

SonarSplat: Novel View Synthesis of Imaging Sonar via Gaussian Splatting

Advaith V. Sethuraman¹ Max Rucker¹ Onur Bagoren¹
 Pou-Chun Kung¹ Nibarkavi N.B. Amutha¹ Katherine A. Skinner¹

Abstract—In this paper, we present SonarSplat, a novel Gaussian splatting framework for imaging sonar that demonstrates realistic novel view synthesis and models acoustic streaking phenomena. Our method represents the scene as a set of 3D Gaussians with acoustic reflectance and saturation properties. We develop a novel method to efficiently rasterize Gaussians to produce a range/azimuth image that is faithful to the acoustic image formation model of imaging sonar. In particular, we develop a novel approach to model azimuth streaking in a Gaussian splatting framework. We evaluate SonarSplat using real-world datasets of sonar images collected from an underwater robotic platform in a controlled test tank and in a real-world river environment. Compared to the state-of-the-art, SonarSplat offers improved image synthesis capabilities (+3.2 dB PSNR) and more accurate 3D reconstruction (77% lower Chamfer Distance). We also demonstrate that SonarSplat can be leveraged for azimuth streak removal.

Index Terms—Mapping, Deep Learning for Visual Perception, Marine Robotics

I. INTRODUCTION

ACOUSTIC sensors, such as imaging sonar, are commonly used for infrastructure inspection, large-area mapping, and target detection in underwater environments [1], [2], [3]. While optical sensors are severely range-limited due to water column effects on light propagation, acoustic sensors can capture data at long ranges to provide critical information about subsea environments. However, acoustic phenomena like elevation ambiguity, azimuth streaking, and multi-path reflections make sonar interpretation difficult for both operators and computer vision algorithms.

Radiance field methods offer a promising approach to addressing these challenges. Neural radiance fields (NeRFs) have demonstrated the potential for high-fidelity data synthesis, denoising, and 3D reconstruction using optical imagery [4], [5], [6], [7]. The benefits of neural rendering have been recognized by the underwater perception community, with prior work exploring neural rendering for 3D object reconstruction using underwater cameras [8], [9] and sonar data [10], [11], [12], [13], [14]. While these innovations are promising, training and processing NeRFs is costly and time-consuming, making deploying them in real-time or on resource-constrained devices

Manuscript received: May 4, 2025; Revised August 10, 2025; Accepted October 11, 2025.

This paper was recommended for publication by Editor Sven Behnke upon evaluation of the Associate Editor and Reviewers' comments.

This work was supported by the National Science Foundation under Award No. 2337774.

¹Authors are with the Department of Robotics, University of Michigan, Ann Arbor, MI, USA. {advaiths, mruck, obagoren, pckung, nibah, kskin}@umich.edu

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE

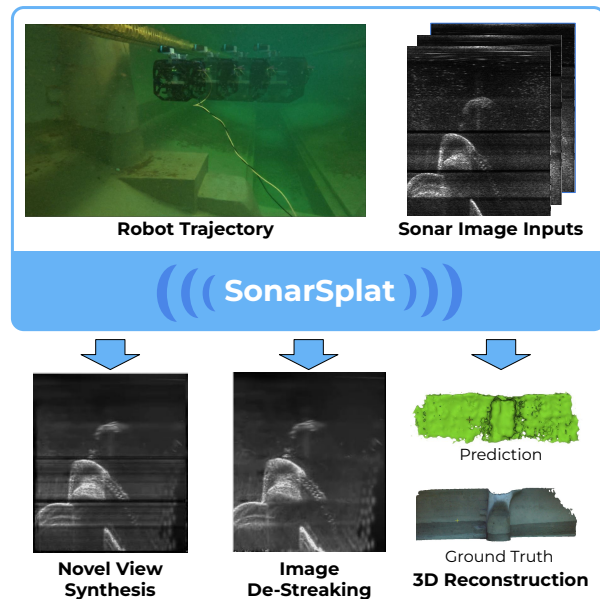


Fig. 1. We present SonarSplat, a novel 3D Gaussian splatting framework that enables novel view synthesis of realistic sonar images, image de-streaking, and 3D reconstruction. We model the acoustic reflectivity and azimuth streak probabilities for all Gaussians, yielding a 3D scene representation that can be queried for a variety of downstream tasks.

difficult. More recently, 3D Gaussian splatting (3DGS) [15] was developed as a faster alternative to NeRF. Gaussian splatting has been leveraged for underwater applications as well [16], [17], [18]. Most relevant to our proposed work, ZSplat proposes a Gaussian splatting framework for RGB-sonar fusion [19]. However, ZSplat leverages the fusion of sonar data to improve the rendering of RGB images and does not focus on enabling high-fidelity data synthesis for sonar imagery or providing evaluation for the quality of rendered sonar images or 3D reconstruction. Thus, there is a clear gap in the literature for a framework capable of efficient and effective sonar image synthesis and 3D reconstruction using Gaussian splatting.

