

AONeuS: A Neural Rendering Framework for Acoustic-Optical Sensor Fusion

Mohamad Qadri*
mqadri@andrew.cmu.edu
Carnegie Mellon University
Pittsburgh, USA

Kevin Zhang*
kzhang24@umd.edu
University of Maryland
College Park, USA

Akshay Hinduja
ahinduja@andrew.cmu.edu
Carnegie Mellon University
Pittsburgh, USA

Michael Kaess
kaess@cmu.edu
Carnegie Mellon University
Pittsburgh, USA

Adithya Pediredla
Adithya.K.Pediredla@dartmouth.edu
Dartmouth College
Hanover, USA

Christopher A. Metzler†
metzler@umd.edu
University of Maryland
College Park, USA




Figure 1: AONeuS Experimental Results: Under the restricted baseline operating conditions commonly encountered in underwater construction and navigation, camera-only reconstruction techniques (NeuS [Wang et al. 2021b]) and sonar-only reconstruction techniques (NeuSIS [Qadri et al. 2023]) struggle to accurately recover 3D surface geometry. This is due to the highly underdetermined nature of their respective measurement processes; cameras lack depth information, and imaging sonars do not capture elevation information. We have developed a multimodal acoustic-optical neural surfaces reconstruction framework (AONeuS) that effectively combines data from these complementary modalities.

ABSTRACT

Underwater perception and 3D surface reconstruction are challenging problems with broad applications in construction, security, marine archaeology, and environmental monitoring. Treacherous operating conditions, fragile surroundings, and limited navigation control often dictate that submersibles restrict their range of motion and, thus, the baseline over which they can capture measurements. In the context of 3D scene reconstruction, it is well-known that smaller baselines make reconstruction more challenging. Our work develops a physics-based multimodal acoustic-optical neural surface reconstruction framework (AONeuS) capable of effectively integrating high-resolution RGB measurements with low-resolution depth-resolved imaging sonar measurements. By fusing these complementary modalities, our framework can reconstruct accurate high-resolution 3D surfaces from measurements captured over heavily-restricted baselines. Through extensive simulations and in-lab experiments, we demonstrate that AONeuS dramatically outperforms recent RGB-only and sonar-only inverse-differentiable-rendering-based surface reconstruction methods.

CCS CONCEPTS

• Computing methodologies → 3D imaging.

KEYWORDS

3D reconstruction, implicit neural representations, multimodal sensing, robotics, imaging sonar, underwater imaging, signed distance functions, neural rendering, inverse rendering, sensor fusion, underwater sensing

1 INTRODUCTION

The 3D reconstruction of underwater environments is an important problem with applications in myriad fields, including underwater construction, marine ecology, archaeology, mapping, inspection, and surveillance [Albiez et al. 2015; Lin et al. 2023; Negahdaripour 2018; Wang et al. 2019b]. The underwater robots applied to this task are typically equipped with both imaging sonars (i.e., acoustic cameras) and optical cameras [Lensgraf et al. 2021; Liu et al. 2023a]. These sensors capture complementary information about their operating environments.

Forward-look imaging sonars consist of a uniform linear array of transducers which, through beamforming, recover both range and azimuth information (but not elevation). (3D imaging sonars

*Both authors contributed equally to the paper.

†Corresponding author.

which record both azimuth and elevation also exist, but can be prohibitively expensive.) Unlike light-based sensors, imaging sonars are highly robust to scattering and low-light conditions. Unfortunately, imaging sonars generally have poor spatial resolution; sonar images of an object of interest often appear textureless and hard to recognize; and imaging sonar measurements can suffer from complex artifacts caused by multipath reflections and the variable speed of sound passing through inhomogeneous water [Tinh and Khanh 2021].

By contrast, optical cameras have high spatial resolution and can resolve object appearance in great detail. However, in turbid water light scattering and absorption can severely restrict the range and contrast of optical cameras [Jaffe 2014]. Moreover, to recover depth information passive optical sensors rely on large displacements/baselines between measurements. In constrained operating environments, such measurements are often inaccessible.

By leveraging the complementary strengths and weaknesses of cameras and imaging sonars, acoustic-optical sensor fusion promises to enable robust and high-resolution underwater perception and scene reconstruction [Ferreira et al. 2016; Menna et al. 2018]. Existing contour matching based acoustic-optical reconstruction methods can already reconstruct accurate high-resolution 3D surfaces [Babae and Negahdaripour 2015]. Unfortunately, these methods require a 360-degree view of the scene and are inapplicable in the small-baseline operating conditions prevalent in real-world unmanned underwater vehicle operation. Alternatively, one can reconstruct the scene from optical and acoustic measurements independently and then fuse the result [Kim et al. 2019]. However, this simple approach provides limited benefits over camera-only surface reconstruction.

In this paper, we develop an inverse-differentiable-rendering-based approach to acoustic-optical sensor fusion that can form dense 3D surface reconstructions from camera and sonar measurements captured across a small baseline. Our work consists of four key contributions.

- We develop a physics-based multimodal acoustic-optical neural surface framework which simultaneously integrates RGB and imaging sonar measurements. Our approach extends the neural surfaces 3D reconstruction framework [Wang et al. 2021b] by combining a unified representation of the scene geometry with modality-specific (acoustic and optical) representations of appearance.
- We conduct experiments on both synthetic and experimentally-captured datasets and demonstrate our method can effectively reconstruct high-fidelity surface geometry from noisy measurements captured over limited baselines.
- We theoretically support our strong empirical performance by analyzing the conditioning of the acoustic-optical forward model. We show that the forward process associated with triangulating a point in 3D from acoustic-optical measurements is better conditioned and easier to invert the unimodal forward models.
- We release a public dataset and open-source implementation of our method.

2 RELATED WORK

Camera Imaging. Myriad works have investigated the use of (optical) cameras for underwater 3D imaging. Johnson-Roberson et al. [2010] proposed a feature-based stereo-optical SLAM system for building 3D models. Iscar et al. [2017] presented a comprehensive evaluation of different monocular and stereo software and hardware systems targeting underwater imaging. However, these techniques, which do not use sonar, have limited depth resolution in small baseline scenarios. Roznere et al. [2023] proposed a multi-view photometric stereo method for non-stationary underwater robots 3D reconstruction that integrates ORB-SLAM [Mur-Artal et al. 2015] with traditional photometric stereo. However, this technique requires active illumination and is sensitive to backscattering in turbid waters.

Sonar Imaging. 3D reconstruction from sonar imagery is an important and widely studied problem. Over the last decade a variety of 3D reconstruction methods have been proposed based on space carving [Aykin and Negahdaripour 2015, 2016a], classical point-cloud processing algorithms [Teixeira et al. 2016; Westman et al. 2020a], generative modeling [Aykin and Negahdaripour 2015, 2016b; Negahdaripour et al. 2017; Westman and Kaess 2019], convex optimization [Westman et al. 2020b], graph-based processing [Wang et al. 2019a, 2018], and supervised machine learning [Arnold and Wehbe 2022; DeBortoli et al. 2019; Wang et al. 2021a].

Last year, two research groups employed neural rendering to enable breakthrough 3D sonar imaging performance. Qadri et al. [2023] developed a Neural Implicit Surface Reconstruction Using Imaging Sonar (NeuSIS) method which forms high-fidelity 3D surface reconstructions from forward imaging sonar measurements by combining neural surface representations with a novel acoustic differentiable volumetric renderer. Similarly, Reed et al. [2023] employed neural rendering to recover 3D volumes from synthetic aperture sonar measurements. The former method relies upon a large number of sonar images captured over a large baseline while the latter applies to synthetic aperture sonar, not forward imaging sonar, and relies on access to raw time-based sonar measurements. These methods represent the state-of-the-art in 3D surface reconstruction with sonar.

To date, no method has effectively recovered a dense 3D scene from 2D sonar images captured over a limited baseline. Without additional constraints, e.g., optical measurements, or strong priors the reconstruction problem is hopelessly underdetermined.

Neural Rendering. In their breakthrough neural radiance fields (NeRF) paper, Mildenhall et al. [2020] combined neural signal representations with differentiable volume rendering to perform novel view synthesis. The underlying differentiable volume rendering concept has since been extended to represent and recover scene geometry. The Implicit Differentiable Renderer (IDR) approach, introduced in [Yariv et al. 2020], represents geometry as the zero-level set of a neural network and uses differentiable surface rendering to fit the parameters of a neural network. IDR requires object masks for supervision. Later methods, like Neural Surfaces (NeuS) [Wang et al. 2021b], Unified Surfaces (UNISURF) [Oechsle et al. 2021], and Volume Signed Distance Functions (VolSDF) [Yariv et al. 2021] combine an implicit surface representation with differentiable volume

rendering to recover 3D geometry from images without the need for object masks. Recent work has sought to reduce the number of training images required [Long et al. 2022] and to accelerate rendering to enable real-time applications [Yariv et al. 2023]. The inverse-differentiable-rendering framework has also been extended to handle measurements from a diverse range of sensors. Cross-spectral radiance fields (X-NeRF) were proposed in [Poggi et al. 2022] to model multispectral, infrared, and RGB images. Transient neural radiance fields were proposed in [Malik et al. 2023] to model the measurements from a single-photon lidar. Time-of-flight radiance fields were proposed in [Attal et al. 2021] to model the measurements from a continuous wave time-of-flight sensor. Polarization-aided decomposition of radiance, or PANDORA, was proposed in [Dave et al. 2022] to model polarimetric measurements of light. Radar neural radiance fields (RaNeRF) were proposed in [Liu et al. 2023b] to model inverse synthetic aperture radar measurements.

