Analysis of refraction aware Neural Radiance Fields for 3D reconstruction of through the water scenes

Markus Brezovsky[1], Anatol Günthner[2], Frederik Schulte[3], Lukas Winiwarter[3], Boris Jutzi[2], Gottfried Mandlburger[1]

[1] Department of Geodesy and Geoinformation, TU Wien, Vienna, Austria - (markus.brezovsky, gottfried.mandlburger)@tuwien.ac.at
[2] Institute of Photogrammetry and Remote Sensing, Karlsruhe Institute of Technology, Karlsruhe, Germany - (anatol.guenthner, boris.jutzi)@kit.edu
[3] Unit of Geometry and Surveying, Faculty of Engineering Sciences, University of Innsbruck, Innsbruck, Austria – (frederik.schulte, lukas.winiwarter)@uibk.ac.at

Keywords: NeRFrac, Underwater Photogrammetry, Neural Radiance Fields, Bathymetry, UAV Imaging, Refraction Correction.

Abstract

Neural Radiance Fields (NeRFs) synthesize novel views based on images acquired from different camera positions to represent 3D scenes. However, since they assume linear light paths, they are unsuitable for underwater environments where refraction causes nonlinear ray trajectories, resulting in blurred scene reconstructions due to the absence of physical light path modeling. Developments, such as NeRFrac, try to explicitly model refractive surfaces by incorporating Snell's law into NeRF frameworks. Nevertheless, the predominant objective within the computer vision community, the generation of high-quality renderings, assessed through metrics, e.g. including the Peak Signal-to-Noise Ratio (PSNR) persists. However, the main task in photogrammetric and geospatial applications is geometric reconstruction in the form of 3D point clouds. Therefore, this work investigates the possibilities of extracting 3D point clouds from refraction-aware NeRF implementations, specifically evaluating the NeRFrac codebase with the use of nine images.

1. Introduction

Geometric reconstruction of underwater environments is essential for various applications, including bathymetric surveying, ecological monitoring, underwater archaeology, and infrastructure inspection. Traditional photogrammetric approaches often struggle in such scenarios due to the complex interaction of light with water, particularly refraction at the water-air interface. This phenomenon causes non-linearities in the paths of light rays, making standard linear photogrammetric models inadequate for geometric reconstructions.

In recent years, Neural Radiance Fields (NeRFs) Mildenhall et al. (2020) have revolutionized the field of 3D reconstruction by synthesizing novel views using neural representations. However, standard NeRF implementations inherently assume linear ray trajectories, rendering them unsuitable for underwater scenes, where refraction at interfaces causes notable distortions and blurring in reconstructed images and geometries.

To overcome these limitations, recent advances have introduced refraction-aware NeRF frameworks such as NeRFrac Zhan et al. (2023), which explicitly incorporate physical models like Snell's law to compute accurate refractive ray trajectories. While these models try to enhance rendering quality, the primary evaluation metrics used within the computer vision community, such as the Peak Signal-to-Noise Ratio (PSNR) and the Structural Similarity Index (SSIM), focus predominantly on perceptual and radiometric rather than geometric accuracy. In contrast, disciplines like photogrammetry and remote sensing place a higher value on geometric outputs, such as geometrically correct and dense 3D point clouds.

Our current research aims to bridge the gap between visually appealing renderings and geometrically correct 3D reconstructions from NeRF scenes. Specifically, this paper investigates the extraction process of 3D point clouds from refraction-aware NeRF implementations, using NeRFrac as a case study. Our developments target the challenges caused by underwater photogrammetric conditions, making substantial progress towards practical applications in bathymetric mapping.

1.1 BathyNeRF Research Project

This research was carried out within the scope of the *BathyNeRF* project, a joint transnational research initiative that focuses on the development of neural NeRF-based methods for reconstructing underwater topography and submerged vegetation from aerial UAV imagery. The central objective of the project is to extend existing NeRF algorithms towards accurate refraction-aware 3D reconstruction in shallow aquatic environments, addressing challenges related to multimedia light propagation and complex water-surface dynamics.

In this context, the refractive NeRFrac framework Zhan et al. (2023) was adopted and extended to better handle refractive effects specific to UAV-based bathymetric photogrammetry.

In the first experimental studies, Guenthner et al. (2025) shows the general applicability of NeRFrac for UAV-based bathymetric reconstruction under field conditions. Our research focuses on evaluating the 3D reconstruction of submerged topography using masked water regions and refractive ray modeling, based on UAV imagery collected at the Pielach River test site (Mandlburger et al., 2025b). The study illustrated that explicit refraction modeling significantly improves the geometric consistency of the reconstructed underwater scene when compared to classical non-refractive approaches.