In this paper, we propose SonarSplat (Fig. 1), a novel Gaussian splatting framework for imaging sonar that enables efficient and high-quality novel view synthesis, sonar image de-streaking, and 3D reconstruction for underwater robotic applications. To the best of our knowledge, SonarSplat is the first sonar-only method that leverages Gaussian splatting to perform these tasks. Our main contributions are as follows:

- We develop a novel 3D Gaussian splatting framework for rasterization of range/azimuth sonar images from known and arbitrary viewpoints.
- We develop a novel method for learning the probability of azimuth streaking in a differentiable manner, which al-

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

lows for removal of unwanted azimuth streaking artifacts from rendered images.

- We introduce a novel densification strategy called the Elevation Sampling Densification Strategy (ESDS) that places Gaussians on elevation arcs and yields better novel view synthesis and 3D reconstruction.
- We perform quantitative and qualitative evaluation of our method on real imaging sonar data collected from a robotic platform with a variety of sonar ranges and environments.

For further information, please visit our project website at <https://umfieldrobotics.github.io/sonarsplat3D/>.

II. RELATED WORK

A. Radiance Field Methods

NeRFs represent scenes implicitly using neural networks and have been explored for a variety of robotic applications, including novel view synthesis [4], data generation [20], Simultaneous Localization and Mapping (SLAM) [21], [22], manipulation [23], and image dehazing [7]. Despite their success as a scene representation, a key challenge for the practical application of NeRFs is the long training time and the slow rendering speeds, primarily due to the computationally expensive volumetric rendering process [4]. A recent development in radiance field methods is Gaussian splatting, which uses more efficient *rasterization* instead of volumetric rendering to produce a scene representation. This has led to the use of radiance field methods in a variety of real-time applications [15], [24], [25].

Despite their impressive performance and rapid adoption, radiance field methods were developed with optical sensors in mind, limiting their direct application to non-visual sensors. Recent developments have integrated non-visual sensors, including radar, LIDAR, sonar, thermal cameras, and event cameras into NeRF and Gaussian splatting frameworks for various vision tasks [26], [27], [21], [28], [10], [29]. We are inspired by the faster rendering speeds and high fidelity of Gaussian splatting methods and propose a novel Gaussian splatting framework for imaging sonar.

B. Perception Challenges for Imaging Sonar

Imaging sonar is a popular underwater sensor for perception tasks like object detection, segmentation, reconstruction, and inspection [10], [30], [1], [31], [32]. However, the sensor is susceptible to multi-path effects, nonlinear noise patterns, and phenomena like azimuth streaking [33]. Azimuth streaking, in particular, can make sonar image understanding difficult by suppressing returns from specific ranges, causing black bars in the sonar image [33]. In addition, the sensor inherently loses information from 3D into 2D when capturing the acoustic returns, causing an effect known as *elevation angle ambiguity* [34]. This ambiguity has made the successful application of imaging sonar to 3D vision tasks challenging [32], [35], [36].

C. Novel View Synthesis for Imaging Sonar

Prior work has incorporated the sonar image formation model into signed distance field (SDF) and NeRF formulations [10], [13], [12]. Neusis presents a volumetric rendering framework for dense 3D reconstruction of objects using an imaging sonar [10]. Neusis-NGP is an extension of Neusis that leverages multi-resolution hash encodings to produce a bathymetric heightmap of the environment [12]. Differentiable Space Carving (DSC) proposes a faster method that models echo probabilities to carve out a mesh of the scene [13]. Through our experiments, we demonstrate that these methods render less realistic sonar imagery compared to SonarSplat according to quantitative metrics including PSNR, SSIM, and LPIPS. This limits their use in data generation and augmentation. Additionally, our experiments indicate that querying these methods is time consuming, which limits their use in real-time applications.

To the best of our knowledge, ZSplat is the only Gaussian splatting paper that incorporates imaging sonar data, with demonstrations on underwater data. However, ZSplat operates with a simplified sonar model that does not model several important sonar-specific phenomena. More importantly, ZSplat only splats the *opacities* of Gaussians and does not consider material properties like acoustic reflectance. This leads to poor sonar renderings, which limits its application in sonar novel view synthesis. We demonstrate that SonarSplat, which more accurately models sonar image formation, produces more realistic novel view synthesis, restores sonar images to remove azimuth streaking artifacts, and produces 3D reconstructions under a unified framework.

III. TECHNICAL APPROACH

Figure 2 shows an overview of SonarSplat, which uses sonar images and their corresponding sensor poses to optimize a scene representation. The scene is represented as a set of 3D Gaussian primitives, each parametrized with a mean, covariance, opacity, acoustic reflectivity, and azimuth streaking probability. First, the Gaussians that fall within the view frustum of the sonar are rasterized into a range/azimuth image. Concurrently, we render an azimuth streak probability image using the per-Gaussian azimuth streak probabilities. Based on the relative probabilities in each range bin, we use a novel adaptive gain module to adjust the acoustic returns and capture azimuth streaking effects. Finally, we produce a sonar image with azimuth streaks that can be optimized by taking a loss with respect to the ground truth sonar image. We provide further details of each of these components in the remainder of this section.