Recent research on neural fields has also addressed the small baseline 3D reconstruction problem we tackle in our paper, but leveraging natural small hand motions to improve depth reconstruction instead of using multiple imaging modalities [Chugunov et al. 2024, 2023, 2022].

Several recent works have modeled light scattering within the neural rendering framework to improve reconstructions through water [Levy et al. 2023; Sethuraman et al. 2023], haze [Chen et al. 2024], and fog [Ramazzina et al. 2023].

Multimodal Imaging. To overcome the disadvantages inherent to using a single sensing modality, numerous multimodal sensing algorithms have been developed [Bijelic et al. 2020; Kim et al. 2009; Lindell et al. 2018; Nishimura et al. 2020]. Most related to our work, Babaee and Negahdaripour [2015] reconstruct 3D objects from RGB and sonar imagery by matching occluding contours across RGB images and imaging sonar measurements, performing stereo matching, and interpolating the curves in 3D space. Unfortunately, this method is inapplicable to the small-baseline setting; it fundamentally requires 360-degree views of the scene. Cardaillac and Ludvigsen [2023] similarly use a camera and an imaging sonar for 3D reconstruction by matching features between the acoustic and optical measurements. However, this matching can be frail and prone to errors. Kim et al. [2019] reconstructs a scene from optical and acoustic measurements independently using classical methods like COLMAP [Schonberger and Frahm 2016] and then fuses the result. As we demonstrate in section 6.3, this approach provides limited benefits over a purely optical approach and is not competitive with state-of-the-art neural rendering based methods.

Outside of sonar, several neural-rendering based approaches to sensor fusion have recently been developed. In Multimodal Neural Radiance Field, Zhu et al. [2023] use neural rendering to combine RGB, thermal, and point cloud data. Similarly, [Kim et al. 2023] use neural rendering to combine multispectral measurements of different polarizations and Carlson et al. [2023] fuse sparse lidar and RGB measurements to build 3D occupancy grid of unbounded scenes.

To our knowledge, ours is the first work to perform acoustic-optical sensor fusion with neural rendering.




Figure 2: Acoustic-Optical Measurement Processes. (a) RGB measurement process and example measurement. Pixels along a common ray passing through the camera center map to the same image pixel on the image plane. (b) Sonar measurement process and example measurement. In a sonar image, the azimuth θ and range r of the imaged object are resolved. However, the elevation information ϕ is lost; all objects located along the elevation arc (in blue) map to the same pixel.

3 BACKGROUND

3.1 Imaging Sonars

Imaging sonars are active sensors that emit acoustic pulses and measure the intensity of the reflected wave. They produce a 2D acoustic image in which the range and azimuth of the imaged object are resolved. However, the object’s elevation remains ambiguous. I.e., the reflecting object can be located anywhere on the elevation arc (fig. 2) and the intensity of a pixel in a sonar image is proportional to the cumulative reflected acoustic energy from all reflecting points along the elevation arc.

3.2 Image Formation Model of an Imaging Sonar

Similar to [Qadri et al. 2023], we use the following sonar image formation model:

$$I^{\text{son}}(r_i, \theta_i) = \int_{\phi_{\min}}^{\phi_{\max}} \int_{r_i - \epsilon}^{r_i + \epsilon} \frac{E_e}{r} T(r, \theta_i, \phi) \sigma(r, \theta_i, \phi) dr d\phi, \quad (1)$$

where $\phi_{\min}, \phi_{\max}$ are the minimum and maximum elevation angles, E_e is the acoustic energy emitted by the sonar. $T = e^{-\int_0^r \sigma(r', \theta_i, \phi_i) dr'}$ is the transmittance term, and σ is the particle density. (See [Qadri et al. 2023] for more details.)

3.3 Image Formation Model of an Optical Camera

We adopt the optical camera image formation model proposed by [Wang et al. 2021b] where a pixel intensity at (x, y) is approximated

by:

$$I^{\text{cam}}(x, y) = \int_0^\infty T(t)\sigma(t)c(\mathbf{p}(t), \mathbf{v})dt, \quad (2)$$

where the integral is over the ray starting at the camera center and passing through pixel (x, y) . T, σ are the transmittance and density values at point $\mathbf{p}(t)$, and $c(\mathbf{p}(t), \mathbf{v})$ is the color of a point viewed from direction $\mathbf{v}$.

4 PROBLEM STATEMENT




Figure 3: Acoustic-Optical Measurement Ambiguities. (a) Two RGB measurements captured over a limited baseline struggle to localize a point along the depth-axis. (b) Two sonar measurements captured over a limited baseline struggle to localize a point along the x-axis. Because they have orthogonal ambiguities, RGB and sonar measurements are highly complementary.

Our goal in this work is to reconstruct the 3D surface of an underwater object using a small collection of RGB and sonar measurements captured over a limited baseline. Specifically, we assume access to two datasets, $\mathcal{D}^{\text{cam}} = \{I_i^{\text{cam}}, p_i^{\text{son}}\}$ and $\mathcal{D}^{\text{son}} = \{I_i^{\text{son}}, p_i^{\text{son}}\}$, consisting of RGB/sonar images and their respective poses.

Given a large dataset captured over a sufficiently diverse range of poses (e.g., thousands of images captured from 360-degrees [Qadri et al. 2023]), existing unimodal (camera-only/sonar-only) surface reconstruction methods are already effective [Qadri et al. 2023; Wang et al. 2021b]. In this work, we focus on the small baseline operating conditions—pervasive in underwater robotics—where optical cameras record insufficient information to recover depth information (see fig. 3(a)) and imaging sonars record insufficient information to recover elevation information (see fig. 3(b)).

Specifically, we introduce a physics-based multimodal inverse-differentiable-rendering framework that integrates information from both acoustic and optical sensors to generate accurate 3D reconstructions. Our approach automatically exploits the complementary information (elevation/range) provided by each sensor.

Because our shared pool-based testing facility does not allow us to introduce turbidity, in this work we focus on the clear-water setting. Modeling the effects of light scattering in our forward model [Chen et al. 2024; Levy et al. 2023; Ramazzina et al. 2023; Sethuraman et al. 2023] would likely improve our system’s in-the-wild performance.

5 METHOD

5.1 Acoustic-Optical NeuS

Our AONeS reconstruction framework is illustrated in fig. 4. Following Qadri et al. [2023]; Wang et al. [2021b], we represent the object’s surface using a Signed Distance Function (SDF), $N(\mathbf{x})$, which outputs the distance of each 3D point $\mathbf{x} = (X, Y, Z)$ to the nearest surface. Distinct from these works, we use two separate rendering neural networks ($\mathbf{M}^{\text{cam}}$ and $\mathbf{M}^{\text{son}}$) that approximate the optical and acoustic outgoing radiance at each spatial coordinate $\mathbf{x}$. This choice is motivated by the fact that different materials have different acoustic and optical reflectance properties. For example, glass is invisible to optical cameras but visible to imaging sonar, and PVC is invisible to imaging sonar but visible to optical cameras.

In this work, we sample and sum points along acoustic and optical rays to approximate the rendering integrals defined by eq. (1) and eq. (2). Our rendering functions can be expressed as

$$\hat{I}^{\text{son}}(r, \theta) = \sum_{\mathbf{x} \in \mathcal{A}_{p^{\text{son}}}} \frac{1}{r(\mathbf{x})} T[\mathbf{x}] \alpha[\mathbf{x}] \mathbf{M}^{\text{son}}(\mathbf{x}), \text{ and} \quad (3)$$

$$\hat{I}^{\text{cam}}(x, y) = \sum_{\mathbf{x} \in \mathcal{R}_{p^{\text{cam}}}} T[\mathbf{x}] \alpha[\mathbf{x}] \mathbf{M}^{\text{cam}}(\mathbf{x}), \quad (4)$$

where $\mathcal{A}_{p^{\text{son}}}$ is the set of sampled points along the acoustic arc at pixel p^{son} and $\mathcal{R}_{p^{\text{cam}}}$ is the set of sampled points along optical ray passing through pixel p^{cam} . $\mathbf{M}^{\text{son}}$ and $\mathbf{M}^{\text{cam}}$ are the predicted radiances at $\mathbf{x}$.

The computation of the discrete transmittance and opacity terms in eq. (3) and eq. (4) requires sampling along both acoustic and optical rays. For any such spatial sample $\mathbf{x}_s$, (i.e., any point along an acoustic or optical ray), the discrete opacity at $\mathbf{x}_s$ can be approximated as

$$\alpha[\mathbf{x}_s] = \max \left(\frac{\Phi_q(N(\mathbf{x}_s)) - \Phi_q(N(\mathbf{x}_{s+1}))}{\Phi_q(N(\mathbf{x}_s))}, 0 \right), \quad (5)$$

where $\Phi_q(x) = (1 + e^{-qx})^{-1}$ is the Sigmoid function and q is a trainable parameter. The discrete transmittance is modeled as

$$T[\mathbf{x}_s] = \prod_{\mathbf{x}_r \mid r < s} (1 - \alpha[\mathbf{x}_r]). \quad (6)$$

5.1.1 Loss Function. Our loss function comprises the sonar and camera intensity losses:

$$\mathcal{L}_{\text{int}}^{\text{son}} \equiv \frac{1}{N_{\mathcal{P}^{\text{son}}}} \sum_{p \in \mathcal{P}^{\text{son}}} \|\hat{I}(p) - I(p)\|_1, \text{ and} \quad (7)$$

$$\mathcal{L}_{\text{int}}^{\text{cam}} \equiv \frac{1}{N_{\mathcal{P}^{\text{cam}}}} \sum_{p \in \mathcal{P}^{\text{cam}}} \|\hat{I}(p) - I(p)\|_1, \quad (8)$$

where $\mathcal{P}^{\text{cam}}$ and $\mathcal{P}^{\text{son}}$ is the set of sampled pixels in the camera and sonar images respectively. We additionally use the eikonal loss as an implicit regularization to encourage smooth reconstructions:

$$\mathcal{L}_{\text{eik}} \equiv \frac{1}{|\mathbf{X}|} \sum_{\mathbf{x} \in \mathbf{X}} (\|\nabla N(\mathbf{x})\|_2 - 1)^2, \quad (9)$$

where $\mathbf{X}$ is the set of all sampled points.

