3D Reconstruction of Underwater Shipwrecks: 3D Gaussian Splatting and Structure from Motion for the Melania shipwreck

Dario Billi [1], Valeria Croce [2], Andrea Piemonte [1], Gabriella Caroti [1]

[1] Department of Civil and Industrial Engineering, ASTRO Laboratory, University of Pisa, L.go Lucio Lazzarino, 56122 Pisa, Italy; dario.billi@phd.unipi.it; andrea.piemonte@ing.unipi.it, gabriella.caroti@ing.unipi.it
[2] LISPEN EA7515, Arts et Métiers Institute of Technology, 2 Cours des Arts et Métiers, 13617 Aix-en-Provence, France; valeria.croce@ensam.eu

Keywords: Underwater, Photogrammetry, 3D-Gaussian Splatting, Melania shipwreck, Artificial Intelligence, Agisoft Metashape

Abstract

Underwater cultural heritage documentation presents significant challenges due to optical distortions inherent to aquatic environments. This study presents a dual approach for 3D reconstruction of underwater environments using two image-based 3D reconstruction techniques: Structure from Motion (SfM) and 3D Gaussian Splatting (3DGS). The goal is to evaluate the respective efficacy of each in terms of geometric accuracy and texture fidelity. The comparison is performed on the pilot case study of the Melania shipwreck, off the coast of Vada in Livorno, Italy. The results demonstrate that, while SfM proves more effective in structural reconstruction of submerged remains, 3DGS might enhance texture photo-realism and more effectively mitigate underwater lighting distortions. Conducted within the framework of the *Sub-IA* project, this research contributes to the development of integrated methodologies for the high-fidelity and realistic, image-based digital documentation of underwater heritage.

1. Introduction

Documenting submerged archaeological sites presents significant challenges due to the inherent conditions of underwater environments. Optical distortions, such as spectrally dependent attenuation and backscatter, severely impair the performance of conventional computer vision and photogrammetric techniques.

Recent advances in neural rendering offer promising alternatives (Croce et al., 2023) to Structure-from-Motion (SfM) for real-time 3D reconstruction. Among these, 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) allows rendering 3D scenes by means of anisotropic Gaussians, enabling efficient and photorealistic rendering even under complex lighting and environmental conditions.

This work compares SfM and 3DGS through a real-world case study: the Melania shipwreck, located in the Tuscan Archipelago. The research explores how these technologies can be integrated to generate high-fidelity, visually realistic 3D models suitable for both structural analysis and public dissemination.

2. Related work

2.1 Underwater 3D Reconstruction with Gaussian Splatting

Recent studies have successfully adapted the 3D Gaussian Splatting (3DGS) technique, originally developed for terrestrial environments, to underwater settings: WaterSplatting (Li et al., 2024) introduced a pipeline that combines the explicit geometric representation of 3DGS with volumetric rendering designed to simulate water-induced scattering. By incorporating a separate volumetric field to model the optical properties of water, this method effectively reduces visual distortion and improves rendering quality when compared to NeRF-based approaches (Mildenhall et al., 2020). Building on a similar concept, SeaSplat (Yang et al., 2024) enhanced underwater 3D representation by introducing a physically grounded image formation model that realistically accounts for light absorption and scattering. This enables real-time, photorealistic rendering and allows for accurate color correction of submerged objects, effectively overcoming the typical blue-green cast seen in underwater imagery.

RecGS (Zhang et al., 2024) used a recurrent neural framework to address the removal of dynamic light patterns, commonly known as water caustics, caused by surface refraction, to capture the temporal evolution of light artifacts, leading to more temporally stable and visually coherent reconstructions.

UW-GS (Wang et al., 2024) introduced a distance-aware color modeling strategy and dynamically adjusted the density of Gaussian points based on scene depth. Binary motion masks are also employed to accurately reconstruct dynamic scene in complex underwater environments.

In terms of performance, as well as focusing on the rendering speed, Gaussian Splashing (Mualem et al., 2024) leveraged the inherent efficiency of 3DGS to achieve reconstructions at up to 140 frames per second. It also demonstrated enhanced depth estimation capabilities, making it particularly suitable for real-time applications such as exploration or in-field operations.

For scenes involving motion or environmental change, UDR-GS (Du et al., 2024) proposed a depth-regularized optimization framework. The method used depth information that is not distorted by the water medium to obtain more geometrically consistent reconstructions of floating or moving objects.

Finally, Water-Adapted 3D Gaussian Splatting (Fan et al., 2025) refined the standard approach by introducing an adaptive point distribution that responds to local scene complexity. Combined with multi-scale rendering based on depth cues, this technique significantly improves both geometric accuracy and visual coherence.