A. Imaging Sonar Formation Model

To adapt the optical Gaussian splatting framework for imaging sonar, we must incorporate the acoustic image formation model into the rasterization process. Imaging sonar is a time-of-flight sensor that provides the echo intensities at a given range and azimuth bin.

A sonar image is defined as $I \in \mathbb{R}^{N_a \times N_r}$ where N_a denotes the number of azimuth bins and N_r denotes the

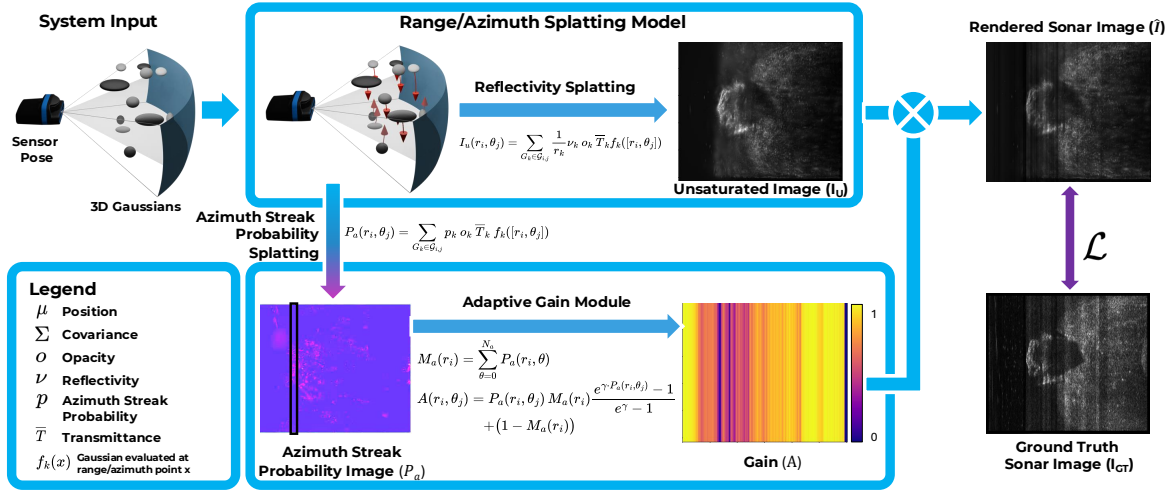


Fig. 2. Overview of SonarSplat. Our method takes as input a sensor pose and an initial set of 3D Gaussians representing the scene. Then, we transform the 3D Gaussians into the sensor’s image space. We splat the reflectivity parameter ν_k to get the unsaturated image I_u . Additionally, we optimize and splat per-Gaussian azimuth streaking probabilities p_k . All of the probabilities in a range bin r_i are considered in our novel Adaptive Gain Module, which adjusts the gain A applied to I_u . Finally, we produce \hat{I} by multiplying the gain A by I_u . All parameters are optimized using gradient descent by taking losses with respect to the sonar image pixel values. All sonar images shown are polar (range/azimuth) coordinates.

number of range bins. In spherical coordinates, the sonar has a vertical field of view (ϕ_{min}, ϕ_{max}), a horizontal field of view ($\theta_{min}, \theta_{max}$), and a maximum range, R_{max} . Finally, $\frac{R_{max}}{N_r}$ is the range resolution, ϵ_r , and $\frac{\theta_{max} - \theta_{min}}{N_a}$ is the azimuthal resolution, ϵ_a .

We leverage the sonar rendering equation introduced in [10], which calculates the intensity of return at bin (r_i, θ_j) in the sonar image I by:

$$I(r_i, \theta_j) = \int_{\phi_{min}}^{\phi_{max}} \int_{\theta_j - \epsilon_a}^{\theta_j + \epsilon_a} \int_{r_i - \epsilon_r}^{r_i + \epsilon_r} \frac{E_e}{r} T(r, \theta, \phi) \sigma(r, \theta, \phi) dr d\theta d\phi \quad (1)$$

where E_e is the intensity of emitted sound, r is the range of the point, $T(r, \theta, \phi)$ is the transmittance term, and $\sigma(r, \theta, \phi)$ is the density of the point. $T(r, \theta, \phi)$ is the probability that a ray travels to point (r, θ, ϕ) without hitting another particle.