In addition to these experiments, Schulte et al. (2025) developed a simulation-based validation framework to systematically investigate the performance of geometric reconstruction under controlled conditions. Using synthetically generated underwater scenes, they analyzed the sensitivity of a standard NeRF model to different acquisition geometries, water depths, and refractive indices. This simulation approach allowed a targeted evaluation of reconstruction errors and highlighted the benefits of physically-based refraction modeling for stable geometry recovery, particularly in shallow water conditions.

Both studies demonstrate the fundamental potential of refraction-aware NeRF frameworks for bathymetric mapping and provide valuable insights into the capabilities and current limitations of the NeRFrac approach under real-world and simulated conditions.

In this paper, we complement this by enabling the export of the implicit geometry within a refractive NeRF as a 3D point cloud. This opens pathways towards downstream analyses that rely on this data representation, such as biomass estimation, water surface and water bed modeling, and the reconstruction of submerged objects. Furthermore, we investigate a training strategy that limits training data solely to submerged areas, ensuring a consistent refractive field can be learned, and the reconstruction of non-submerged areas does not hinder the modeling of refractive effects.

2. Related Work

Optical 3D underwater mapping has traditionally relied on active and passive remote sensing methods (Mandlburger, 2022). Active approaches such as airborne LiDAR bathymetry (ALB) enable accurate 3D point cloud generation, even for semi-transparent underwater surfaces, thanks to multi-return capabilities of laser systems Shan and Toth (2018). Passive methods like Structure-from-Motion (SfM) and Dense Image Matching (DIM) depend on image orientation and dense correspondences but are challenged by refractive distortions at the air-water interface Förstner and Wrobel (2016); Mandlburger (2019).

In recent years, Deep Learning (DL) techniques, and particularly Neural Radiance Fields (NeRFs), have emerged as promising methods for reconstructing complex 3D scenes Mildenhall et al. (2020). Unlike traditional active (e.g., LiDAR) or passive (e.g., photogrammetric) remote sensing methods, NeRFs do not directly measure depth or disparity. Instead, they rely on passive multi-view imagery and represent scenes as continuous volumetric functions that map spatial position and viewing direction to color and opacity. By sampling points along camera rays and applying volumetric rendering, NeRFs synthesize novel photorealistic views with high consistency across perspectives. These models inherently build on the passive sensing principle of image-based reconstruction but leverage learned priors and implicit scene representations to infer geometry and appearance. Their effectiveness, however, still depends on well-registered multi-view image data sets, typically acquired by passive sensor systems such as UAV-borne cameras. These acquisition platforms are often used in large-scale or high-resolution surveys, where the number of captured images and their variability in geometry and appearance present significant computational and modeling challenges for NeRF frameworks.

To cope with these demands, the original NeRF approach has been extended to improve efficiency and robustness. Mip-NeRF Barron et al. (2021) introduced a multiscale representation based on conical frustums, allowing for better handling of scale variations, improved anti-aliasing, and faster training. Mip-NeRF merges coarse and fine sampling stages into a single model, significantly improving efficiency while preserving geometric detail. NeRF-W Martin-Brualla et al. (2021) extends NeRF to handle uncontrolled image collections with varying lighting, exposure, and transient occlusions. This is achieved by augmenting the model with additional latent variables that explain per-image appearance variations. Although NeRF-W improves robustness under uncontrolled acquisition, geometric accuracy may still degrade in poorly observed regions or under unfavorable camera pose distributions.

Beyond these, several specialized NeRF variants aim to address scaling and scene complexity. BungeeNeRF Xiangli et al. (2022) combines multi-scale imagery, such as satellite and UAV data, using a hierarchical network structure to enable city-scale scene rendering. Despite promising qualitative results, such large-scale variants often lack explicit mechanisms to control geometric accuracy, particularly in highly refractive or underwater environments. In the context of underwater photogrammetry, most NeRF models implicitly assume simple scene geometries with near-spherical camera distributions, which are not achievable in UAV-based bathymetric surveys where primarily straight flight lines with nadir and oblique views are available. These limitations are further amplified when imaging through air-water interfaces, where light rays are refracted, leading to nonlinear ray paths that are not modeled in conventional NeRF frameworks.