2.2 Underwater Photogrammetry and Shipwreck Modelling

Underwater photogrammetry has emerged as a critical methodology for the non-invasive documentation and analysis of submerged archaeological sites and benthic ecosystems.

Recent advancements have addressed both technical challenges of underwater data acquisition and subsequent data processing for accurate 3D reconstructions. Alongside advanced neural rendering techniques, photogrammetry continues to play a fundamental role in the documentation of underwater environments, especially for archaeological shipwrecks. Recent research has emphasized the importance of meticulous image processing, both before and after 3D reconstruction.

The study titled Image Pre-Processing Strategies for Enhancing Photogrammetric 3D Reconstruction of Underwater Shipwreck Datasets (Calantropio et al., 2020) demonstrated how pre-processing steps, such as contrast enhancement, filtering, and lighting correction, can substantially improve the sharpness and overall quality of 3D models obtained in low-visibility environments. Complementing this research, the work by Adams et al. (2024) focused on post-processing strategies that enhance the interpretability of photogrammetric outputs. This is particularly relevant for underwater archaeology, where small structural features can hold significant scientific and historical value. Together, traditional and neural approaches offer a diverse and powerful toolkit for advancing the study and preservation of underwater cultural heritage.

Drap et al. (2015) presented a photogrammetric approach integrated with semantic knowledge representation for deep-sea surveys, exemplified by the case study of the Xlendi wreck off Malta. Their methodology involved real-time low-resolution 3D model generation from submarine-acquired images, followed by high-resolution modeling in laboratory conditions. Furthermore, the extraction and identification of artefacts are enhanced through the application of ontologies, specifically the CIDOC-CRM model, bridging logical frameworks with photogrammetric workflows.

The restoration and enhancement of underwater images, especially in complex lighting conditions, remain significant challenges due to the dynamic illumination setups in deep-sea environments. Song et al. (2023) proposed a novel method exploiting light field constraints within the camera frustum, capturing illumination and backscatter information volumetrically. Their system improves the accuracy of albedo restoration for multi-light setups, supporting consistent texturing in large-scale 3D underwater maps, for both archaeological documentation and marine biology studies.

Complementary to image restoration techniques, an extensive review by Vlachos and Skarlatos (2021) analyzed developments in underwater image color correction over the past fifteen years. Their survey underscores the evolution from early model-based corrections to modern AI-driven approaches, emphasizing the need for robust quality evaluation metrics.

In ecological monitoring contexts, Nocerino et al. (2020) demonstrated a rigorous underwater photogrammetry pipeline developed for the Moorea Island Digital Ecosystem Avatar project. By integrating underwater geodetic networks and rigorous camera calibration procedures, their work highlights the importance of minimizing geometric deformation and quantifying uncertainty propagation in multi-temporal 3D reef monitoring.

Doležal et al. (2019) further introduced a virtual reality (VR) educational tool designed to teach maritime archaeology students the fundamentals of underwater photogrammetry. The VR environment allows users to simulate marker placement, distance measurement, and photographic acquisition, promoting experiential learning and reducing errors during real-world deployments.

3. Materials and data collection

3.1 The Melania Shipwreck

The wreck of the cargo ship Melania, built in 1938 and sunk in February 1970 (Figure 1), lies at a depth of 7 to 12 meters off the coast of Vada, in Livorno, Italy. In March and April of 1970, the wreck was divided into three sections by the divers of the Italian Navy using explosive charges, due to the risk posed to navigation. The structure rests on a seabed largely covered by *Posidonia oceanica*, adding ecological and visual complexity to the survey.