B. Range/Azimuth Splatting Model

To evaluate this integral, prior methods use sampling and quadrature [10], [11], [13]. However, motivated by the increased efficiency of splatting techniques, we explore a point-based approach adapted to the sonar image formation model. To this end, we represent the scene with a set of 3D Gaussian primitives $\mathcal{G} = \{G_k\}_{k=1}^N$. Formally, these Gaussians are parametrized with a mean $\mu_k \in \mathbb{R}^3$, scale $s_k \in \mathbb{R}^3$, orientation $q_k \in \mathbb{H}$, opacity $o_k \in \mathbb{R}^+$, and azimuth streak probability $p_k \in \mathbb{R}^+$. We take inspiration from other Time-of-Flight novel view synthesis work and use spherical harmonics coefficients sh_k to model view-dependent acoustic reflectivity ν_k [28]. Specifically, $\nu_k = SH(sh_k, \vec{v})$, where \vec{v} is the viewing direction.

$$G_k = (\mu_k, s_k, q_k, o_k, sh_k, p_k) \quad (2)$$

where the covariance $\Sigma_k = R_k S_k S_k^T R_k^T$ for rotation matrix R_k given by quaternion q_k and scaling matrix S_k given by s_k .

To render sonar images, we first transform the means and covariances of the Gaussians into spherical coordinates. Let $\mu_k = [x_k, y_k, z_k]^T$ and Σ_k be the mean and covariance in the

sensor’s frame, respectively. The transformation to spherical coordinates is as follows [37]:

$$\mu_s = \begin{bmatrix} r_k \\ \theta_k \\ \phi_k \end{bmatrix} = \begin{bmatrix} \|\mu_k\| \\ \arctan(y_k, x_k) \\ \arctan(z_k, \sqrt{x_k^2 + y_k^2}) \end{bmatrix} \quad (3)$$

The covariance matrix Σ_k is transformed by using a first-order linearization of the coordinate transformation, which the Jacobian describes:

$$J_s = \begin{bmatrix} \frac{x_k}{\|\mu_k\|} & \frac{y_k}{\|\mu_k\|} & \frac{z_k}{\|\mu_k\|} \\ \frac{-y_k}{x_k^2 + y_k^2} & \frac{x_k}{x_k^2 + y_k^2} & 0 \\ \frac{-x_k \cdot z_k}{\|\mu_k\|^2 \sqrt{x_k^2 + y_k^2}} & \frac{-y_k \cdot z_k}{\|\mu_k\|^2 \sqrt{x_k^2 + y_k^2}} & \frac{\sqrt{x_k^2 + y_k^2}}{\|\mu_k\|^2} \end{bmatrix} \quad (4)$$

and the covariance matrix becomes

$$\Sigma_s = J_s \Sigma_k J_s^T \quad (5)$$

To convert the spherical coordinate mean and covariance into range/azimuth image space, we first specify the intrinsic matrix K of the sonar,

$$K = \begin{bmatrix} \frac{1}{\epsilon_r} & 0 & 0 \\ 0 & \frac{1}{\epsilon_a} & \frac{N_a}{2} \\ 0 & 0 & 1 \end{bmatrix} \quad (6)$$

and then we perform the transformation:

$$\mu'_k = K \mu_s, \quad \Sigma'_k = J_K W \Sigma_s W J_K^T \quad (7)$$

where μ'_k and Σ'_k are the mean and covariance transformed to range/azimuth space, J_K is the Jacobian of K , and W is the sensor’s view matrix, similar to the camera splatting formulation [15].

To find the intensity of acoustic returns at a given range/azimuth bin $[r_i, \theta_j]$, we consider $\mathcal{G}_{i,j} \subset \mathcal{G}$, the set of Gaussians transformed to range/azimuth image space that overlap with the range/azimuth bin $[r_i, \theta_j]$. The rasterization equation to obtain the sonar image for SonarSplat is similar

to that of [15], with a few key differences. Specifically, we use single-channel acoustic reflectance ν instead of per-channel color c and we form the image by splatting in the range/azimuth plane using $[r, \theta]$ coordinates instead of traditional camera coordinates. SonarSplat's rasterization equation is given by:

$$I_u(r_i, \theta_j) = \sum_{G_k \in \mathcal{G}_{i,j}} \frac{1}{r_k} \nu_k o_k \bar{T}_k f_k([r_i, \theta_j]) \quad (8)$$

$$f_k(x) = \exp\left(-\frac{1}{2} \|x - \mu'_k\|_{\Sigma'_k}^2\right) \quad (9)$$

$$\bar{T}_k = \frac{1}{|\mathcal{P}_k|} \sum_{p \in \mathcal{P}_k} T_k^p \quad (10)$$

where r_k is the range of G_k from the sensor, $f_k(x)$ is the evaluation of Gaussian G_k at range/azimuth point $[r_i, \theta_j]$ with the norm as the Mahalanobis distance, \bar{T}_k is the average transmittance of the Gaussian from the sensor origin to the Gaussian along the range. To compute \bar{T}_k , 3D Gaussians are splatted onto the azimuth/elevation plane to model occlusions in the range dimension. We define \mathcal{P}_k as the set of azimuth/elevation pixels that the Gaussian G_k overlaps and we calculate T_k^p as the transmittance of G_k at pixel p according to [15]. The subscript indicates that I_u is the *unsaturated* image, meaning that it does not account for the azimuth streaking that causes saturation across the image. Our sonar images are gain adjusted, so $\frac{1}{r_k}$ is dropped.