Recent work on modeling refraction in NeRFs has made progress, but also revealed key limitations. LB-NeRF Fujitomi et al. (2022) uses deformable neural fields to approximate refractive effects but lacks a physically grounded refraction model. Ref-NeRF Verbin et al. (2022) improves appearance modeling for reflective surfaces but does not explicitly address refraction. Ref²-NeRF Kim et al. (2024) incorporates secondary reflections but only models refractive effects implicitly, limiting its geometric accuracy in complex underwater scenarios.

In this context, NeRFrac **?** represents a physically motivated advancement by explicitly integrating Snell's law into the NeRF rendering pipeline. By modeling refractive ray bending at air-water interfaces, NeRFrac enables physically correct ray paths and improves geometric reconstruction in submerged environments. This physics-based approach is essential for achieving accurate 3D mapping in bathymetric applications where refraction plays a dominant role.

3. Methods

In the following, we make use of a refraction-aware NeRF and investigate the applicability of physically-based refraction modeling for UAV-based bathymetric 3D reconstruction and geometry extraction. The approach extends the standard NeRF framework by explicitly accounting for refraction of light rays at the air-water interface using Snell's law. To this end, NeRFrac introduces a two-stage ray model, where primary rays are first traced to an estimated water surface, refracted, and then enter the underwater scene with an updated direction vector.

3.1 Refractive Modeling

The refraction-aware NeRF pipeline involves a two-stage process to model refracted rays. Initially, viewing rays are traced from the camera to a virtual refractive surface. This surface is determined by a refractive field. The offset network then computes water depth corrections to these intersection points, effectively relocating the initial sampling positions of the rays to reflect physical refraction phenomena.

Subsequently, secondary rays originate from these corrected positions along directions determined by Snell's law, guided by estimated normals of the refractive surface. These refracted rays sample the geometry of the underwater scene.

3.2 Masking of water area

To mitigate the ambiguities introduced by refraction effects and to enhance the network's capability to reconstruct underwater geometry, we implemented a targeted training strategy focused on water-covered areas. For this purpose, binary masks were created to delineate regions containing water from other irrelevant parts of the scene.

The masks were generated manually using QGIS 3.40 QGIS Development Team (2025). Polygons were drawn that approximate the water-covered regions within the scene. These binary masks (with values 0 and 1) were then rasterized and downsampled to the same resolution as the training images to match the input data dimensions.

During training, these masks were used for ray selection. Only rays that intersect water regions were selected for the loss calculation and hence optimization. This selective sampling ensures that the NeRFrac model primarily learns refractive geometries and ignores areas without water (Figure 1), which do not exhibit significant refractive effects. The masking approach thus increases the stability of the model and may allow for more focused optimization of the refracted geometry below the water surface.

Figure 1. RGB image with water mask (overlain in white transparent).

3.3 3D reconstruction

To extract explicit 3D geometric information from the trained NeRFrac model, a customized point cloud export procedure was implemented, as illustrated in Figure 2. The goal of this procedure is to convert the implicit volumetric scene representation, stored in the neural network weights, into discrete 3D point sets while respecting the physically-based refractive ray paths estimated during training.

The point cloud generation consists of the following steps, which are schematically summarized in Figure 2:

3.3.1 Multi-view ray sampling: For point cloud generation, dense sets of rays are sampled across multiple camera poses to ensure broad coverage of the reconstructed scene. For each camera view, a quasi-uniform 2D grid of sampling locations is defined across the image plane. At each selected grid location, rays are cast through the scene according to the camera's intrinsic and extrinsic parameters.

To improve the stability of local surface normal estimation, which is required for correct refraction modeling, not only single rays are considered, but small local neighborhoods of rays are sampled simultaneously. Concretely, 3×3 patches of adjacent pixels are used, such that each central ray is accompanied by eight neighboring rays. This local redundancy allows for the estimation of *local surface orientation* in subsequent processing stages.

3.3.2 Refractive ray calculation: Each sampled ray is processed by the trained NeRFrac model. The model first predicts the refractive surface intersection point, i.e., the location where the incoming ray intersects the water surface, corrected by the learned surface offset field. Based on the estimated surface normal at this intersection point, the refracted ray direction is computed according to Snell's law. This physically-based calculation accounts for the bending of light as it transitions from air into water, producing refracted rays that subsequently penetrate into the submerged scene along physically plausible paths. The accuracy of this refraction modeling critically depends on the quality of the predicted surface geometry and, in particular, the reliability of the estimated surface normals. Inaccuracies in surface orientation can lead to erroneous refraction angles, and thus degrade the geometric consistency of the reconstructed underwater scene.