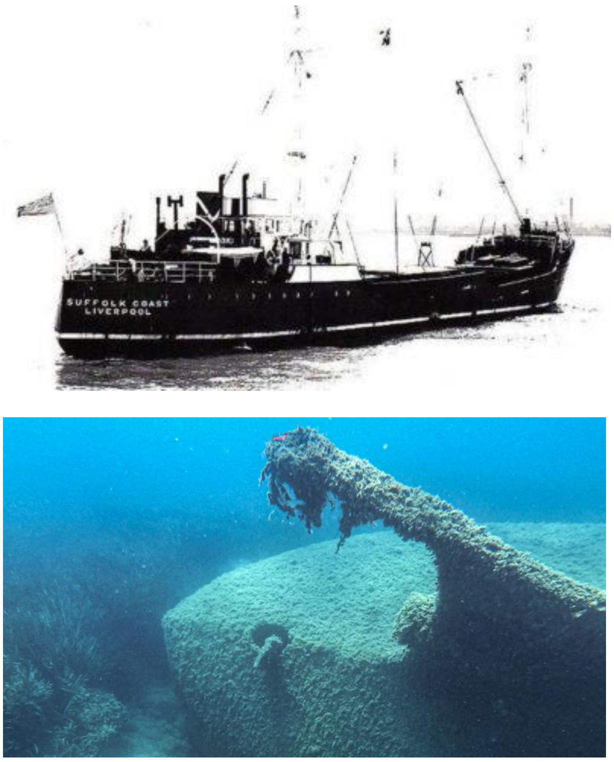

Figure 1. The MV Suffolk Coast, a British merchant ship, photographed during its service under the Liverpool flag. It was later renamed *Melania*.
Source: Daily Nautica – "Melania, the Forgotten Wreck" (above), and stern detail of the Melania wreck (below).

3.2 Survey Techniques

Underwater surveys were conducted using GoPro cameras mounted on divers. Multiple video sequences were collected and frames were extracted using FFmpeg. A total of 2649 images were exported from the video sequences, each with a resolution of 4000 × 3000 pixels at 300 dpi. A color checker was included in the scene to enable color correction in Adobe Lightroom. This step is crucial to mitigate the effects of spectral attenuation, particularly the loss of red tones at depth. Tests are conducted using an NVIDIA GeForce RTX 4090 GPU (24 GB VRAM) and an AMD Ryzen 97950X 16-Core CPU.

3.3 Software and environment employed:

- Agisoft Metashape v.2.0.0 (SfM software)
- Lightroom Classic v.14.0.1 (color correction)
- 3D Gaussian-Splatting (latest code update on Aug. 2024)
- Anaconda environment v.conda 23.7.4 (prompt for 3DGS)
- SuperSplat (latest code update on Apr. 2025)

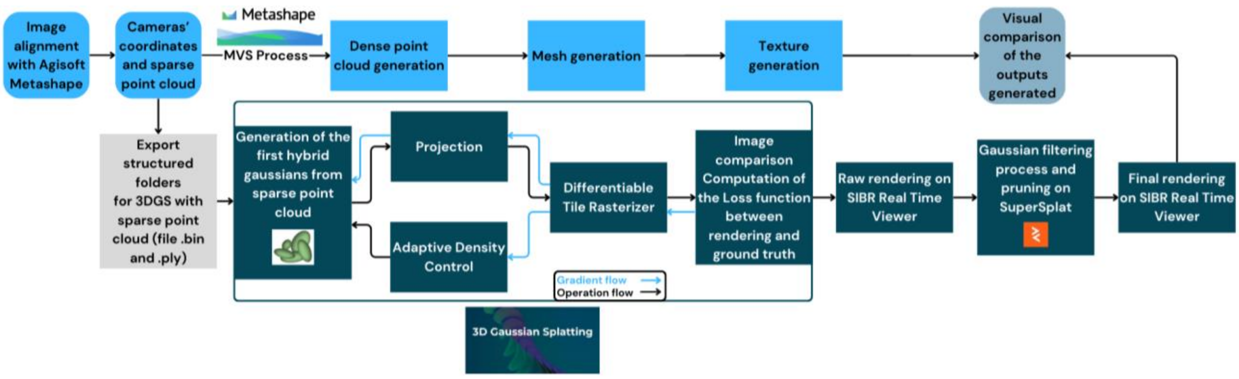

Figure 2. Overview of the proposed methodology

4. Methodology

The methodology (Figure 2) begins for both SfM and 3DGS with a shared input generated through the creation of a sparse point cloud from colour-corrected images. This initial step is fundamental, as it addresses the photogrammetric challenge by estimating both the internal and external parameters of the cameras used to capture the imagery. Typically, inside 3DGS this process is carried out using COLMAP, an open-source tool well-regarded for its performance in 3D reconstruction through Structure-from-Motion (SfM) techniques. However, in our case, COLMAP encountered difficulties in correctly aligning all images, resulting in incomplete reconstructions and inaccuracies in camera parameter estimation. To overcome these limitations, we opt for an executable script that processes input data from SfM analysis conducted with Agisoft Metashape. This approach provides two key benefits: it guarantees complete image alignment and produces a high-quality sparse point cloud, which serves as the common foundation for the two parallel processes. The common parameters for the alignment are shown in Table 1:

Alignment settings	
Accuracy	High
Limit key points	0
Limit tie points	0
Generic preselection	No
Reference preselection	No
Adaptive camera model fitting	No
Exclude stationary tie points	Yes
Guided image matching	No

Table 1. Common alignment parameters

On the one hand, the photogrammetric reconstruction is performed. The process is run via Agisoft Metashape and includes: dense point cloud generation, mesh construction, and texture mapping. The parameters for model generation are shown in Table 2.