We also utilize an ℓ_1 loss term as an additional prior term which biases the network towards reconstructions that minimize the total




Figure 4: AONeuS Reconstruction Framework. A shared surface geometry SDF network N is used in combination with rendering specific neural rendering modules. For each sampled point x along an acoustic or optical ray, N outputs its signed distance, its gradient as well as 2 features vectors F^{son} and F^{cam} all serving as input to their respective rendering networks. D^{son} and D^{cam} are respectively the directions of the acoustic and optical rays.

opacity of the scene (for example in cases where the object is on the seafloor and only specific sides can be imaged):

$$\mathcal{L}_{\text{reg}} \equiv \frac{1}{|X|} \sum_{x \in X} \|\alpha[x]\|_1. \quad (10)$$

Hence, our total loss is

$$\mathcal{L} = \alpha(t) \mathcal{L}_{\text{int}}^{\text{son}} + (1 - \alpha(t)) \mathcal{L}_{\text{int}}^{\text{cam}} + \lambda_{\text{eik}} \mathcal{L}_{\text{eik}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}. \quad (11)$$

The network is trained with the ADAM optimizer.

5.1.2 Weight Scheduling. The weights assigned to the sonar and camera intensity losses (respectively $\alpha(t)$ and $1 - \alpha(t)$ in eq. 11) impact the reconstruction quality as they determine which measurements the network should emphasize throughout training. We adopt a simple two-step weighting scheme:

$$\alpha(t) = \begin{cases} 1 & \text{if } t < E_t, \\ \lambda & \text{if } E_t < t < E_e. \end{cases} \quad (12)$$

In the early iterations, $t < E_t$, the sonar measurements are used exclusively and serve to "mask" the object; i.e., update the weights of the SDF network N to bias it towards reconstructions in which the geometry of the object are better constrained in the depth direction. This process establishes an initialization for the later iterations.

In later iterations, $t > E_t$, more emphasis is placed on the camera measurements. These measurements constrain the x and y directions and help resolve the elevation ambiguity inherent in sonar data. In this phase, sonar measurements receive less weight and




Figure 5: Simulation setup. We visualize the orientations of the camera and the sonar relative to the scene, here a turtle, in our simulation. The optical axis of the camera, the z axis of its own local coordinate frame, is aligned with the Z axis of the world coordinate frame. The elevation axis of the sonar, the z axis of its own local coordinate frame, is aligned with the $-X$ axis of the world coordinate frame. Additionally, we visualize the image planes of the camera and sonar, the 2D planes onto which they project the 3D scene points.

act as a depth regularizer. See section A.3 of the appendix for a comparison of our mixing procedure against different weighting schemes.

6 EXPERIMENTAL RESULTS

In this section, we evaluate the AONeuS on both synthetic and experimentally captured data. Hyperparameters for our experiments can be found in section A.3 of the appendix material.

6.1 Notation

We use uppercase X, Y, Z to refer to the world coordinate system and lowercase x, y, z to refer to sensor-specific coordinate systems. See fig. 5.

6.2 Results on Synthetic Data

To generate synthetic measurements, we used a custom-made sonar simulator as well as Blender [Community 2022] for RGB measurements to collect simulated sonar-camera datasets for various objects. The objects are assumed to be acoustically diffuse. We investigate the impact specular reflections have on the performance of our method in section A.1 of the appendix, which also lists the values of the simulation parameters we used. The sonar and camera are approximately collocated, and are translated linearly over a short baseline along the X axis of the world frame for a distance of 1.2 m with the sonar's azimuthal plane parallel the YZ plane in the world frame. The sonar's azimuthal plane is oriented orthogonal to the direction of motion to ensure the trajectory was non-degenerate; multiple measurements captured from positions within the azimuthal plane of the sonar would be highly redundant and uninformative [Negahdaripour 2018].

For each object, the trajectory is sub-sampled into smaller baselines: 0.96 m, 0.72 m, 0.48 m, and 0.24 m for analysis. We scaled the meshes so that the objects are approximately ~ 1 m in size and the sensors are placed about 1.5 m–2 m away from the object. The

elevation aperture of the sonar is 12° . We benchmark our method against two methods: NeuS [Wang et al. 2021b] and NeuSIS [Qadri et al. 2023], executing all methods 9 times with randomly initialized seeds. To ensure we had reasonable camera-only results, we provided NeuS with masks of the object. This information is not required by nor provided to NeuSIS and AONeuS.

In fig. 8(a), we compare the reconstruction performance of all three techniques for a total of five scenes. We could observe that AONeuS consistently reconstructs the scene geometry better than NeuS and NeuSIS. Further, we can also observe that NeuS (camera-only) incorrectly reconstructs the depth axis (Z-axis) whereas NeuSIS (sonar-only) can reconstruct only the depth-axis accurately. The proposed AONeuS was able to recover underlying scene geometry along all the axes.

In fig. 8(b), we show the results for the turtle mesh for various baselines. To visualize the ambiguities associated with camera and sonar modalities and the benefit of the fusion algorithm, we rendered the reconstructed meshes with a virtual camera pointing in Y-axis. Hence, the rendered images are projections of the reconstructed mesh on ZX-plane. As we decrease the baseline (top to bottom), for NeuS, we observe an increasing loss of features along depth direction: the back legs of the turtle are progressively lost and depth-reconstruction worsens with decreasing baselines. For sonar-only methods, significant ambiguities along the elevation axis can be seen across all baselines: due to the limited translation of the sonar, the collected measurements are not enough to constrain and resolve the turtle shell adequately. Our framework AONeuS integrates orthogonal information from both imaging modalities to yield reconstructions of higher quality across all baselines: all features of the turtle including its shell and its back legs are clearly discernible. These observations are further supported by the quantitative analysis in table 1 where we report the mean and standard deviation of the Chamfer L_1 distance, precision, and recall of the reconstructions over nine trials. The results demonstrate that AONeuS outperforms the existing methods, particularly with reduced baselines. Note that recall of NeuSIS appears to be slightly better than AONeuS but that is only because the NeuSIS generates a large blob that covers most part of the object. The per-baseline

Table 1: Metrics for synthetic turtle data. Best metrics are bolded.

		NeuS	NeuSIS	AONeuS
1.2m	Chamfer ↓	0.123 ± 0.028	0.130 ± 0.013	0.075 ± 0.006
	Precision ↑	0.653 ± 0.095	0.566 ± 0.043	0.862 ± 0.042
	Recall ↑	0.526 ± 0.134	0.836 ± 0.022	0.825 ± 0.056
0.96m	Chamfer ↓	0.139 ± 0.024	0.134 ± 0.011	0.079 ± 0.005
	Precision ↑	0.602 ± 0.076	0.531 ± 0.031	0.840 ± 0.017
	Recall ↑	0.470 ± 0.132	0.816 ± 0.031	0.807 ± 0.017
0.72m	Chamfer ↓	0.205 ± 0.027	0.135 ± 0.011	0.081 ± 0.005
	Precision ↑	0.423 ± 0.051	0.537 ± 0.062	0.810 ± 0.032
	Recall ↑	0.279 ± 0.071	0.768 ± 0.029	0.792 ± 0.022
0.48m	Chamfer ↓	0.249 ± 0.045	0.139 ± 0.012	0.088 ± 0.006
	Precision ↑	0.337 ± 0.062	0.470 ± 0.028	0.791 ± 0.023
	Recall ↑	0.189 ± 0.074	0.706 ± 0.028	0.770 ± 0.023
0.24m	Chamfer ↓	0.406 ± 0.087	0.146 ± 0.009	0.111 ± 0.017
	Precision ↑	0.223 ± 0.060	0.450 ± 0.028	0.690 ± 0.045
	Recall ↑	0.107 ± 0.049	0.587 ± 0.042	0.679 ± 0.042




Figure 6: Experimental hardware setup. (a) Test water tank used to conduct the experiments and its dimensions. (b) Test object. (c) Bluefin Hovering Autonomous Underwater Vehicle (HAUV) and its mounted hardware (Didson imaging sonar and Doppler Velocity Log (DVL)). (d) FLIR Blackfly S GigE camera used for image capture and its watertight enclosure.

quantitative results for the remaining four meshes can be found in section A.5 of the appendix.