3.3.3 Volume sampling and weighting: Along each refracted ray, 3D samples are distributed to probe the learned volumetric scene representation. These samples correspond to discrete 3D positions along the refracted ray path where the neural network is queried for *density* (opacity) and *color*. The sampling is performed in a hierarchical two-stage fashion: first, a coarse stratified sampling distributes a fixed number of samples evenly along the ray within a predefined near-far range. In a second refinement stage, importance sampling is applied to increase the sample density in regions of higher predicted opacity, focusing computational effort on areas likely to contain objects.

Each sampled location yields an opacity value (alpha), which indicates the likelihood of a surface to exist at that position along the ray. Additionally, for every sample, the transmittance, i.e., the accumulated opacity up to that point, is computed. Combining opacity and transmittance gives a rendering weight for each sample, which reflects its effective contribution to the final rendered pixel color during neural volume rendering.

3.3.4 Weight-based filtering: Since not all sampled 3D points contribute equally to the geometry of the scene, a filtering step is applied to retain only the most relevant samples. The rendering weights serve as a proxy for surface likelihood. Points with low rendering weight, i.e., located in transparent or empty regions, are discarded. Filtering can be performed either by applying an absolute threshold to the weight values or by retaining only a specified top-k% fraction of the highest-weighted samples. This selection ensures that primarily surface-relevant 3D points enter the exported point cloud.

3.3.5 Export and coordinate back-transformation: The retained 3D points are initially expressed in the normalized coordinate system used during NeRFrac training. In principle, these coordinates can be transformed back into the original world coordinate system by reversing the normalization operations applied during LLFF-style data preparation, including scene recentering, uniform scaling, and rotational alignment. However, the accuracy and consistency of this back-transformation remain uncertain at the current stage. Further work is required to fully validate and, if necessary, refine the transformation pipeline to ensure metric consistency and enable a reliable comparison with external photogrammetric reference data.

3.3.6 Current limitation: Since the NeRFrac model operates in a normalized scene space, the density, scale, and completeness of the extracted point cloud can be influenced by factors such as the number of training views, their spatial distribution, and the depth range covered during training. Additionally, the selection of sampling density, filtering thresholds, and batch sizes during extraction directly affect point cloud resolution and noise characteristics. As the export pipeline remains subject to ongoing development, no comprehensive parameter sensitivity analysis has yet been completed.

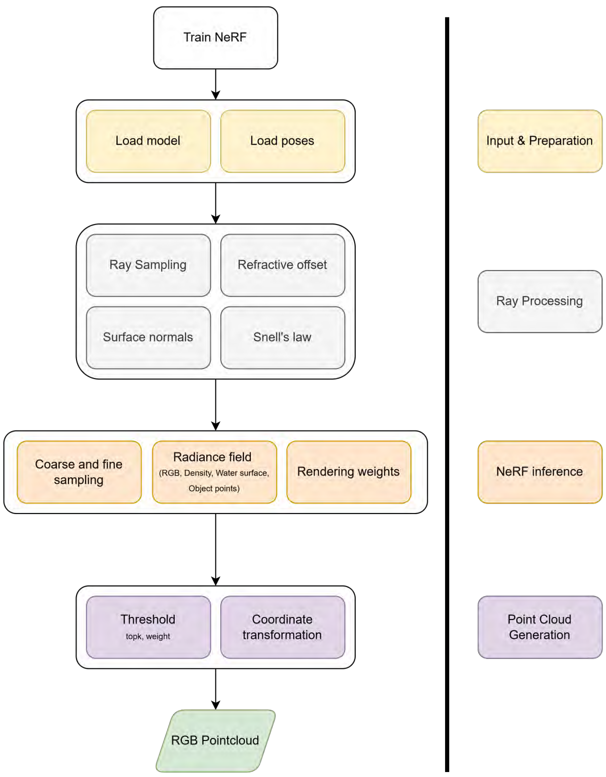

Figure 2. 3D reconstruction workflow

4. Experimental Setup

To evaluate the performance of the refractive NeRF approach, experiments were conducted on real UAV-based data sets acquired over the Pielach River at the Neubacher Au site in Lower Austria. Data acquisition was carried out in October 2024 under favorable environmental conditions, which ensured high water clarity and good visibility of the submerged topography.

