figure 10. Data completeness result for the 3DGS-PostShot point cloud. Red points represent the missing parts.

3.5 Additional tests and results

Nerfacto, despite being one of the methods tested in this paper, failed to generate a coherent result using underwater data.

As shown in Figure 11, the resulting point cloud is extremely noisy and lacks geometric fidelity. The reconstruction fails to accurately represent the surveyed objects and structures, producing a dispersed and incoherent cloud. A notable issue is the presence of large, irregular blue patches, which correspond to the rendered water surface appearing in every image. These artifacts introduce further distortion, making the model unsuitable for both qualitative visualization and quantitative analysis.



Figure 11. Point cloud generated using the Nerfacto NeRF pipeline.

Due to these limitations, the Nerfacto-derived model was excluded from the subsequent metric comparisons (e.g., M3C2 analysis, density and roughness evaluation, and completeness assessment). Nevertheless, a simple Cloud to Cloud (C2C) comparison has still been done, achieving an overall RMSE average value of 13 cm. These results highlight the current limitations of general-purpose NeRF pipelines in underwater applications, especially in complex conditions with light scattering and inconsistent textures.

It is therefore an interesting point to note that 3DGS was able to generate results without any specific modifications for an underwater setup, contrary to NeRF. Some of this may be explained by the way the radiance fields are represented by each method, i.e. either implicitly or explicitly.

Further tests were also carried out on a dataset acquired in turbid water conditions. The images were captured during the second edition of the SUNRISE (Seashore and UNDERwater documentation of aRchaeological herItage palimpSests and Environment) Summer School, held in Puglia, Italy, in

September 2024 in collaboration with the Italian Society of Photogrammetry and Topography (SIFET) and with the support of the International Society of Photogrammetry and Remote Sensing (ISPRS) and CIPA Heritage Documentation. It consists of 169 images of 5 Roman columns (Figure 12), however, the tests output a very rough and noisy point cloud, therefore, we decided not to include the results in this contribution, leaving an in-depth analysis for future works.




Figure 12. Example of images acquired with turbid water.

The suspended particulate matter in the water, in fact, has greatly affected the generation of the point cloud, not allowing the creation of a coherent model. Image enhancing strategies or the use of deep learning algorithms for feature matching such as SuperGlue (Sarlin et al., 2020) could overcome this problem. This, in essence, also showed the dependence of neural radiance methods on conventional photogrammetry and SfM; indeed, these novel methods still require a proper photogrammetric setup in order to perform an equally proper image orientation before scene rendering may even take place. In this sense, the radiance fields-based methods emerge not as a competition but rather as a complement to conventional photogrammetry. If anything, radiance fields present an alternative to MVS-based methods.

4. Discussions

The performance of NeRF techniques in underwater environments is still being studied, despite its reputation for being able to capture intricate features and challenging lighting circumstances. The initial experiments demonstrate that radiance field-based methods, such as NeRF and 3D Gaussian Splatting, are applicable to underwater environments, though not without limitations.

Key challenges such as unstable lighting conditions, reduced visibility, image noise, and geometric distortions due to light refraction across media interfaces continue to affect the quality and consistency of the reconstructions. Nevertheless, the results

are promising in terms of both geometric accuracy and visual fidelity.

Among the evaluated methods, SeaThru-NeRF yielded the most accurate geometric results, thanks to its integration of underwater-specific correction mechanisms. These corrections help mitigate the effects of light attenuation and scattering, enhancing both colour fidelity and spatial precision. 3DGS, on the other hand, achieved the most photorealistic renderings and produced stable outputs in extremely short timeframes, although it presented limitations in terms of point cloud density and noise levels. Nerfacto did not achieve results of comparable quality, likely due to its lack of domain-specific adjustments and the complex optical conditions of the underwater environment.

An additional advantage of these techniques lies in their high computational efficiency, which enables the generation of detailed 3D models in significantly shorter processing times compared to traditional photogrammetric methods. This efficiency makes them particularly attractive for rapid documentation, preliminary site analysis, or integration into automated pipelines.

However, the results also highlight the necessity of refining both the input data and the reconstruction algorithms. In particular, radiance field models remain sensitive to image quality, acquisition geometry, and environmental noise. Current implementations lack robustness when dealing with turbidity, occlusions, and varying light conditions, especially in shallow water sites.

Although these approaches are not yet ready to replace conventional workflows, they represent a valuable complementary tool, especially in complex or time-sensitive underwater scenarios. The ability of these methods to produce realistic renderings, even from suboptimal imagery, suggests great potential for enhancing the documentation and communication of underwater heritage sites.