6.3 Results on Experimentally-Captured Data

We also perform real-world experiments on an object (fig. 6b) submerged in a water tank (fig. 6a). Please check the supplementary video for more visualizations of the setup. The object is made of standard wooden plywood of $\sim 0.013\text{m}$ thickness and has an acoustic impedance of $2.5 \times 10^6 \text{ kg/m}^2\text{s}$. Our object is also covered in a thin layer of insulating foam which increases its surface roughness and makes it more diffuse. We used a SoundMetrics DIDSON imaging sonar mounted on a Bluefin Hovering Autonomous Underwater Vehicle (HAUV) (fig. 6c) to capture two sonar datasets of the test object with two different elevation apertures 14° and 28° . The sonar operates at a frequency of 1.8MHz corresponding to a wavelength of 0.833mm .¹ The vehicle uses an IMU and a Doppler Velocity Log (DVL) to measure sonar pose information. We asynchronously capture optical images of the same object using a FLIR

¹assuming the speed of sound in water to equal $\sim 1500\text{m/s}$

Blackfly S Gige 5MP camera (fig. 6d) with camera pose information computed with COLMAP. To reduce the water-glass-air effect inside the camera, we used a checkerboard, underwater, to measure our camera’s underwater intrinsic parameters. The sonar and camera trajectories were aligned post-capture. Similar to the simulation setup, both camera and sonar followed an approximately 1.2m non-degenerate linear trajectory (~ 120 total sonar and RGB measurements), which we later sub-sampled into the same 5 base-lines. Because any additional measurements captured along that trajectory are highly redundant, adding additional measurements (without added viewpoint diversity) has a limited effect on reconstruction accuracy. We benchmarked our method against three algorithms: The COLMAP based sensor fusion method introduced in [Kim et al. 2019]², NeuS [Wang et al. 2021b], and NeuSIS [Qadri et al. 2023]. For each dataset and sensor baseline, we executed each method six times with randomly initialized seeds except that of [Kim et al. 2019], which is deterministic.

Qualitatively, we observe in fig. 9 that AONeuS outputs a more complete shape across baselines compared to sonar-only (NeuSIS) and camera-only (NeuS) only methods: the hole, two legs, and cross-bar are clearly discernible. Conversely, when using only sonar, parts of the object are not well reconstructed as we can observe, for example, with the long leg with NeuSIS at 14° . Similarly, camera-only methods result in the loss of features such as the hole accompanied with significant introduced depth errors. We quantify the results in table 10 of the appendix, in section A.5, where we report the mean and standard deviation of the Chamfer L_1 distance, precision, and recall against the ground truth mesh computed over six trials with different random seeds for training. We observe that the fusion of the acoustic and optical signals generates higher quality reconstruction, even with very short baselines measuring only 24 cm, as indicated by the mean value of each metric. When comparing AONeuS with sonar-only methods (NeuSIS), we note that, despite the increased elevation ambiguity introduced by the 28° elevation aperture, our technique is able to leverage camera information and its constraints in the x and y axes to resolve spatial locations that are otherwise under-constrained when solely relying on sonar. Techniques that rely on a camera only (NeuS) exhibit a decrease in performance as the sensor baseline is reduced. Complementing camera with sonar information introduces constraints in the depth direction easing the resolution of depth which is known to be difficult to resolve with limited camera motion. We additionally emphasize the standard deviation of the Chamfer distance over the random seeds: the fusion of both modalities result in outputs that are more robust to the randomness of the algorithm (i.e. network initialization, point samples, etc.).

7 DISCUSSION AND ANALYSIS

7.1 Distribution of per-Axis Errors

In fig. 10, we visualize the per-axis deviations from the ground truth for the synthetic turtle scene at 0.24 m baseline. We compute per-axis deviations by first determining the closest vertex in the dense ground truth mesh and taking absolute differences in x , y , and z coordinates. We histogram these deviations along all three

axes and show them along rows in the fig. 10. We have repeated this procedure for NeuS, NeuSIS, and AONeuS and show them along the columns.

From the data, we can observe (1) NeuS has large deviations along Z axes, (2) NeuSIS has large deviations along X axes. These results are consistent with the ambiguities associated with their respective measurement processes. AONeuS has low spread on all axes as it captures the best of both camera (NeuS) and sonar (NeuSIS) imaging modalities. Further analysis of the reconstruction error can be found in section A.4 of the appendix.

7.2 Multimodal Sensing is Better Conditioned

The strong empirical performance of our multimodal reconstructions can be explained in terms of system conditioning. Given point correspondences between measurements, it is far easier to triangulate a point using multimodal acoustic-optical measurements than camera-only or sonar-only measurements.

Previous works have analyzed the conditioning of an imaging system consisting of an acoustic-optical stereo pair [Negahdaripour et al. 2009], e.g., a sonar (only) on the left and a camera (only) on the right. In this work, we follow a similar analysis to characterize a multimodal stereo pair, e.g., one co-located sonar and camera pair on the left and another co-located sonar and camera pair on the right. Our results follow closely from Negahdaripour’s analysis of a stereo pair of 2D imaging sonars Negahdaripour [2018]; our stereo camera measurements introduce four additional linear constraints to the system’s forward model.

Consider a point $P = [X, Y, Z]^T$ that is observed by an acoustic-optical sensor from two positions. The sonar’s azimuthal plane is the yz plane, in its own coordinate system. The camera’s image plane is the $z = f$ plane, in its own coordinate system. Without loss of generality, assume the sensor’s coordinate system at its initial location is the world coordinate system and its coordinate system at its second position is described by a rotation $\mathbf{R}$ and translation $\mathbf{t} = [t_x, t_y, t_z]$. That is, the coordinate of point P in the new coordinate system is $P' = \mathbf{R}P + \mathbf{t}$.

Under this model, the acoustic-optical sensor records 8 measurements:

$$\begin{aligned} x_c &= f \frac{[1, 0, 0]P}{[0, 0, 1]P}, & y_c &= f \frac{[0, 1, 0]P}{[0, 0, 1]P}, \\ R &= \|\mathbf{P}\|, & \theta &= \tan^{-1} \left(\frac{[0, 1, 0]P}{[0, 0, 1]P} \right), \\ x'_c &= f \frac{\mathbf{r}_1^t P + t_x}{\mathbf{r}_3^t P + t_z}, & y'_c &= f \frac{\mathbf{r}_2^t P + t_y}{\mathbf{r}_3^t P + t_z}, \\ R' &= \sqrt{\|\mathbf{R}\|^2 + \|\mathbf{t}\|^2 + 2\mathbf{t}^t \mathbf{R}P}, & \theta' &= \tan^{-1} \left(\frac{\mathbf{r}_2^t P + t_y}{\mathbf{r}_3^t P + t_z} \right), \end{aligned} \quad (13)$$

where $\mathbf{r}_i$ denotes the i^{th} row of $\mathbf{R}$.

Following Negahdaripour [2018]; Negahdaripour et al. [2009], we can turn each of these measurements into seven linear constraints

²COLMAP outputs a sparse pointcloud. Hence, a mesh was computed using the ball pivoting algorithm [Bernardini et al. 1999].




Figure 7: System Conditioning. Histograms of the condition numbers of the camera-only (top), sonar-only (middle), and multimodal (bottom) forward models. Median condition numbers are highlighted in red. The acoustic-optical multimodal forward model is generally better conditioned and easier to invert (triangulation).

and one non-linear constraint on P .

$$A_{multi}P = b \text{ and } \|P\|^2 = R^2 \text{ with}$$

$$A_{multi} = \begin{bmatrix} (-f, 0, x_c) \\ (0, -f, y_c) \\ (0, -1, \tan(\theta)) \\ x'_c r_3^t - f r_2^t \\ y'_c r_3^t - f r_2^t \\ \tan(\theta') r_3^t - r_2^t \\ t^t R \end{bmatrix} \text{ and } b = \begin{bmatrix} 0 \\ 0 \\ 0 \\ f t_x - x'_c t_z \\ f t_y - y'_c t_z \\ t_y - \tan(\theta') t_z \\ \frac{(R')^2 - R^2 - \|t\|^2}{2} \end{bmatrix}. \quad (14)$$

One can similarly form camera-only, A_{cam} , and sonar-only, A_{son} , forward models by considering only rows 1, 2, 4, and 5 and rows 3, 6, and 7, respectively, of A_{multi} . By inverting these systems, one can triangulate P in space. Here we perform Monte Carlo sampling to compare the conditioning of A_{cam} , A_{son} , and A_{multi} . We sample P uniformly in a 1 m^3 cube centered at $(0, 0, 1.5)$ with edges parallel to the x , y , and z axis; we assume $f = 100 \text{ mm}$; we sample t_x , t_y , and t_z uniformly in the range 0 cm to 10 cm ; and we sample the yaw, pitch, and roll between measurements uniformly in the range -5° to 5° .

For each realization of these parameters, we compute the condition number, κ , of A_{cam} , A_{son} , and A_{multi} . We repeat this process 50,000 times to form histograms, illustrated in fig. 7. The condition number of the multimodal system is generally much lower and the system is thus easier to invert; multimodal triangulation is easier as each sensor does indeed contribute complementary constraints.

8 LIMITATIONS

Currently, our approach is limited to clear-water settings and does not accurately model effects such as optical scattering and water

absorption. Hence, using more sophisticated optical physical models will be important when moving towards open-sea experiments. Similar to other differentiable rendering methods, our approach is too slow for online reconstruction (currently $\sim 30 \text{ min}$) per scene on an Nvidia 3090 GPU. Enabling real-time reconstruction is an important direction for future work. Finally, we assume that our robot's trajectory is nearly orthogonal to the sonar's azimuthal plane: When the trajectory is within this plane additional sonar measurements provide little information.

9 CONCLUSION

We have introduced and validated a multimodal inverse-differentiable-rendering framework for reconstructing 3D surface information from camera and sonar measurements. Our framework combines camera and sonar information using a unified surface representation module and separate modality-specific appearance modules and rendering functions. By extracting information from these complementary modalities, our framework is able to offer breakthrough underwater sensing capabilities for restricted baseline imaging scenarios. We have demonstrated that AONeuS can accurately reconstruct the geometry of complex 3D objects from synthetic as well as noisy, real-world measurements captured over severely restricted baselines.

While we demonstrate the first neural fusion of camera and sonar measurements, there are many interesting directions to explore this amalgamation. In section 5.1.2, we introduced a heuristic for weighing camera and sonar measurements. A structured way of combining the camera and sonar data, which is aware of the uncertainties [Goli et al. 2023; Jiang et al. 2023] in the complementary imaging systems could result in faster convergence rates and better reconstructions.