The images were collected using a DJI Zenmuse P1 camera providing 45-megapixel resolution and a DJI M350 RTK multicopter. Flight planning was designed to capture a wide range of viewing geometries in the study area. In addition to nadir flights, multiple oblique flight lines were performed to introduce angular diversity into the data set. The acquisition combined near-nadir and oblique viewing directions, covering forward-, backward-, left- and right-oriented perspectives. This was achieved using the Smart Oblique Mode functionality of the UAV system in combination with additional manually planned flight patterns. The complete flight configuration is visualized in Figure 3, which shows the distribution of camera positions and view orientations across the site. Ground Control Points (GCPs), placed throughout the area and indicated by purple crosses in the figure, were used to geo-reference and assess the accuracy of the reconstruction. The course of the Pielach River, which constitutes the primary target of bathymetric reconstruction, is highlighted in light blue.

In addition to the UAV imagery, airborne laser bathymetry (ALB) reference data was simultaneously acquired over the same area. This data set provides an independent water depth reference for the Pielach River and allows quantitative evaluation of the NeRFrac-based bathymetric reconstructions. The corresponding ALB-derived water depth map is presented in Figure 4, which illustrates the variation in the topography of the riverbed within the area of interest.

As a second independent reference, a dense multi-view stereo (MVS) point cloud was generated from the full UAV data set using standard photogrammetric processing workflows i.e., not accounting for refraction effects at the water surface. This data set serves as an additional reference for evaluating the geometric accuracy of the NeRFrac reconstructions. A representative section of the generated MVS point cloud is shown in Figure 5, clearly visualizing both submerged and emergent structures along the riverbed and its adjacent gravel banks.

To systematically analyze the influence of image configuration on reconstruction quality, several data set subsets with varying image geometries were derived from the full acquisition. In total, four configurations were defined, each containing nine images: two configurations contain only nadir images with varying baseline distances, while the other two include only oblique images, again with differing baselines. These configurations allow an isolated evaluation of the effect of both viewing direction and acquisition geometry on NeRFrac's reconstruction capabilities. The variants of the resulting data set are visualized in Figure 6, which shows the spatial arrangement and the baseline distances for each subset.

The entire data set, including reference and derived data products, is based on the publicly available Pielach benchmark data set Mandlburger et al. (2025a,b).

5. Results and Discussion

The following sections present and discuss the results obtained from applying the refractive NeRF (NeRFrac) approach to UAV imagery of shallow water environments. Specifically, we focus on qualitative aspects of mask training, the visual quality of 2D renderings, and the geometric consistency of derived point clouds.

5.1 Mask-based Training and 2D Rendering

The incorporation of binary water masks during training proved essential to constrain the NeRFrac model's focus on water-

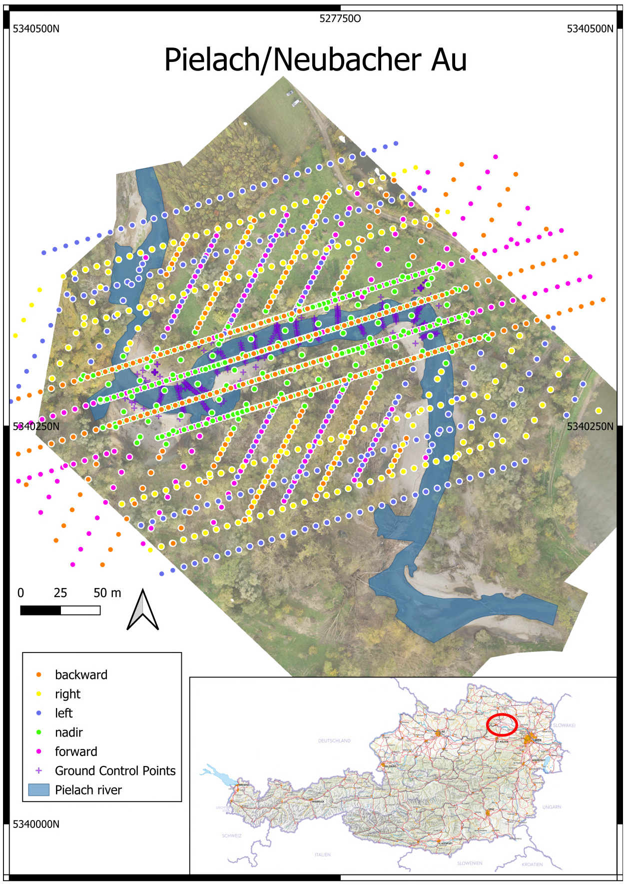

Figure 3. Flight configuration for the Pielach/Neubacher Au study area in Lower Austria. Colored symbols represent different viewing directions. Ground Control Points (GCPs) are shown as purple crosses. The light blue area indicates the river course of the Pielach. Used coordinate system: ETRS89/UTM33N (EPSG: 25833)

covered regions. By explicitly guiding the training process towards refractive regions, the masks allowed the model to better represent the refraction effects encountered at the air-water interface. Consequently, areas outside the mask were ignored, substantially reducing irrelevant data and stabilizing the training process.