On the other hand, for 3DGS, a set of anisotropic Gaussian ellipsoids is generated from the 3D points, with each ellipsoid modeled as a 3D Gaussian distribution. These ellipsoids are then projected onto 2D images from multiple viewpoints using the recovered camera poses. Within the camera frustum, the differentiable Gaussian functions are rendered into images through rasterization. A loss function is then computed by comparing the rendered images to the ground truth images, leading to adjustments of each Gaussian's parameters, including

Mesh settings	
Source data	Depth Maps
Surface type	Arbitrary
Quality	High
Face count	High
Interpolation	Enabled
Depth filtering	Mild

Table 2. Agisoft Metashape parameters to generate mesh model

position, size, and orientation. An adaptive density control method is also applied to optimize the properties of the Gaussian ellipsoids, refining their spatial distribution, scale, orientation, quantity, color, and opacity.

The 3DGS model was generated using a script executed via the Anaconda Prompt, with all parameters left at their default values except for the image resolution, which was reduced to 1/8 of the maximum value in order to shorten the processing time. This adjustment was necessary due to the large number and high quality of the input images. The training process was carried out for 30,000 iterations.

4.1 3DGS Rendering evaluation metrics

For novel view synthesis using 3DGS, visual quality is evaluated using widely adopted metrics (Table 3). The most common among these are Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) (Wang et al., 2004), and Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al., 2018).

Metric	Range	Interpolation
SSIM	> 0.90	Excellent structural similarity
	0.85 – 0.90	High quality
	0.80 – 0.85	Good quality
	<0.80	Noticeable structural degradation
PSNR	> 27	Very high visual fidelity
	23 – 27	High quality
	20 – 23	Medium/acceptable quality
	< 20	Perceptible degradation
LPIPS	< 0.10	Excellent perceptual similarity
	0.10 – 0.15	High perceptual quality
	0.15 – 0.20	Medium quality
	> 0.20	Low perceptual fidelity/perceptible error

Table 3. Metric quality thresholds for underwater scenes estimated using as reference, WaterSplatting (Li et al., 2024) and SeaSplat (Yang et al., 2024)

PSNR measures the similarity between rendered images and corresponding ground truth images, where higher PSNR values indicate better image fidelity. However, PSNR mainly focuses on mean squared error and does not account for human sensitivity to different frequency components or perceptual distortions. As a result, PSNR may not always accurately reflect subjective visual quality.

SSIM evaluates structural similarity by considering brightness, contrast, and structural information between two images. Its values range from −1 to 1, with scores closer to 1 indicating greater similarity.

LPIPS assesses image similarity based on feature representations extracted by deep neural networks pre-trained for image classification. The original LPIPS framework employed networks such as SqueezeNet (Iandola et al., 2016), VGG (Simonyan and Zisserman, 2015), and AlexNet (Krizhevsky et al., 2012) as feature extractors. Unlike traditional metrics such as PSNR and SSIM, which rely on pixel-based comparisons, LPIPS leverages deep learning to measure perceptual similarity more closely aligned with human visual judgments. Lower LPIPS scores correspond to higher perceived similarity between images.

By utilizing deep feature distances rather than raw pixel differences, LPIPS offers a more perceptually relevant assessment of image quality, making it particularly suitable for evaluating the outputs of neural rendering techniques like 3DGS. In the final phase of the workflow, the SuperSplat platform was employed to manually refine the 3D Gaussian