The sonar we have used in our implementations are forward-looking sonars. Fusion algorithms for side-scan sonars, synthetic-aperture sonars, sonars of different ranges and wavelengths, could be an interesting forward direction. Similarly, extending the technique for various geometries and materials including multi-object scenes, dynamic scenes, cluttered scenes and scattering media (murky water) would make AONeuS more practical. Finally, on-the-fly reconstructions could allow one to select the best next underwater view to improve reconstruction accuracy and further reduce the required baseline and acquisition time.

ACKNOWLEDGMENTS

The authors would like to thank Tianxiang Lin and Jui-Te Huang for their help with data collection; Sarah Friday for providing an animation showcasing the real experimental setup; and Shahriar Negahdaripour for helpful discussions, feedback, and guidance. M.Q., A.H., and M.K. were supported in part by ONR grant N00014-21-1-2482. A.P. was supported by NSF grant 2326904. K.Z. and C.A.M. were supported in part by AFOSR Young Investigator Program award no. FA9550-22-1-0208, ONR award no. N00014-23-1-2752, and a seed grant from SAAB, Inc.

REFERENCES

- Jan Albiez, Sylvain Joyeux, Christopher Gaudig, Jens Hilljegerdes, Sven Kroffke, Christian Schoo, Sascha Arnold, Geovane Mimoso, Pedro Alcantara, Rafael Saback, et al. 2015. Flatfish-a compact subsea-resident inspection auv. In *OCEANS 2015-MTS/IEEE Washington*. IEEE, IEEE, 201 Waterfront Street National Harbor, Maryland 20745 USA, 1–8.
- Sascha Arnold and Bilal Wehbe. 2022. Spatial Acoustic Projection for 3D Imaging Sonar Reconstruction. In *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, Philadelphia, 3054–3060. <https://doi.org/10.1109/ICRA46639.2022.9812277>
- Benjamin Attal, Eliot Laidlaw, Aaron Gokaslan, Changil Kim, Christian Richardt, James Tompkin, and Matthew O’ Toole. 2021. TöRF: Time-of-Flight Radiance Fields for Dynamic Scene View Synthesis. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., Virtual, Web, 26289–26301.
- Murat D Aykin and Shahriar Negahdaripour. 2015. On 3-D target reconstruction from multiple 2-D forward-scan sonar views. In *OCEANS 2015-Genova*. IEEE, Institute of Electrical and Electronics Engineers, Genova, Italy, 1–10.
- Murat D Aykin and Shahriar Negahdaripour. 2016a. Three-dimensional target reconstruction from multiple 2-d forward-scan sonar views by space carving. *IEEE Journal of Oceanic Engineering* 42, 3 (2016), 574–589.
- Murat D Aykin and Shahriar S Negahdaripour. 2016b. Modeling 2-D lens-based forward-scan sonar imagery for targets with diffuse reflectance. *IEEE journal of oceanic engineering* 41, 3 (2016), 569–582.
- Mohammadreza Babae and Shahriar Negahdaripour. 2015. 3-D object modeling from 2-D occluding contour correspondences by opti-acoustic stereo imaging. *Computer Vision and Image Understanding* 132 (2015), 56–74.
- Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Cláudio Silva, and Gabriel Taubin. 1999. The ball-pivoting algorithm for surface reconstruction. *IEEE transactions on visualization and computer graphics* 5, 4 (1999), 349–359.
- Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix Heide. 2020. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Institute of Electrical and Electronics Engineers, Virtual, Web, 11682–11692.
- Alexandre Cardaillac and Martin Ludvigsen. 2023. Camera-Sonar Combination for Improved Underwater Localization and Mapping. *IEEE Access* 11 (2023), 123070–123079. <https://doi.org/10.1109/ACCESS.2023.3329834>
- Alexandra Carlson, Manikandasirram S. Ramanagopal, Nathan Tseng, Matthew Johnson-Roberson, Ram Vasudevan, and Katherine A. Skinner. 2023. CLONeR: Camera-Lidar Fusion for Occupancy Grid-Aided Neural Representations. *IEEE Robotics and Automation Letters* 8, 5 (2023), 2812–2819. <https://doi.org/10.1109/LRA.2023.3262139>
- W. Chen, W. Yifan, S. Kuo, and G. Wetzstein. 2024. DehazeNeRF: Multiple Image Haze Removal and 3D Shape Reconstruction using Neural Radiance Fields. In *3DV*. IEEE Computer Society, Davos, Switzerland, no–pagination.
- Ilya Chugunov, David Shustin, Ruyi Yan, Chenyang Lei, and Felix Heide. 2024. Neural Spline Fields for Burst Image Fusion and Layer Separation. *CVPR* (2024).
- Ilya Chugunov, Yuxuan Zhang, and Felix Heide. 2023. Shakes on a Plane: Unsupervised Depth Estimation from Unstabilized Photography. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 13240–13251.
- Ilya Chugunov, Yuxuan Zhang, Zhihao Xia, Xuaner Zhang, Jiawen Chen, and Felix Heide. 2022. The Implicit Values of A Good Hand Shake: Handheld Multi-Frame Neural Depth Refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2852–2862.
- Blender Online Community. 2022. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam. <http://www.blender.org>
- Akshat Dave, Yongyi Zhao, and Ashok Veeraraghavan. 2022. *PANDORA: Polarization-Aided Neural Decomposition of Radiance*. Springer Nature Switzerland, Tel Aviv, Israel, 538–556. https://doi.org/10.1007/978-3-031-20071-7_32
- Robert DeBortoli, Fuxin Li, and Geoffrey A Hollinger. 2019. Elevatenet: A convolutional neural network for estimating the missing dimension in 2d underwater sonar images. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Institute for Electrical and Electronics Engineers, Macau, China, 8040–8047.
- Fausto Ferreira, Diogo Machado, Gabriele Ferri, Samantha Dugelay, and John Potter. 2016. Underwater Optical and Acoustic Imaging: A Time for Fusion? A Brief Overview of the State-of-the-Art. In *OCEANS 2016 MTS/IEEE Monterey*. Institute of Electrical and Electronics Engineers, Monterey, USA, 1–6. <https://doi.org/10.1109/OCEANS.2016.7761354>
- Lily Goli, Cody Reading, Silvia Sellán, Alec Jacobson, and Andrea Tagliasacchi. 2023. Bayes’ Rays: Uncertainty Quantification for Neural Radiance Fields. arXiv:2309.03185 [cs.CV]
- Eduardo Iscar, Katherine A Skinner, and Matthew Johnson-Roberson. 2017. Multi-view 3D reconstruction in underwater environments: Evaluation and benchmark. In *OCEANS 2017-Anchorage*. IEEE, Institute of Electrical and Electronics Engineers, Anchorage, Alaska, 1–8.
- Jules S Jaffe. 2014. Underwater optical imaging: the past, the present, and the prospects. *IEEE Journal of Oceanic Engineering* 40, 3 (2014), 683–700.
- Wen Jiang, Boshu Lei, and Kostas Daniilidis. 2023. FisherRF: Active View Selection and Uncertainty Quantification for Radiance Fields using Fisher Information. arXiv:2311.17874 [cs.CV]
- Matthew Johnson-Roberson, Oscar Pizarro, Stefan B Williams, and Ian Mahon. 2010. Generation and visualization of large-scale three-dimensional reconstructions from underwater robotic surveys. *Journal of Field Robotics* 27, 1 (2010), 21–51.
- Jason Kim, Meungsuk Lee, Seokyeong Song, Byeongjin Kim, and Son-Cheol Yu. 2019. 3-D Reconstruction of Underwater Objects Using Image Sequences from Optical Camera and Imaging Sonar. In *OCEANS 2019 MTS/IEEE SEATTLE*. Institute of Electrical and Electronics Engineers, Seattle, USA, 1–6. <https://doi.org/10.23919/OCEANS40490.2019.8962558>
- Youngchan Kim, Wonjoon Jin, Sunghyun Cho, and Seung-Hwan Baek. 2023. Neural Spectro-polarimetric Fields. In *SIGGRAPH Asia 2023 Conference Papers* (, Sydney, NSW, Australia.) (SA ’23). Association for Computing Machinery, New York, NY, USA, Article 109, 11 pages. <https://doi.org/10.1145/3610548.3618172>
- Young Min Kim, Christian Theobalt, James Diebel, Jana Kosecka, Branislav Miscusik, and Sebastian Thrun. 2009. Multi-view image and ToF sensor fusion for dense 3D reconstruction. In *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*. Institute of Electrical and Electronics Engineers, Kyoto, Japan, 1542–1549. <https://doi.org/10.1109/ICCVW.2009.5457430>
- D. Langer and M. Hebert. 1991. Building Qualitative Elevation Maps from Side Scan Sonar Data for Autonomous Underwater Navigation. In *1991 IEEE International Conference on Robotics and Automation Proceedings*. Institute of Electrical and Electronics Engineers, Sacramento, USA, 2478–2483 vol.3. <https://doi.org/10.1109/ROBOT.1991.131997>
- Samuel Lensgraf, Amy Sniffen, Zachary Zitzewitz, Evan Honnold, Jennifer Jain, Weifu Wang, Alberto Li, and Devin Balkcom. 2021. Droplet: Towards Autonomous Underwater Assembly of Modular Structures. In *Robotics: Science and Systems XVII*. Robotics: Science and Systems Foundation, Virtual, Web, no–pagination. <https://doi.org/10.15607/RSS.2021.XVII.054>
- Deborah Levy, Amit Peleg, Naama Pearl, Dan Rosenbaum, Derya Akkaynak, Simon Korman, and Tali Treibitz. 2023. SeaThru-NeRF: Neural Radiance Fields in Scattering Media. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Institute of Electrical and Electronics Engineers, Vancouver, Canada, 56–65.
- Tianxiang Lin, Akshay Hinduja, Mohamad Qadri, and Michael Kaess. 2023. Conditional GANs for Sonar Image Filtering with Applications to Underwater Occupancy Mapping. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, Institute of Electronics and Electrical Engineers, London, England, 1048–1054.
- David B. Lindell, Matthew O’Toole, and Gordon Wetzstein. 2018. Single-Photon 3D Imaging with Deep Sensor Fusion. *ACM Transactions on Graphics* 37, 4 (July 2018), 113:1–113:12. <https://doi.org/10.1145/3197517.3201316>
- Afei Liu, Shuanghui Zhang, Chi Zhang, Shuaifeng Zhi, and Xiang Li. 2023b. RaNeRF: Neural 3-D Reconstruction of Space Targets From ISAR Image Sequences. *IEEE Transactions on Geoscience and Remote Sensing* 61 (2023), 1–15. <https://doi.org/10.1109/TGRS.2023.3298067>
- Haowen Liu, Monika Roznere, and Alberto Quattrini Li. 2023a. Deep Underwater Monocular Depth Estimation with Single-Beam Echosounder. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, Institute of Electrical and Electronics Engineers, London, England, 1090–1097.
- Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. 2022. SparseNeuS: Fast Generalizable Neural Surface Reconstruction from Sparse Views. In *Computer Vision – ECCV 2022*, Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (Eds.). Vol. 13692. Springer Nature Switzerland, Cham, 210–227. https://doi.org/10.1007/978-3-031-19824-3_13
- Anagh Malik, Parsa Mirdehghan, Sotiris Niousias, Kyros Kutulakos, and David Lindell. 2023. Transient Neural Radiance Fields for Lidar View Synthesis and 3D Reconstruction. In *Advances in Neural Information Processing Systems*, Vol. 36. Curran Associates, Inc., New Orleans, USA, 71569–71581.
- Fabio Menna, Panagiotis Agraftotis, and Andreas Georgopoulos. 2018. State of the art and applications in archaeological underwater 3D recording and mapping. *Journal of Cultural Heritage* 33 (2018), 231–248.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I*. Springer-Verlag, Berlin, Heidelberg, 405–421. https://doi.org/10.1007/978-3-030-58452-8_24
- Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. 2015. ORB-SLAM: a versatile and accurate monocular SLAM system. *IEEE transactions on robotics* 31, 5 (2015), 1147–1163.
- S.K. Nayar, K. Ikeuchi, and T. Kanade. 1991. Surface Reflection: Physical and Geometrical Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13, 7 (July 1991), 611–634. <https://doi.org/10.1109/34.85654>
- Shahriar Negahdaripour. 2018. Application of forward-scan sonar stereo for 3-D scene reconstruction. *IEEE journal of oceanic engineering* 45, 2 (2018), 547–562.
- Shahriar Negahdaripour, Victor M Milenkovic, Nikan Salarieh, and Mahsa Mirzargar. 2017. Refining 3-D object models constructed from multiple FS sonar images by space carving. In *OCEANS 2017-Anchorage*. IEEE, Institute of Electrical and