C. Azimuth Streak Modeling (ASM)

Azimuth streaking, as shown in Fig. 4, is a phenomena that is frequently observed in sonar images. It is a form of saturation that occurs when the sonar receiver receives strong returns from incident angles close to 0 at a specific range bin r_i [33]. In our model, we assign each Gaussian a probability that it will cause an azimuth streak. The benefit of defining this probability per-Gaussian is that it can be optimized to produce azimuth streaks that are multi-view consistent. Let p_k be the probability that Gaussian k will contribute to an azimuth streaking phenomena. We can then splat these per-Gaussian probabilities into a range/azimuth bin to obtain an azimuth streak probability image P_a :

$$P_a(r_i, \theta_j) = \sum_{G_k \in \mathcal{G}_{i,j}} p_k o_k \bar{T}_k f_k([r_i, \theta_j]) \quad (11)$$

$$M_a(r_i) = \sum_{\theta=0}^{N_a} P_a(r_i, \theta) \quad (12)$$

where $M_a(r_i)$ represents the probability of an azimuth streak occurring at range interval r_i . The azimuth streak probabilities across the range bins are then used to compute the final image using an adaptive gain mechanism.




Fig. 3. Illustration of the proposed Elevation Sampling Densification Strategy (ESDS). We randomly sample N_p pixels (shown as green sectors) based on the magnitude of \mathcal{L} then randomly place N_g Gaussians (shown as green ellipses) on the elevation arc of that pixel.

1) *Adaptive Gain Module*: We introduce a novel adaptive gain term, A , that is produced by transforming the azimuth streak probability image. We begin designing this gain term with a few observations. First, when no Gaussian in a range bin r_i has a high probability of azimuth streaking, the gain should be unity. Second, suppose a single Gaussian G_k exhibits a high probability of azimuth streaking across the range bin r_i . In that case, we wish to adaptively *suppress* the other Gaussians in r_i and assign a high gain to G_k . Finally, if multiple Gaussians in r_i have high probabilities of contributing to an azimuth streak, higher gain should be assigned to those Gaussians. Following these observations, the adaptive gain term is generated for each pixel: (r_i, θ_j) :

$$A(r_i, \theta_j) = P_a(r_i, \theta_j) M_a(r_i) \frac{e^{\gamma \cdot P_a(r_i, \theta_j)} - 1}{e^\gamma - 1} + (1 - M_a(r_i)) \quad (13)$$

where γ is a scaling factor that dictates the steepness of the adaptive gain. This behavior aligns with the insight presented in [33]: if a certain percentage of returns exceeds a threshold, an operation is applied to all of the values in a range interval. However, our proposed model differs because we do not restrict the gain to a fixed function (quadratic) but rather offer a family of curves for our splatting model to explore during optimization. The final sonar image $\hat{I}(r_i, \theta_j)$ with azimuth streaks is then computed by:

$$\hat{I}(r_i, \theta_j) = A(r_i, \theta_j) \cdot I_u(r_i, \theta_j) \quad (14)$$

D. Model Optimization

We optimize SonarSplat using gradient descent and by taking losses $(\mathcal{L}_{l1}, \mathcal{L}_{ssim})$ between \hat{I} and the ground truth image I_{GT} [15]. Next, we note that large Gaussians cause increased training times and more distortion to the model presented in Sec. III-B. To prevent Gaussians from becoming too large, we use $\mathcal{L}_{size} = \text{ReLU}(s - s_{init})$, which encourages the scale of the Gaussians s to remain smaller than their initialized scale s_{init} [38]. Additionally, we find it useful for reconstruction tasks to supervise the opacities of Gaussians via an opacity loss \mathcal{L}_o , similar to [19]. We compute \mathcal{L}_o between the predicted opacity map \hat{I}_o and I_{GT} thresholded at τ_o to isolate strong returns. The threshold parameter is found experimentally for each dataset. \mathcal{L}_o then becomes an $l1$ loss between \hat{I}_o and I_{GT}^o . The final loss becomes:

$$\mathcal{L} = \lambda_{l1} \mathcal{L}_{l1} + (1 - \lambda_{l1}) \mathcal{L}_{ssim} + \lambda_o \mathcal{L}_o + \lambda_{size} \mathcal{L}_{size} \quad (15)$$

where λ_{l1} , λ_o , and λ_{size} are the weights for the $l1$, opacity, and size losses respectively.




Fig. 4. We present our robotic platform and selected images from our diverse set of 9 sequences used in our experiments. We show RGB images in top right corners for visualization. Our evaluation datasets focus on smaller objects and larger-scale structures such as the test-tank basin. Azimuth streaks clearly occur in the *infra_360* and *piling_1* sequences.