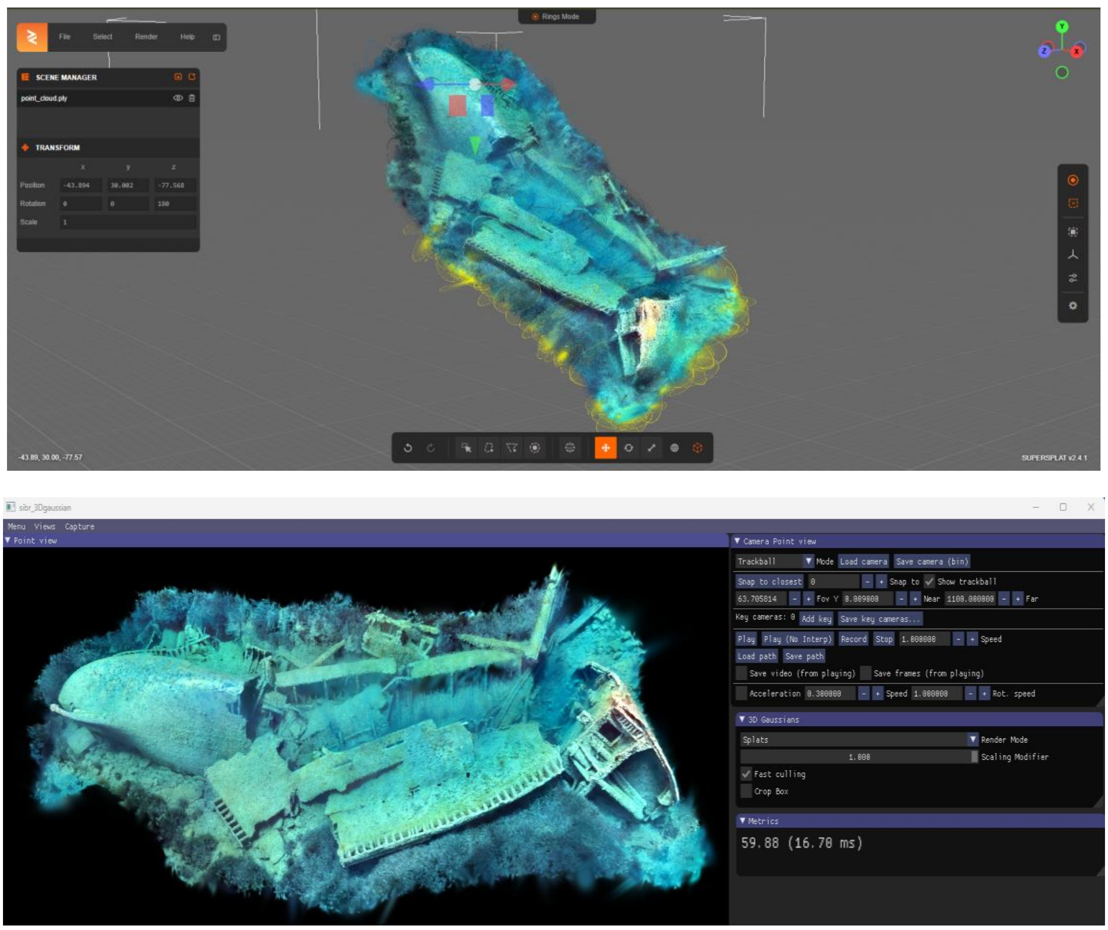

Figure 3. The manual ellipsoids filtering process implemented in SuperSplat (*above*), and SIBR viewer inside 3DGS environment (*below*).

Splat (3DGS) model. This open-source browser-based tool, developed on the PlayCanvas engine, enabled the efficient inspection and selective editing of individual splats, which was particularly useful for removing noise and artifacts such as misaligned or floating ellipsoids generated during the reconstruction process.

Thanks to its native support for the PLY format and its intuitive interface, SuperSplat allowed for precise selection, translation, and deletion of splats directly within the 3D environment without the need for local installations. This made it an ideal solution for the final cleanup and preparation of the dataset for visualization and dissemination.

Gaussian splat. As for 3DGS, rendering is executed through SIBR viewer and post-processed using SuperSplats for final manual cleanup (Figure 3).

The comparison of Figure 4 reveals a clear trade-off between geometric fidelity and textural realism. In terms of geometric accuracy, SfM provides a detailed and reliable structural representation, making it particularly effective for archaeological measurements and modeling tasks. The resulting mesh accurately captures the shape and dimensions of the wreck, preserving the essential geometric characteristics of the scene.