- Electronics Engineers, Anchorage, USA, 1–9.
- Shahriar Negahdaripour, Hicham Sekkati, and Hamed Pirsiavash. 2009. Opti-acoustic stereo imaging: On system calibration and 3-D target reconstruction. *IEEE Transactions on image processing* 18, 6 (2009), 1203–1214.
- Mark Nishimura, David B Lindell, Christopher Metzler, and Gordon Wetzstein. 2020. Disambiguating monocular depth estimation with a single transient. In *European Conference on Computer Vision*. Springer, Springer, Virtual, Web, 139–155.
- Michael Oechsle, Songyou Peng, and Andreas Geiger. 2021. UNISURF: Unifying Neural Implicit Surfaces and Radiance Fields for Multi-View Reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. Institute of Electrical and Electronics Engineers, Virtual, Web, 5589–5599.
- M. Poggi, P. Ramirez, F. Tosi, S. Salti, S. Mattoccia, and L. Stefano. 2022. Cross-Spectral Neural Radiance Fields. In *2022 International Conference on 3D Vision (3DV)*. IEEE Computer Society, Los Alamitos, CA, USA, 606–616. <https://doi.org/10.1109/3DV57658.2022.00071>
- Mohamad Qadri, Michael Kaess, and Ioannis Gkioulekas. 2023. Neural implicit surface reconstruction using imaging sonar. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, Institute of Electrical and Electronics Engineers, London, USA, 1040–1047.
- Andrea Ramazzina, Mario Bijelic, Stefanie Walz, Alessandro Sanvito, Dominik Scheuble, and Felix Heide. 2023. ScatterNeRF: Seeing Through Fog with Physically-Based Inverse Neural Rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. Institute of Electrical and Electronics Engineers, Paris, France, 17957–17968.
- Albert Reed, Juhyeon Kim, Thomas Blanford, Adithya Pediredla, Daniel Brown, and Suren Jayasuriya. 2023. Neural Volumetric Reconstruction for Coherent Synthetic Aperture Sonar. *ACM Transactions on Graphics* 42, 4 (July 2023), 113:1–113:20. <https://doi.org/10.1145/3592141>
- Monika Roznere, Philippos Mordohai, Ioannis Rekleitis, and Alberto Quattrini Li. 2023. 3-D Reconstruction Using Monocular Camera and Lights: Multi-View Photometric Stereo for Non-Stationary Robots. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, Institute of Electronics and Electrical Engineers, London, England, 1026–1032.
- Johannes L Schonberger and Jan-Michael Frahm. 2016. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. Institute of Electrical and Electronics Engineers, Las Vegas, USA, 4104–4113.
- Advait Venkatramanan Sethuraman, Manikandasriram Srinivasan Ramanagopal, and Katherine A Skinner. 2023. Waternerf: Neural radiance fields for underwater scenes. In *OCEANS 2023-MTS/IEEE US Gulf Coast*. IEEE, Institute of Electrical and Electronics Engineers, Biloxi, USA, 1–7.
- Pedro V Teixeira, Michael Kaess, Franz S Hover, and John J Leonard. 2016. Underwater inspection using sonar-based volumetric submaps. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Institute of Electrical and Electronics Engineers, Daejeon, South Korea, 4288–4295.
- Nguyen Dinh Tinh and T Dang Khanh. 2021. A new imaging geometry model for multi-receiver synthetic aperture sonar considering variation of the speed of sound in seawater. *IEIE Transactions on Smart Processing and Computing* 10, 4 (2021), 302–308.
- K. E. Torrance and E. M. Sparrow. 1967. Theory for Off-Specular Reflection From Roughened Surfaces*. *Journal of the Optical Society of America* 57, 9 (Sept. 1967), 1105. <https://doi.org/10.1364/JOSA.57.001105>
- Jinkun Wang, Tixiao Shan, and Brendan Englot. 2019b. Underwater terrain reconstruction from forward-looking sonar imagery. In *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, Institute of Electrical and Electronics Engineers, Montreal, Canada, 3471–3477.
- Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. 2021b. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. *Advances in Neural Information Processing Systems* 34 (2021), 27171–27183.
- Yusheng Wang, Yonghoon Ji, Dingyu Liu, Hiroshi Tsuchiya, Atsushi Yamashita, and Hajime Asama. 2021a. Elevation angle estimation in 2d acoustic images using pseudo front view. *IEEE Robotics and Automation Letters* 6, 2 (2021), 1535–1542.
- Yusheng Wang, Yonghoon Ji, Hanwool Woo, Yusuke Tamura, Atsushi Yamashita, and Hajime Asama. 2019a. Three-dimensional underwater environment reconstruction with graph optimization using acoustic camera. In *2019 IEEE/SICE International Symposium on System Integration (SII)*. IEEE, Institute of Electrical and Electronics Engineers, Paris, France, 28–33.
- Yusheng Wang, Yonghoon Ji, Hanwool Woo, Yusuke Tamura, Atsushi Yamashita, and Asama Hajime. 2018. 3D occupancy mapping framework based on acoustic camera in underwater environment. *IFAC-PapersOnLine* 51, 22 (2018), 324–330.
- Eric Westman, Ioannis Gkioulekas, and Michael Kaess. 2020a. A theory of fermat paths for 3d imaging sonar reconstruction. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Institute of Electrical and Electronics Engineers, Las Vegas, USA, 5082–5088.
- Eric Westman, Ioannis Gkioulekas, and Michael Kaess. 2020b. A volumetric albedo framework for 3D imaging sonar reconstruction. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, Institute of Electrical and Electronics Engineers, Virtual, Web, 9645–9651.
- Eric Westman and Michael Kaess. 2019. Wide aperture imaging sonar reconstruction using generative models. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Institute of Electrical and Electronics Engineers, Macau, China, 8067–8074.
- Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. 2021. Volume Rendering of Neural Implicit Surfaces. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., Virtual, Web, 4805–4815.
- Lior Yariv, Peter Hedman, Christian Reiser, Dor Verbin, Pratul P. Srinivasan, Richard Szeliski, Jonathan T. Barron, and Ben Mildenhall. 2023. BakedSDF: Meshing Neural SDFs for Real-Time View Synthesis. In *ACM SIGGRAPH 2023 Conference Proceedings (SIGGRAPH '23)*. Association for Computing Machinery, New York, NY, USA, 1–9. <https://doi.org/10.1145/3588432.3591536>
- Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. 2020. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems* 33 (2020), 2492–2502.
- Haidong Zhu, Yuyin Sun, Chi Liu, Lu Xia, Jiajia Luo, Nan Qiao, Ram Nevatia, and Cheng-Hao Kuo. 2023. Multimodal Neural Radiance Field. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. Institute of Electrical and Electronics Engineers, London, England, 9393–9399. <https://doi.org/10.1109/ICRA48891.2023.10160388>




Figure 8: Simulated Small-Baseline Results. In (a), we visualize the reconstructions in simulation for all objects at 0.4x baseline, or 0.48 meters. In (b), we visualize the reconstructions in simulation for the turtle across all baseline settings.