E. Elevation Sampling Densification Strategy (ESDS)

Splitting and cloning Gaussians in image space is unintuitive for imaging sonar given the elevation angle ambiguity. It is possible that a Gaussian can be split then immediately occluded by high opacity Gaussians in front of it. To address this, we introduce an alternative strategy for densification of Gaussians called the Elevation Sampling Densification Strategy (ESDS) shown in Fig. 3. Given predicted image \hat{I} and ground truth I_{GT} , we sample N_p pixels, where each pixel’s probability is proportional to the magnitude of the loss at that pixel. We initialize N_g Gaussians on each elevation arc specified by the sampled pixel locations. This primarily serves to provide SonarSplat enough opportunities to resolve the elevation angle ambiguity and any possible occlusions by placing Gaussians at different elevation angles. Low-opacity Gaussians are pruned similar to [15].

F. Implementation Details

SonarSplat is built on the *gsplat* repo [39]. We train our method on an NVIDIA RTX 3090 GPU with 24 GB of vRAM. We calculate the frames per second (FPS) of all methods on the same NVIDIA A6000 GPU with 50 GB of vRAM. For every pixel in a sonar image with intensity above a threshold τ_{init} , we initialize N_{init} Gaussians on the elevation arc specified by the pixel’s range, azimuth, and the sensor’s vertical field of view. This process is repeated for all sonar images in a dataset to produce the initial set of 3D Gaussians. To properly optimize for the azimuth streaking probabilities p_k of each Gaussian, we first identify where azimuth streaks occur by calculating the average intensity of the range interval r_i . Then, for the first N_s iterations, we train only on pixels in range bins that exceed a certain average intensity threshold τ_s . This avoids incorrectly supervising the splat in azimuth streaked regions. Importantly, p_k is not optimized during this interval. After N_s iterations, we optimize only p_k in order to capture the azimuth streaking effects. We do not optimize μ_k , Σ_k , o_k , or ν_k during this interval. This allows the Gaussians to fit the images before optimizing the azimuth streaking parameters.

IV. RESULTS & EXPERIMENTS

A. Robotic Platform and Datasets

We note that real, public datasets of posed sonar images suitable for novel view synthesis and reconstruction experiments

TABLE I

NOVEL VIEW SYNTHESIS PERFORMANCE ON REAL WORLD HELD-OUT VALIDATION DATA (7 DATASETS). NOTE THAT ZSPAT* IS ZSPAT WITH ONLY THE SONAR LOSS ACTIVATED FOR TRAINING. FIRST, SECOND, AND THIRD BEST ARE SHOWN IN RED, ORANGE, AND YELLOW.

Method	Sonar Image Synthesis			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FPS \uparrow
Neusis [10]	14.68	0.04	0.64	0.03
Neusis-NGP [12]	19.56	0.33	0.64	0.01
DSC [13]	14.70	0.07	0.66	0.03
ZSplat* [19]	10.06	0.02	0.69	155.32
SonarSplat (ours)	22.79	0.41	0.48	13.45

do not exist. Rather, we collect sequences using a BlueRobotics BlueROV2 platform, shown in Fig. 4. Our robotic platform is equipped with a Blueprint Subsea Oculus M750-d imaging sonar operating at 1.2MHz, capturing sequences between 3-20m in range, with a 130° horizontal field of view and a 20° vertical field of view. The Oculus M750-d has a range resolution varying from 0.5-4 cm/pixel and angular resolution of 0.6° with a beam separation of 0.25° [40]. We tilt the sonar both at 30° and 0° downwards for various sequences. We use this platform to collect 9 trajectories inspecting a range of objects and in both a test tank and river environment. We use 7 sequences for evaluating image synthesis, and 2 sequences with higher coverage of the inspected object for evaluating 3D reconstruction. Our 3D reconstruction datasets are called Concrete Piling and *Monohansett Shipwreck*. Sensor poses are estimated by filtering a Doppler Velocity Logger and IMU, similar to [41], [10]. For Concrete Piling, we construct ground truth for 3D reconstruction using RGB images and MAST3R, scaling the reconstruction to metric scale [42]. For *Monohansett Shipwreck*, we use an underwater LiDAR scan of the shipwreck as ground truth.

B. Baselines

For baselines, we compare to Neusis [10], Neusis-NGP [12], DSC [13], and ZSplat* [19]. Since we are focused on sonar-only rendering, we compare to a modified version, ZSplat*, trained with only the sonar loss since there is no other work that has explored Gaussian splatting for imaging sonar [19]. All baselines were trained using official code and hyperparameters were tuned to the best of our ability.

C. Novel View Synthesis

We report validation PSNR, SSIM, and LPIPS following [15], using the *lpipsPyTorch* library. For our validation




Fig. 5. We present qualitative novel view synthesis results from selected datasets and selected baselines. Note that SonarSplat produces more realistic sonar image synthesis compared to baselines and is better able to capture finer details in the environment. Note that ZSplat* indicates ZSplat trained only with the sonar loss. Images are cartesian.

set, we take every 8th image from the dataset as introduced in [4]. We report SonarSplat’s performance compared to baselines in Tab. I. We find that SonarSplat consistently synthesizes more realistic sonar images, quantified by its higher PSNR and SSIM, and lower LPIPS values. Qualitative results from selected sequences are presented in Fig. 5. We also report the rendering speed in FPS of each method in Tab. I. SonarSplat and ZSplat* both leverage Gaussian splatting, which offers significantly faster rendering speeds compared to neural rendering methods. ZSplat* is optimized for fast rendering, which will be a focus of future work.