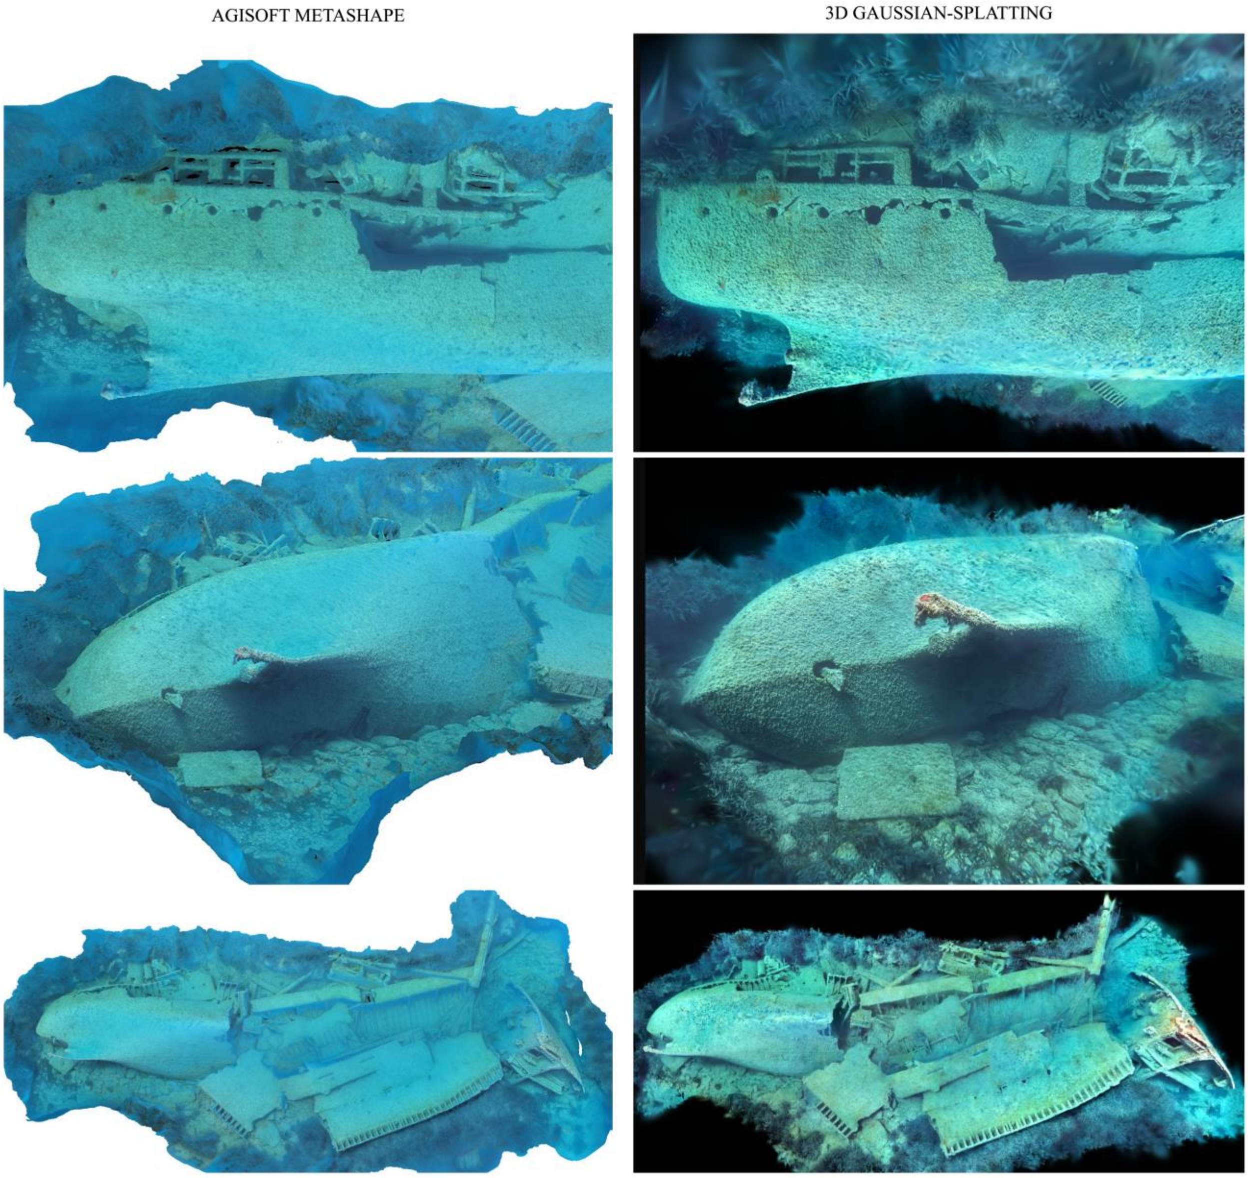

Figure 4. Visual comparison between Agisoft Metashape and 3D Gaussian-Splatting models.

5. Results and discussion

Figure 4 displays a visual comparison between the textured mesh obtained via Agisoft Metashape and the corresponding

Conversely, with respect to textural realism, 3DGS demonstrates a superior capacity to mitigate underwater lighting distortions. The reconstructed colors appear more vivid and closer to real-world perception, effectively addressing the

https://doi.org/10.5194/isprs-archives-XLVIII-2-W10-2025-19-2025 | © Author(s) 2025. CC BY 4.0 License.

typical limitations encountered in SfM-based texture rendering under aquatic conditions.
While the SfM model retains high geometric completeness, it exhibits muted colors and a general loss of contrast. In contrast, the 3DGS model produces visually more vibrant and perceptually engaging imagery, particularly in areas most affected by depth-induced light attenuation.