Figure 9: Experimental Results with "h" Object. As the baseline diminishes, NeuS exhibits increasing amount of distortion along the depth direction as can be seen at the intersection of the short piling and crossbar at the 0.72m and 0.96m baselines. NeuSIS similarly generates reconstructions with significant errors (for example, the long piling is poorly reconstructed with the 14° elevation). Conversely, AONeuS consistently produces faithful reconstructions across a range of baselines.




Figure 10: Per-axis error distributions. At 0.2x baseline for the turtle example, we plot the distributions of deviations from the ground truth mesh along all three axes for NeuS, NeuSIS, and AONeUS reconstructions. The NeuS reconstruction has larger Z errors, noticeable from the long tail, whereas NeuSIS reconstruction has larger X errors. AONeUS has tighter distributions along all three axes compared to NeuS or NeuSIS showing that the proposed technique takes the best of both of the imaging modalities.

A SUPPLEMENT

A.1 Effect of Acoustic Specularity on the Reconstruction Accuracy

We investigate the reconstruction accuracy of AONeuS when imaging acoustically-specular objects. Consider the following sonar measurement formation model [Langer and Hebert 1991] that has a diffuse and specular component,

$$I_r = \underbrace{C_{dl} \cos \alpha}_{\text{diffuse}} + \underbrace{C_{sl} G(\alpha) \frac{1}{\cos \alpha} \exp\left(-\frac{\alpha^2}{2\sigma_\alpha^2}\right)}_{\text{specular}}, \quad (15)$$

where I_r is the intensity of the reflection, α is the angle of incidence, C_{dl} and C_{sl} represent how strong the diffuse and specular components are, relatively, σ_α is the standard deviation of slope of the microfacet distribution which models the surface roughness, and $G(\alpha)$ is a geometric attenuation factor [Nayar et al. 1991] which represents how the microfacets might occlude each other. For sonar,

$$G(\alpha) = \min(1, 2 \cos^2(\alpha)). \quad (16)$$

We note that this is a modification of the extension of the Torrance-Sparrow reflection model [Torrance and Sparrow 1967] developed in [Nayar et al. 1991] that excludes the direct specular spike component of the reflection and retains only the diffuse and specular lobe components of the reflection. The rationale is that for a sonar, the wave source and receiver are at the same position, so the sensor receives only reflected in a direction towards it. Therefore, the receiver will only receive the specular spike if the surface is close to perfectly normal to the sensor, which rarely happens in practice. The reflection geometry is described in more detail in fig. 11. For the experiments in section 6.2, we set $C_{dl} = 1$ and $C_{sl} = 0$.




Figure 11: Comparing the received intensity for different incidence angles. In (a), the sensor receives all components of the reflection: the diffuse lobe, specular lobe, and the specular spike. In (b), the sensor only receives the specular lobe and the diffuse lobes, because those reflection components have portions which are reflected under the normal to the surface, where the sensor is, whereas the specular spike is completely reflected above the normal of the surface, away from the sensor.

In the following experiments, we focus on the specular component of eq.15 (i.e. set $C_{dl} = 0$, $C_{sl} = 1$) and investigate the effect of

varying the parameter σ_α (which models the width of the specular lobe for purely acoustic specular reflections) on the reconstruction quality of AONeuS. All experiments were ran on the 0.24 m baseline. We visualize the qualitative results of the experiments in fig. 12, which show that AONeuS can generate accurate reconstructions even from specular imaging sonar data. Of note is that as the width of the specular lobe decreases and becomes more narrow, the quality of the reconstruction decreases, implying that performing 3D reconstruction becomes more difficult. The qualitative trends are also validated by the quantitative results in table 2 to table 4. Here, we note that for all three meshes we consider, AONeuS performs best when $\sigma_\alpha = 1$, or when the specular lobe is the widest, and for certain geometries, like the turtle, the reconstruction performance degrades at $\sigma_\alpha = 0.1$.




Figure 12: Reconstructions by AONeuS from RGB images and specular imaging sonar measurements. As the width of the specular lobe decreases, i.e. as σ_α decreases, we observe that the performance of AONeuS decreases.

Table 2: Quantitative metrics, airplane mesh, specular data.

	Chamfer L1 ↓	Precision ↑	Recall ↑
$\sigma_\alpha = 1.0$	0.141 ± 0.026	0.558 ± 0.059	0.667 ± 0.058
$\sigma_\alpha = 0.5$	0.150 ± 0.020	0.503 ± 0.064	0.649 ± 0.039
$\sigma_\alpha = 0.1$	0.149 ± 0.027	0.445 ± 0.079	0.629 ± 0.054

Table 3: Quantitative metrics, lobster mesh, specular data.

	Chamfer L1 ↓	Precision ↑	Recall ↑
$\sigma_\alpha = 1.0$	0.186 ± 0.053	0.440 ± 0.142	0.439 ± 0.174
$\sigma_\alpha = 0.5$	0.183 ± 0.035	0.468 ± 0.113	0.441 ± 0.120
$\sigma_\alpha = 0.1$	0.206 ± 0.040	0.329 ± 0.077	0.473 ± 0.142

Table 4: Quantitative metrics, turtle mesh, specular data.

	Chamfer L1 ↓	Precision ↑	Recall ↑
$\sigma_\alpha = 1.0$	0.116 ± 0.018	0.686 ± 0.045	0.691 ± 0.045
$\sigma_\alpha = 0.5$	0.104 ± 0.009	0.717 ± 0.048	0.720 ± 0.045
$\sigma_\alpha = 0.1$	0.191 ± 0.027	0.533 ± 0.079	0.615 ± 0.061

A.2 Hyperparameters

Table 5: List of hyperparameters.

Parameter	Sonar dataset 1 14° elevation angle	Sonar dataset 2 28° elevation angle	Simulation
E_t	4000	4000	2000
E_e	8000	8000	5000
λ	0.3	0.3	0.3
λ_{eik}	0.01	0.1	0.1
λ_{reg}	0.1	1	0

A.3 Ablation Study (Weighting Schemes)

Table 6: Various weighting scheme. Experimenting on the real object with 0.24m baseline. Values are averaged over 6 trials. Constant: fixed weights. Linear: Weights change linearly from initial to end values. Step: Weights are switched from start to end value at iteration E_t .

Camera + sonar at 14° elevation					
Mode	$\alpha(t)$ start	$\alpha(t)$ end	Chamfer L1 ↓	Precision ↑	Recall ↑
Constant	0.5	0.5	0.120 ± 0.026	0.586 ± 0.055	0.624 ± 0.124
Constant	0.7	0.7	0.141 ± 0.046	0.582 ± 0.078	0.524 ± 0.133
Constant	0.3	0.3	0.093 ± 0.009	0.626 ± 0.049	0.741 ± 0.035
Linear	1	0	0.181 ± 0.036	0.513 ± 0.040	0.388 ± 0.088
Linear	0	1	0.108 ± 0.019	0.597 ± 0.056	0.670 ± 0.092
Step	0	0.7	0.113 ± 0.010	0.542 ± 0.042	0.695 ± 0.014
Step (Ours)	1	0.3	0.085 ± 0.009	0.706 ± 0.063	0.758 ± 0.041
Camera + sonar at 28° elevation					
Mode	$\alpha(t)$ start	$\alpha(t)$ end	Chamfer L1 ↓	Precision ↑	Recall ↑
Constant	0.5	0.5	0.131 ± 0.014	0.437 ± 0.048	0.604 ± 0.061
Constant	0.7	0.7	0.121 ± 0.008	0.513 ± 0.038	0.585 ± 0.036
Constant	0.3	0.3	0.141 ± 0.010	0.380 ± 0.042	0.595 ± 0.046
Linear	1	0	0.140 ± 0.027	0.498 ± 0.076	0.590 ± 0.101
Linear	0	1	0.144 ± 0.015	0.388 ± 0.068	0.594 ± 0.098
Step	0	0.7	0.139 ± 0.015	0.454 ± 0.068	0.572 ± 0.086
Step (Ours)	1	0.3	0.108 ± 0.005	0.538 ± 0.022	0.695 ± 0.036

A.4 Variance of the Density Field Over Realizations of the Algorithm




Figure 13: Distribution of Reconstructions of the Real Object. Top: Different Reconstructions from different random seeds for the 0.24m baseline with the sonar elevation at 14°. Bottom: Histogram distribution of the standard deviation of the density field vs. voxel count. Neus’s distribution is heavily tailed while NeuSIS’s distribution exhibits a large mean and variance. AONeUS’s distribution is light-tailed with a small mean and therefore it is better constrained.

A.5 Additional tables

Tables 7 to 10 provide additional quantitative metrics for our synthetic experiments.

Table 7: Quantitative metrics for the airplane mesh.