The results rendered illustrate the impact of mask-based training. Figure 7 (middle) shows the binary water mask image, while Figure 7 (bottom) depicts the corresponding rendering outcome. The distinct brownish appearance observed in the rendered images is attributable to the model's exclusive focus on submerged, sediment-rich areas defined by the mask. This results in a relatively homogeneous reflectance pattern, which effectively highlights the underwater topography but lacks the visual diversity present in above-water regions.

Importantly, the NeRF model was not provided with any training information outside of the masked water area. As a result, regions beyond the Pielach river, such as vegetation or background structures, are only inferred by extrapolation. This explains the uniform or oversimplified appearance of these areas in the output. Although the mask-guided approach improves the physical plausibility of underwater refraction modeling, it inherently limits the range of scene features the model can learn.

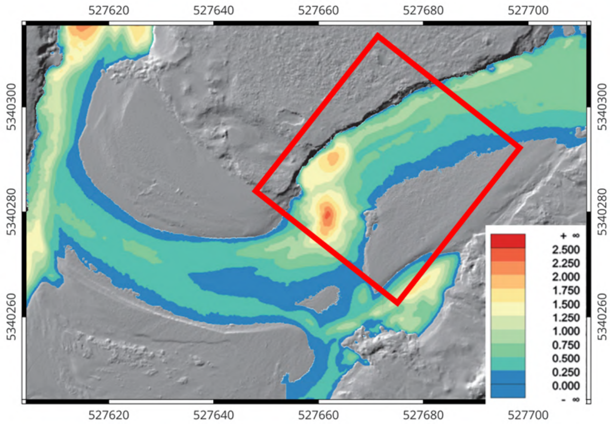

Figure 4. Airborne laser bathymetry (ALB) derived water depth map of the Pielach river used as reference data set for evaluating bathymetric reconstruction accuracy. The red rectangle shows the region of interest for the pointcloud export.

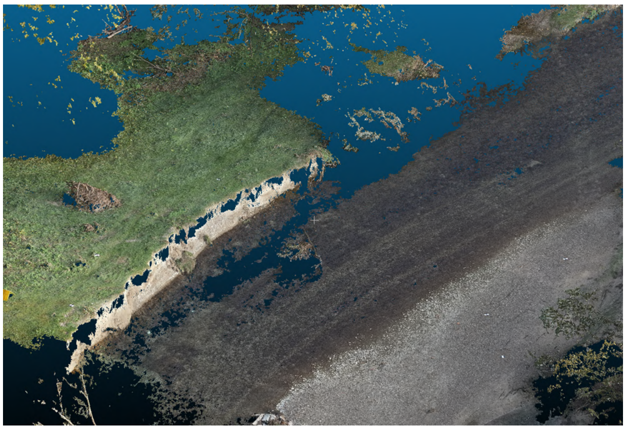

Figure 5. Reference multi-view stereo (MVS) point cloud generated from the UAV image data. The data set includes both submerged riverbed structures and dry gravel surfaces along the shoreline.

5.2 Point Cloud Export

The extraction of 3D point clouds from the trained NeRFrac model represents a key advantage of explicitly modeling refractive effects within the reconstruction process. The NeRFrac framework allows sampling along physically plausible refracted ray paths, which can be directly utilized to reconstruct submerged geometries. Figure 8 shows a representative NeRFrac-derived point cloud generated using the custom export procedure described previously. In this study, point cloud extraction was performed on a model trained without the application of water masks, as the export and verification procedures for masked training configurations are still under development and have not yet been properly validated. The current export pipeline therefore focuses exclusively on full-scene models without prior masking.

The point distribution in Figure 8 reflects the refractive ray paths traced through the water column and illustrates the initial reconstruction capabilities of the NeRFrac model under these conditions. Despite the overall promising results, several challenges still affect the geometric consistency and completeness