The rendering metrics computed for the 3DGS model (Table 4) report an SSIM of 0.79, a PSNR of 25.24, and an LPIPS score of 0.29. According to the quality thresholds outlined in Table 3, these values indicate a medium level of structural similarity and visual fidelity, with low perceptual quality.

Specifically, the SSIM value of 0.79 falls into the lower range, signaling noticeable structural degradation, though not severe. The PSNR score of 25.24 falls within the high-quality bracket, indicating good visual fidelity in terms of signal-to-noise ratio and supporting the preservation of photometric and geometric details. However, the LPIPS score of 0.29 exceeds the 0.20 threshold, reflecting low perceptual similarity and a more pronounced visual discrepancy when compared to ground-truth references, based on learned perceptual metrics.

METRICS	
WITHOUT SUPERSPLAT EDITING	
SSIM	0.79
PSNR	25.24
LPIPS	0.29

Table 4. Rendering metrics values

These results show that, while the 3DGS model does not achieve top-tier structural or perceptual scores, it succeeds in maintaining a solid trade-off between visual quality and accuracy, especially in challenging underwater scenarios. The model's ability to achieve a high PSNR is particularly relevant for underwater archaeological imaging, where environmental conditions often compromise image integrity. In this context, although perceptual similarity remains limited, the level of structural and photometric detail achieved represents a significant improvement over previous methods, supporting both the visual interpretability and the scientific utility of the generated documentation.

However, while the current outcomes are encouraging, the model's suitability for high-precision spatial or metric applications should be further assessed. Future developments may focus on hybrid approaches that integrate the geometric precision of Structure-from-Motion (SfM) with the perceptual realism enabled by 3DGS. Such pipelines could offer enhanced tools for rigorous documentation and more accessible visual communication in the study and preservation of underwater cultural heritage.

6. Conclusions

This study investigated the effectiveness of two distinct approaches to 3D reconstruction in underwater environments: traditional photogrammetry based on SfM and the more recent neural rendering method known as 3DGS. Using the Melania shipwreck as a case study, we evaluated the strengths and limitations of each technique, both in terms of geometric accuracy and visual fidelity.

The SfM pipeline, implemented using Agisoft Metashape, proved to be robust in generating detailed and metrically accurate meshes. However, its final output quality is highly dependent on image visibility conditions, which are often compromised underwater due to turbidity, refraction, and selective light absorption.

Conversely, the 3DGS method, tested via the SuperSplat implementation, yielded impressive results in terms of photorealistic rendering and photometric consistency, while also significantly reducing processing times compared to traditional workflows. Nevertheless, this approach relies on an implicit representation, lacking an explicit triangular mesh, which limits its applicability in metric analyses or detailed geometric evaluations.

In conclusion, the findings suggest that:

- Photogrammetry remains the most suitable tool for applications requiring metric precision and models that can be exported to CAD/GIS platforms.
- 3D Gaussian Splatting offers a valuable alternative for immersive visualization, rapid documentation, and digital museum dissemination, with enhanced visual fidelity.
- The future integration of both methods could lead to an ideal hybrid solution, combining the geometric precision of SfM with the photorealistic rendering capabilities of 3DGS.

Further research should explore the optimization of preprocessing techniques for underwater imagery (e.g., dehazing, color correction) and the extension of the pipeline to more complex or dynamic scenarios.

7. Acknowledgements

The authors would like to express their sincere gratitude to all the individuals and institutions involved in the Sub-IA project, including Simonetta Menchelli, Domingo Belcari and Silvia Marini from the Department of Civilisations and Forms of Knowledge, as well as Ferruccio Maltagliati and Jonathan Tempesti from the Department of Biology of the University of Pisa. We also gratefully acknowledge the support of the industrial sponsors: INEOS, MRB Energia, Divesystem, Ratio Computer and UTRTEK-DIR Store.

8. References

Adams, D., Helmholz, P., McMeekin, D., & Woods, A., 2024. Improving Underwater Photogrammetric 3D Reconstruction Processing of Shipwreck Sites. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, X-4-2024, 7–14. DOI: 10.5194/isprs-annals-X-4-2024-7-2024

Calantropio, A., Chiabrando, F., Seymour, B., Kovacs, E., Lo, E., Rissolo, D., 2020. Image Pre-Processing Strategies for Enhancing Photogrammetric 3D Reconstruction of Underwater Shipwreck Datasets. ISPRS Arch. Photogramm. Remote Sens. Spatial Inf. Sci., XLIII-B2-2020, 941–948. https://doi.org/10.5194/isprs-archives-XLIII-B2-2020-941-2020.