5. Conclusions and future perspectives

This study explored the feasibility and performance of radiance field-based methods for underwater 3D reconstruction using real-world archaeological data. The results indicate that such methods, particularly SeaThru-NeRF and 3D Gaussian Splatting, can serve as valuable tools for visualising and documenting underwater cultural heritage. SeaThru-NeRF showed the highest geometric fidelity, thanks to its physically-based corrections tailored for underwater environments, while 3DGS offered visually compelling outputs with exceptional rendering speed.

While these techniques are not yet ready to replace conventional photogrammetric workflows in terms of metric accuracy and completeness, they offer a promising complement, especially in scenarios where rapid processing, low logistics, or limited image acquisition are key constraints. They also open new avenues for real-time visualisation and interpretation, which could greatly benefit both researchers and the wider public.

Looking forward, future perspectives regard improvements in data acquisition strategies, such as optimized lighting configurations could enhance input quality and reduce artefacts. Second, further development of reconstruction algorithms is needed to better manage optical complexities specific to underwater environments, including turbidity, caustics, and multi-path lighting effects.

Another promising direction involves the integration of radiance fields into standardised documentation pipelines, particularly through the fusion of outputs with dense SfM models or multibeam echosounders (MBES), creating unified datasets for archaeological interpretation. Validation on a wider range of submerged sites, varying in depth, typology, and environmental conditions, will be crucial to assess the robustness and scalability of these methods.

Acknowledgements

This study was carried out within the FAIR - Future Artificial Intelligence Research and received funding from the European Union Next-GenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.3 – D.D. 1555 11/10/2022, PE00000013). This manuscript reflects only the authors' views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

This publication is part of the project PNRR-NGEU which has received funding from the MUR – DM 351/2022.

Furthermore, this paper was published within the context of the Digital Twins for Cultural Heritage (Twin4CH) Junior Professorship Chair, financed by the French National Agency for Research (*Agence nationale de la recherche* – ANR).

Authors would like to thank SIFET, Fabio Menna, Loren Clark and Valeria Longhi for acquiring and kindly providing the data for this study.

References

Clini, P., Nespeca, R., Angeloni, R., Coppetta, L., 2024. 3D representation of Architectural Heritage: a comparative analysis of NeRF, Gaussian Splatting, and SfM-MVS reconstructions using low-cost sensors. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-2/W8-2024, 93-99.

Croce, V., Caroti, G., Luca, L. D., Piemonte, A., 2023. Neural radiance fields (nerf): review and potential applications to digital cultural heritage. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII, 25–30.

Drap, P., Seinturier, J., Hijazi, B., Merad, D., Boi, J.-M., Chemisky, B., Seguin, E., Long, L. (2015). The ROV 3D Project: Deep-sea underwater survey using photogrammetry: Applications for underwater archaeology. *Journal on Computing and Cultural Heritage (JOCCH)*, 8 (4), 1-24.

Haitz, D., Hermann, M., Roth, A. S., Weinmann, M., and Weinmann, M.: The Potential of Neural Radiance Fields and 3D Gaussian Splatting for 3D Reconstruction from Aerial Imagery, *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.*, X-2-2024, 97–104.

Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42.

Levy, D., Peleg, A., Pearl, N., Rosenbaum, D., Akkaynak, D., Korman, S., & Treibitz, T. (2023). *SeaThru-NeRF: Neural Radiance Fields in Scattering Media* (arXiv:2304.07743). arXiv. <https://doi.org/10.48550/arXiv.2304.07743>

Li, H., Song, W., Xu, T., Elsig, A. and Kulhanek, J., 2024. Watersplating: Fast underwater 3d scene reconstruction using gaussian splatting. arXiv preprint arXiv:2408.08206.

Menna, F., Agrafiotis, P., & Georgopoulos, A. (2018). State of the art and applications in archaeological underwater 3D recording and mapping. *Journal of Cultural Heritage*, 33, 231-248.

Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R. and Ng, R., 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1), pp.99-106.

Sarlin, P.E., DeTone, D., Malisiewicz, T. and Rabinovich, A., 2020. Superglue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 4938-4947).

Tancik, M., Weber, E., Ng, E., Li, R., Yi, B., Kerr, J., Wang, T., Kristoffersen, A., Austin, J., Salahi, K., Ahuja, A., McAllister, D., & Kanazawa, A. (2023). Nerfstudio: A Modular Framework for Neural Radiance Field Development. *Special Interest Group on Computer Graphics and Interactive Techniques Conference Proceedings*, 1–12. <https://doi.org/10.1145/3588432.3591516>

Yang, D., Leonard, J.J. and Girdhar, Y., 2024. Seasplat: Representing underwater scenes with 3D gaussian splatting and a physically grounded image formation model. arXiv preprint arXiv:2409.17345.

Zhang, T. and Johnson-Roberson, M., 2023. Beyond NeRF underwater: Learning neural reflectance fields for true color correction of marine imagery. *IEEE Robotics and Automation Letters*, 8(10), pp.6467-6474.