		NeuS	NeuSIS	AONeuS
1.2m	Chamfer ↓	0.112 ± 0.018	0.197 ± 0.011	0.117 ± 0.014
	Precision ↑	0.652 ± 0.058	0.295 ± 0.019	0.582 ± 0.053
	Recall ↑	0.650 ± 0.044	0.643 ± 0.025	0.741 ± 0.028
0.96m	Chamfer ↓	0.144 ± 0.021	0.200 ± 0.019	0.134 ± 0.016
	Precision ↑	0.559 ± 0.045	0.291 ± 0.014	0.575 ± 0.017
	Recall ↑	0.579 ± 0.042	0.650 ± 0.043	0.697 ± 0.027
0.72m	Chamfer ↓	0.146 ± 0.021	0.200 ± 0.016	0.141 ± 0.023
	Precision ↑	0.554 ± 0.052	0.289 ± 0.029	0.558 ± 0.035
	Recall ↑	0.599 ± 0.039	0.629 ± 0.067	0.689 ± 0.048
0.48m	Chamfer ↓	0.174 ± 0.016	0.199 ± 0.012	0.146 ± 0.033
	Precision ↑	0.468 ± 0.039	0.287 ± 0.047	0.533 ± 0.087
	Recall ↑	0.516 ± 0.040	0.569 ± 0.076	0.668 ± 0.044
0.24m	Chamfer ↓	0.223 ± 0.046	0.182 ± 0.011	0.166 ± 0.034
	Precision ↑	0.341 ± 0.090	0.358 ± 0.042	0.451 ± 0.103
	Recall ↑	0.413 ± 0.072	0.555 ± 0.069	0.644 ± 0.045

Table 8: Quantitative metrics for the lobster mesh.

		NeuS	NeuSIS	AONeuS
1.2m	Chamfer ↓	0.147 ± 0.017	0.187 ± 0.012	0.105 ± 0.019
	Precision ↑	0.448 ± 0.073	0.328 ± 0.020	0.592 ± 0.110
	Recall ↑	0.454 ± 0.100	0.460 ± 0.037	0.626 ± 0.090
0.96m	Chamfer ↓	0.167 ± 0.027	0.203 ± 0.014	0.142 ± 0.074
	Precision ↑	0.394 ± 0.070	0.302 ± 0.025	0.534 ± 0.230
	Recall ↑	0.398 ± 0.106	0.445 ± 0.033	0.533 ± 0.228
0.72m	Chamfer ↓	0.191 ± 0.019	0.225 ± 0.030	0.128 ± 0.021
	Precision ↑	0.357 ± 0.053	0.275 ± 0.040	0.533 ± 0.112
	Recall ↑	0.417 ± 0.077	0.390 ± 0.040	0.576 ± 0.105
0.48m	Chamfer ↓	0.237 ± 0.039	0.243 ± 0.019	0.143 ± 0.015
	Precision ↑	0.321 ± 0.062	0.257 ± 0.012	0.523 ± 0.069
	Recall ↑	0.392 ± 0.091	0.370 ± 0.017	0.577 ± 0.043
0.24m	Chamfer ↓	0.293 ± 0.044	0.272 ± 0.055	0.186 ± 0.034
	Precision ↑	0.251 ± 0.049	0.225 ± 0.039	0.440 ± 0.116
	Recall ↑	0.277 ± 0.103	0.290 ± 0.076	0.483 ± 0.130

Table 9: Quantitative metrics for the seastar mesh.

		NeuS	NeuSIS	AONeuS
1.2m	Chamfer ↓	0.108 ± 0.020	0.177 ± 0.010	0.088 ± 0.022
	Precision ↑	0.585 ± 0.070	0.335 ± 0.021	0.714 ± 0.106
	Recall ↑	0.722 ± 0.061	0.894 ± 0.031	0.886 ± 0.026
0.96m	Chamfer ↓	0.122 ± 0.030	0.170 ± 0.014	0.089 ± 0.016
	Precision ↑	0.539 ± 0.117	0.352 ± 0.032	0.720 ± 0.063
	Recall ↑	0.689 ± 0.122	0.848 ± 0.022	0.829 ± 0.033
0.72m	Chamfer ↓	0.159 ± 0.035	0.171 ± 0.010	0.126 ± 0.035
	Precision ↑	0.435 ± 0.089	0.352 ± 0.033	0.571 ± 0.118
	Recall ↑	0.550 ± 0.127	0.791 ± 0.026	0.764 ± 0.092
0.48m	Chamfer ↓	0.255 ± 0.041	0.170 ± 0.006	0.143 ± 0.046
	Precision ↑	0.175 ± 0.059	0.372 ± 0.019	0.486 ± 0.121
	Recall ↑	0.201 ± 0.081	0.742 ± 0.039	0.657 ± 0.080
0.24m	Chamfer ↓	0.548 ± 0.144	0.191 ± 0.006	0.196 ± 0.033
	Precision ↑	0.067 ± 0.038	0.344 ± 0.023	0.377 ± 0.088
	Recall ↑	0.071 ± 0.045	0.627 ± 0.069	0.516 ± 0.072

Table 10: Quantitative metrics for the shell mesh.

		NeuS	NeuSIS	AONeuS
1.2m	Chamfer ↓	0.063 ± 0.002	0.077 ± 0.010	0.066 ± 0.006
	Precision ↑	0.847 ± 0.011	0.754 ± 0.066	0.858 ± 0.024
	Recall ↑	0.941 ± 0.018	0.814 ± 0.035	0.844 ± 0.038
0.96m	Chamfer ↓	0.068 ± 0.005	0.078 ± 0.012	0.078 ± 0.006
	Precision ↑	0.833 ± 0.023	0.756 ± 0.072	0.816 ± 0.022
	Recall ↑	0.929 ± 0.020	0.803 ± 0.043	0.794 ± 0.031
0.72m	Chamfer ↓	0.078 ± 0.006	0.091 ± 0.016	0.090 ± 0.006
	Precision ↑	0.769 ± 0.035	0.747 ± 0.086	0.774 ± 0.029
	Recall ↑	0.876 ± 0.044	0.746 ± 0.075	0.730 ± 0.021
0.48m	Chamfer ↓	0.107 ± 0.014	0.107 ± 0.017	0.098 ± 0.009
	Precision ↑	0.620 ± 0.064	0.791 ± 0.077	0.714 ± 0.031
	Recall ↑	0.690 ± 0.084	0.676 ± 0.078	0.691 ± 0.037
0.24m	Chamfer ↓	0.199 ± 0.029	0.168 ± 0.008	0.109 ± 0.010
	Precision ↑	0.309 ± 0.028	0.832 ± 0.041	0.667 ± 0.048
	Recall ↑	0.302 ± 0.029	0.465 ± 0.032	0.640 ± 0.047

Table 11: For the hardware reconstruction of "H" object, we report the mean and standard deviation of the Chamfer L1 distance, precision, and recall (with a threshold of 0.05 m) compared to the ground truth (obtained from a laser scan of the real structure) for various reconstruction techniques. We computed the standard deviation over 6 trials. For all methods, we compute the metrics for the intermediate reconstructions throughout training and report the best results.

Baseline	Metric	NeuS	Kim et al. (2019)	Sonar dataset 1 14° elevation angle		Sonar dataset 2 28° elevation angle	
				NeuSIS (14°)	AONeuS (14°)	NeuSIS (28°)	AONeuS (28°)
1.2m	Chamfer L1 ↓	0.092 ± 0.015	0.177	0.159 ± 0.032	0.092 ± 0.007	0.151 ± 0.014	0.101 ± 0.007
	Precision ↑	0.693 ± 0.066	0.336	0.553 ± 0.074	0.661 ± 0.040	0.513 ± 0.061	0.569 ± 0.066
	Recall ↑	0.679 ± 0.098	0.387	0.417 ± 0.056	0.708 ± 0.046	0.583 ± 0.068	0.702 ± 0.049
0.96m	Chamfer L1 ↓	0.107 ± 0.013	0.182	0.158 ± 0.023	0.088 ± 0.007	0.154 ± 0.012	0.093 ± 0.004
	Precision ↑	0.661 ± 0.048	0.318	0.560 ± 0.068	0.687 ± 0.019	0.500 ± 0.026	0.602 ± 0.045
	Recall ↑	0.563 ± 0.084	0.345	0.420 ± 0.058	0.722 ± 0.039	0.562 ± 0.042	0.723 ± 0.034
0.72m	Chamfer L1 ↓	0.127 ± 0.013	0.178	0.167 ± 0.029	0.095 ± 0.008	0.154 ± 0.013	0.088 ± 0.003
	Precision ↑	0.651 ± 0.047	0.368	0.562 ± 0.077	0.667 ± 0.041	0.502 ± 0.054	0.636 ± 0.025
	Recall ↑	0.500 ± 0.062	0.396	0.402 ± 0.057	0.688 ± 0.008	0.586 ± 0.043	0.748 ± 0.026
0.48m	Chamfer L1 ↓	0.150 ± 0.022	0.179	0.170 ± 0.028	0.089 ± 0.005	0.143 ± 0.007	0.086 ± 0.001
	Precision ↑	0.626 ± 0.055	0.324	0.543 ± 0.045	0.668 ± 0.006	0.547 ± 0.021	0.636 ± 0.022
	Recall ↑	0.415 ± 0.022	0.218	0.395 ± 0.066	0.726 ± 0.021	0.605 ± 0.021	0.757 ± 0.019
0.24m	Chamfer L1 ↓	0.167 ± 0.012	0.198	0.163 ± 0.019	0.085 ± 0.009	0.148 ± 0.017	0.108 ± 0.005
	Precision ↑	0.580 ± 0.031	0.305	0.551 ± 0.058	0.706 ± 0.063	0.481 ± 0.056	0.538 ± 0.022
	Recall ↑	0.363 ± 0.056	0.140	0.385 ± 0.039	0.758 ± 0.041	0.529 ± 0.047	0.695 ± 0.036