Croce, V., Caroti, G., De Luca, L., Piemonte, A., Véron, P., 2023. Neural Radiance Fields (NeRF): Review and Potential Applications to Digital Cultural Heritage. ISPRS Arch.

Photogramm. Remote Sens. Spatial Inf. Sci., XLVIII-M-2-2023, 453–460. https://doi.org/10.5194/isprs-archives-XLVIII-M-2-2023-453-2023.

Doležal, M., Vlachos, M., Secci, M., Demesticha, S., Skarlatos, D., & Liarokapis, F., 2019. Understanding Underwater Photogrammetry for Maritime Archaeology Through Immersive Virtual Reality. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-2/W10, 85–91. DOI: 10.5194/isprs-archives-XLII-2-W10-85-2019

Drap, P., Merad, D., Mahiddine, A., Boi, J.M., Chemisky, B., Long, L., Seinturier, J., 2015. Underwater Photogrammetry and Object Modeling: A Case Study of Xlendi Wreck in Malta. Sensors, 15(12), 30351–30384. https://doi.org/10.3390/s151229802.

Du, Y., Zhang, Z., Zhang, P., Sun, F., Lv, X., 2024. UDR-GS: Enhancing Underwater Dynamic Scene Reconstruction with Depth Regularization. Symmetry, 16(8), 1010. https://doi.org/10.3390/sym16081010.

Fan, X., Wang, X., Ni, H., Yuanxue, X., & Shi, P., 2025. Water-Adapted 3D Gaussian Splatting for Precise Underwater Scene Reconstruction. *Remote Sensing*.

Iandola, F.N., Han, S., Moskewicz, M.W., Ashraf, K., Dally, W.J., Keutzer, K., 2016. SqueezeNet: AlexNet-Level Accuracy with 50x Fewer Parameters and <0.5MB Model Size. arXiv preprint arXiv:1602.07360.

Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. ACM Trans. Graph., 42(4), 1–14. https://doi.org/10.1145/3592433.

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. ImageNet Classification with Deep Convolutional Neural Networks. Adv. Neural Inf. Process. Syst. (NIPS 2012), 25, 1097–1105.

Li, H., Song, W., Xu, T., Elsig, A., Kulhanek, J., 2024. WaterSplatting: Fast Underwater 3D Scene Reconstruction Using Gaussian Splatting. arXiv:2408.08206. https://arxiv.org/abs/2408.08206.

Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R., 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. ECCV 2020. https://arxiv.org/abs/2003.08934.

Mualem, N., Amoyal, R., Freifeld, O., Akkaynak, D., 2024. Gaussian Splashing: Direct Volumetric Rendering Underwater. arXiv:2411.19588. https://arxiv.org/abs/2411.19588.

Nocerino, E., Menna, F., Remondino, F., Toschi, I., & Fassi, F., 2020. Coral Reef Monitoring by Scuba Divers Using Underwater Photogrammetry and Geodetic Surveying. *Remote Sensing*, 12(18), 3036. DOI: 10.3390/rs12183036

Simonyan, K., Zisserman, A., 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv preprint arXiv:1409.1556.

Song, Y., She, M., & Köser, K., 2023. Advanced Underwater Image Restoration in Complex Illumination Conditions. *arXiv preprint arXiv:2309.02217*. https://arxiv.org/abs/2309.02217

Vlachos, M., & Skarlatos, D., 2021. An Extensive Literature Review on Underwater Image Colour Correction. *Sensors*, 21(17), 5690. DOI: 10.3390/s21175690

Wang, H., Anantrasirichai, N., Zhang, F., Bull, D., 2024. UW-GS: Distractor-Aware 3D Gaussian Splatting for Enhanced Underwater Scene Reconstruction. arXiv:2410.01517. https://arxiv.org/abs/2410.01517.

Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., 2004. Image Quality Assessment: From Error Visibility to Structural Similarity. IEEE Trans. Image Process., 13(4), 600–612.

Yang, D., Leonard, J.J., Girdhar, Y., 2024. SeaSplat: Representing Underwater Scenes with 3D Gaussian Splatting and a Physically Grounded Image Formation Model. arXiv:2409.17345. https://arxiv.org/abs/2409.17345.

Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O., 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Salt Lake City, UT, USA, 586–595.

Zhang, T., Zhi, W., Huang, K., Mangelson, J., Barbalata, C., Johnson-Roberson, M., 2024. RecGS: Removing Water Caustic with Recurrent Gaussian Splatting. arXiv:2412.03435.