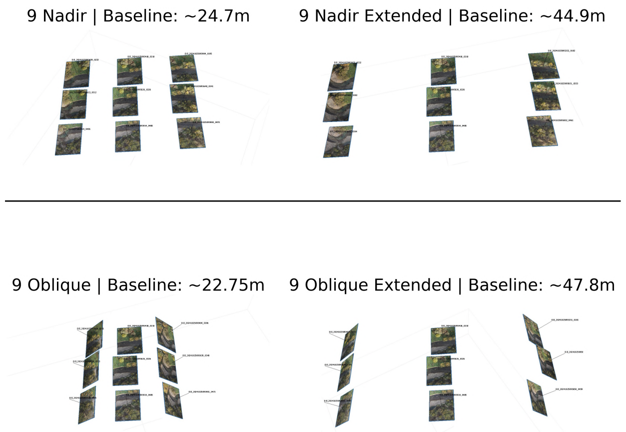

Figure 6. Visualization of the four data set subsets used for NeRFrac evaluation: nadir and oblique configurations with short and extended baselines. The varying acquisition geometries allow systematic evaluation of NeRFrac's reconstruction sensitivity with respect to image configuration.

of the NeRFrac-derived point clouds. In the current state, the exported points remain in a normalized internal coordinate system, which has not yet been transformed back into an external geodetic reference system. As a consequence, the absolute position and scale are not yet metrically consistent with the external data products. Transformation into a global reference frame is currently under development as part of the ongoing optimization of the export pipeline. However, the extracted geometry already shows characteristic terrain features, such as the distinct riverbank slope, which is clearly recognizable in the upper left region of Figure 8. This indicates that the NeRFrac model successfully captures both submerged topography and shoreline discontinuities based on refractive sampling, not only in radiometric but also geometric terms.

For reference and comparison, Figure 5 presents the dense MVS point cloud generated from the UAV images using conventional photogrammetric workflows. This data set is fully georeferenced and provides 3D information of both underwater and terrestrial areas. In contrast to the NeRFrac-derived point cloud, the MVS reconstruction is fully embedded in a global coordinate system and serves as a benchmark for assessing the geometric performance of the NeRFrac approach. However, the MVS reconstruction lacks any consideration of refractive effects, which leads to underestimations of the water depth in the extracted geometries.

The visual comparison of both point clouds highlights the fundamental reconstruction capability of NeRFrac in refractive media, while also revealing current limitations. Although larger-scale topographic features are already well reconstructed, finer details and surface smoothness still show deviations when compared to the MVS reference. Further developments are therefore required, particularly concerning the refinement of the back-transformation, the optimization of the filtering criteria for point extraction, and the incorporation of external reference data to improve metric consistency.

In summary, the presented results demonstrate substantial potential for refractive NeRF approaches in UAV-based optical bathymetry, while also highlighting clear areas for methodological advancement to achieve robust and operationally applicable bathymetric mapping solutions.

Figure 7. Comparison between the original UAV-derived image (top), the binary water mask overlay (middle), and the NeRFrac rendering output (bottom). The brownish appearance of the rendered output is due to sediment-rich underwater regions being exclusively considered during training.

6. Conclusion

This paper investigated the applicability and current limitations of refraction-aware NeRF, using the NeRFrac framework, for 3D reconstruction from UAV imagery over shallow water. The experiments show that incorporating explicit refraction modeling with binary water masks helps to focus the network on relevant underwater regions and allows for the extraction of 3D point clouds.

At the current stage, the presented analysis is limited to qualitative observations. Although initial 3D point clouds have been extracted from the trained NeRFrac model, these data sets are still represented in normalized internal coordinates and have not yet been fully transformed into an external metric reference frame. As a consequence, a quantitative geometric evaluation of 3D accuracy remains open. The first quantitative 2D evaluations of refractive NeRF reconstructions have been conducted and are presented in Guenthner et al. (2025). The ongoing research focuses on further developing the coordinate transformation pipeline and establishing a robust framework for metric 3D evaluation against independent reference data.

Additional limitations remain regarding the underlying assumptions of static water surfaces, as well as the sensitivity of the reconstruction process to variations in image network geometry,

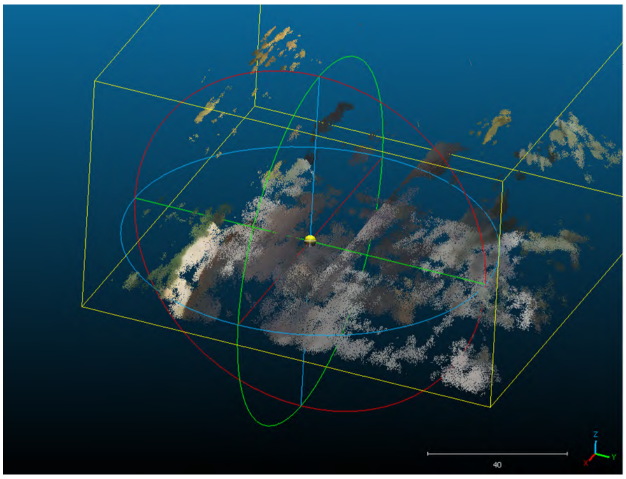

Figure 8. BathyNeRF point cloud derived from NeRFrac export. The data set remains in normalized coordinates; external coordinate transformation is currently under development. Larger terrain structures such as the riverbank are already distinguishable.

sampling density, and filtering strategies. Currently, there are no established guidelines for selecting optimal parameter configurations.