D. Ablation Studies

We perform ablation studies on the view-dependent reflectivity parameter ν_k (VDR), the azimuth streak modeling (ASM), and the Elevation Sampling Densification Strategy (ESDS). Ablation studies are averaged across 7 evaluation datasets and trained with exactly the same hyperparameters as results from Sec. IV-C. Methods that don’t use ESDS use the mcmc strategy in `gsplat` [39]. Results are reported in Tab. II and shown in Fig. 6. Clearly, the addition of azimuth streak modeling, view-dependent acoustic reflectivity, and the elevation sampling densification strategy improves the image synthesis capabilities of SonarSplat.

E. Azimuth Streak Removal/De-Streaking

Using SonarSplat, we can undo the adaptive gain to render images without azimuth streaks present and recover suppressed returns. This vision task is most similar to *image de-stripping* from the remote sensing literature. We present quantitative results in Tab. III using the inverse coefficient of variation (ICV) metric, which is a common no-reference metric used

TABLE II
ABLATION STUDY RESULTS FOR THE PROPOSED ACOUSTIC REFLECTANCE PARAMETER (VDR), AZIMUTH STREAK MODELING (ASM), AND ELEVATION SAMPLING DENSIFICATION STRATEGY (ESDS). BEST IN BOLD.

SonarSplat Rendering Ablations					
VDR	ASM	ESDS	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
\times	\times	\times	21.32	0.35	0.53
\checkmark	\times	\times	21.16	0.34	0.55
\checkmark	\checkmark	\times	22.51	0.40	0.52
\checkmark	\checkmark	\checkmark	22.79	0.41	0.48




Fig. 6. Qualitative comparison of SonarSplat without and with azimuth streak modeling (ASM) for `piling_1`. Note that if ASM is not modeled, rendered views will produce erroneous streaks (indicated with red arrows) instead of capturing the true streaks (indicated with blue arrows).




Fig. 7. Qualitative results for azimuth streak removal capabilities of SonarSplat. Locations of azimuth streaks in the captured images are indicated by blue arrows.

TABLE III
WE REPORT THE INVERSE COEFFICIENT OF VARIATION (ICV) FOR AZIMUTH STREAK REMOVAL. BEST IS IN BOLD.

De-Streaking Performance (ICV) \uparrow		
	Captured Image	De-Streaked Prediction
Average	2.91	18.58

in de-stripping works [43], [44], [45]. We present qualitative results of de-streaked images in Fig. 7. The Azimuth Streak Modeling allows SonarSplat to qualitatively and quantitatively remove streaks effectively and improves the finer features in sonar images.

F. Real-World 3D Reconstruction

3D reconstruction is crucial for inspection and visualization of underwater structures. In a test tank, we demonstrate inspection of Concrete Piling, a structure commonly found in bridges and dams. We also surveyed *Monohansett*




Fig. 8. Qualitative results from 3D reconstruction of SonarSplat compared to baselines. We inspect *Concrete Piling* in a test tank. *Monohansett Shipwreck* was surveyed in Lake Huron, Michigan. Ground truth is created using collected RGB images and LIDAR respectively. Readers are encouraged to zoom in.

Shipwreck, a wooden shipwreck at a depth of 18 feet in Lake Huron, Michigan, measuring 160 feet in length and 30 feet in width.

We show that SonarSplat learns the geometric properties of the scene by converting the optimized splat into a mesh for visualization. We sample a set of points from the probability density functions described by the Gaussians, then query the opacities and use marching cubes to produce a mesh. Specifically, we used between 10-50 samples per Gaussian depending on the size of the scene and the density of the Gaussians. We compare to state-of-the-art sonar reconstruction methods: Neusis, DSC, and ZSplat* [13], [10], [19]. Additionally, we ablate ESDS (denoted as “SonarSplat w/o ESDS”). Qualitative results are shown in Fig. 8 and quantitative results are shown in Tab. IV. Following [10], [13], we report Root Mean Squared (RMS) l_1 Chamfer Distance (CD- l_1) and Hausdorff Distance (HD), which are both in meters. We produce meshes from all methods then randomly sample 30,000 points from the surface for 30 iterations to calculate RMS. We use ICP to align the predictions and ground truth meshes then crop the predictions to the bounding box of the ground truth. Note that neither DSC, Neusis, nor ZSplat* is able to reconstruct the object whereas SonarSplat is able to produce an accurate reconstruction of the piling. For *Monohansett Shipwreck*, we see that SonarSplat captures details like the wooden ribs of the shipwreck more accurately than baselines.