More work is required to refine these methodological aspects, improve metric accuracy, and systematically validate the approach in different aquatic environments. However, the results demonstrate that refractive NeRF models offer considerable potential for underwater photogrammetry and bathymetric mapping.

Acknowledgments

The research presented was carried out within the transnational WEAVE project BathyNeRF funded by the Austrian Science Fund (FWF) [10.55776/PIN1353223] and the German Research Foundation (DFG) [538522540].

References

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. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 5855–5864.

Fujitomi, Y., Saito, S., Li, H., Saito, S., 2022. Lb-nerf: Neural radiance fields with learned bending. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12597–12606.

Förstner, W., Wrobel, B. P., 2016. *Photogrammetric Computer Vision: Statistics, Geometry, Orientation and Reconstruction*. Springer.

Guenthner, A., Brezovsky, M., Schulte, F., Winiwarter, L., Mandlburger, G., Jutzi, B., 2025. Exploring the Potential of Refractive NeRFs for Photogrammetric Bathymetry - First Application to UAV-based Data from the Pielach River. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*.

Kim, J., Lee, S.-W., Lee, S.-H., Kim, M. H., 2024. Ref^2-nerf: Reflection-refraction-aware neural radiance fields. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12345–12354.

Mandlburger, G., 2019. Through-Water Dense Image Matching for Shallow Water Bathymetry. *Photogrammetric Engineering & Remote Sensing*, 85(6), 445–455.

Mandlburger, G., 2022. A REVIEW OF ACTIVE AND PASSIVE OPTICAL METHODS IN HYDROGRAPHY. The International Hydrographic Review, 28, PP. 8-52. *The International Hydrographic Review*, 28. https://ihr.iho.int/articles/a-review-of-active-and-passive-optical-methods-in-hydrography/.

Mandlburger, G., Rhomberg-Kauert, J., Gueguen, L.-A., Mulsow, C., Brezovsky, M., Dammert, L., Haines, J., Glas, S., Himmelsbach, T., Schulte, F., Amon, P., Winiwarter, L., Jutzi, B., Maas, H.-G., 2025a. Mapping shallow inland running waters with UAV-borne photo and laser bathymetry : The Pielach River showcase. *Journal of Applied Hydrography*, 130(3), 42–54.

Mandlburger, G., Rhomberg-Kauert, J., Gueguen, L.-A., Mulsow, C., Brezovsky, M., Dammert, L., Haines, J., Glas, S., Schulte, F., Amon, P., Winiwarter, L., Jutzi, B., Maas, H.-G., 2025b. "pielach river research dataset october 2024 - mapping shallow inland running waters with uav-borne photo and laser bathymetry".

Martin-Brualla, R., Radwan, N., Sajjadi, M. S. M., Barron, J. T., Dosovitskiy, A., Duckworth, D., 2021. Nerf in the wild: Neural radiance fields for unconstrained photo collections. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7210–7219.

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. arXiv:2003.08934.

QGIS Development Team, 2025. QGIS Geographic Information System. QGIS Association.

Schulte, F., Guenthner, A., Brezovsky, M., Mandlburger, G., Jutzi, B., Winiwarter, L., 2025. Simulation and validation of underwater scenes for two-media optical 3D reconstruction. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*.

Shan, J., Toth, C. K., 2018. *Topographic Laser Ranging and Scanning: Principles and Processing*. CRC Press.

Verbin, D., Novotny, D., Pollefeys, M., Snavely, N., Freeman, W. T., Tancik, M., 2022. Ref-nerf: Structured view-dependent appearance for neural radiance fields. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5481–5490.

Xiangli, Y., Xu, L., Pan, X., Zhao, N., Rao, A., Theobalt, C., Dai, B., Lin, D., 2022. Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering. *Proceedings of the European Conference on Computer Vision (ECCV)*, 41–58.

Zhan, Y., Nobuhara, S., Nishino, K., Zheng, Y., 2023. NeRFrac: Neural Radiance Fields through Refractive Surface. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 18356–18366. ISSN: 2380-7504.