TABLE IV

QUANTITATIVE EVALUATION OF 3D RECONSTRUCTION USING RMS l_1 CHAMFER DISTANCE (CD- l_1) AND HAUSDORFF DISTANCE (HD). UNITS ARE METERS. BEST IN BOLD.

3D Reconstruction Performance				
Method	Concrete Piling		Monohansett Shipwreck	
	CD- l_1 ↓	HD ↓	CD- l_1 ↓	HD ↓
Neusis [10]	0.32	0.78	1.75	3.67
DSC [13]	0.29	0.82	1.60	6.77
ZSplat* [19]	2.93	4.88	3.06	8.30
SonarSplat w/o ESDS	0.92	1.77	3.62	12.52
SonarSplat (ours)	0.14	0.57	0.29	1.06

V. CONCLUSION & FUTURE WORK

We propose SonarSplat, the first sonar-only Gaussian splatting framework for imaging sonar in underwater applications. First, we adapt the sonar rendering equation for efficient range/azimuth splatting by evaluating 3D Gaussians. We model azimuth streaking within the Gaussian splatting framework, allowing us to learn per-Gaussian azimuth streaking probabilities and produce high-quality de-streaked images. We also introduce a novel densification strategy (ESDS) to enable

SonarSplat to resolve elevation angle ambiguities and associated occlusions. From our extensive experiments on a diverse set of real-world data in test tank and lake environments, we find that SonarSplat outperforms state-of-the-art methods in novel view synthesis (+3.2 dB PSNR) and 3D reconstruction (77% lower Chamfer Distance), and additionally allows for de-streaking of sonar images.

Future work will focus on leveraging SonarSplat for sonar data synthesis by randomizing scene parameters including acoustic reflectance and azimuth streaking probability. This can yield more diverse datasets for training sonar-based scene understanding algorithms, including multi-view target recognition algorithms [46], [47]. We will also work to improve SonarSplat’s ability to perform dense 3D reconstruction by better modeling additional acoustic phenomena like speckle noise and multi-path reflections. Our choice of Gaussian splatting as a 3D scene representation also opens possibilities for real-time sensor fusion and SLAM [24].

REFERENCES

- [1] H. Johannsson, M. Kaess, B. Englot, F. Hover, and J. Leonard, “Imaging sonar-aided navigation for autonomous underwater harbor surveillance,” in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 4396–4403. 1, 2
- [2] A. V. Sethuraman, A. Sheppard, O. Bagoren, C. Pinnow, J. Anderson, T. C. Havens, and K. A. Skinner, “Machine learning for shipwreck segmentation from side scan sonar imagery: Dataset and benchmark,” *The International Journal of Robotics Research*, vol. 44, no. 3, pp. 341–354, 2024. [Online]. Available: <https://doi.org/10.1177/02783649241266853> 1
- [3] J. McConnell, J. D. Martin, and B. Englot, “Fusing concurrent orthogonal wide-aperture sonar images for dense underwater 3D reconstruction,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 1653–1660. 1
- [4] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis,” in *2020 European Conference on Computer Vision (ECCV)*, ser. Lecture Notes in Computer Science, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, pp. 405–421. 1, 2, 6
- [5] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang, “Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction,” in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 27 171–27 183. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2021/file/e41e164f7485ec4a28741a2d0ea41c74-Paper.pdf 1
- [6] N. Pearl, T. Treibitz, and S. Korman, “NAN: Noise-Aware NeRFs for Burst-Denoising,” in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, June 2022, pp. 12 662–12 671. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/CVPR52688.2022.01234> 1
- [7] W.-T. Chen, W. Yifan, S.-Y. Kuo, and G. Wetzstein, “Dehazenerf: Multi-image haze removal and 3D shape reconstruction using neural radiance fields,” in *2024 International Conference on 3D Vision (3DV)*, 2024, pp. 247–256. 1, 2

- [8] D. Levy, A. Peleg, N. Pearl, D. Rosenbaum, D. Akkaynak, S. Korman, and T. Treibitz, "SeaThru-NeRF: Neural Radiance Fields in Scattering Media," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 56–65. **1**
- [9] A. V. Sethuraman, M. S. Ramanagopal, and K. A. Skinner, "WaterNeRF: Neural Radiance Fields for Underwater Scenes," in *OCEANS 2023 - MTS/IEEE U.S. Gulf Coast*, Sept. 2023, pp. 1–7. **1**
- [10] M. Qadri, M. Kaess, and I. Gkioulekas, "Neural implicit surface reconstruction using imaging sonar," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1040–1047. **1, 2, 3, 5, 7**
- [11] M. Qadri, K. Zhang, A. Hinduja, M. Kaess, A. Pediredla, and C. A. Metzler, "Aoneus: A neural rendering framework for acoustic-optical sensor fusion," in *ACM SIGGRAPH 2024 Conference Papers*, ser. SIGGRAPH '24. New York, NY, USA: Association for Computing Machinery, 2024. [Online]. Available: <https://doi.org/10.1145/3641519.3657446> **1, 3**
- [12] Y. Xie, G. Troni, N. Bore, and J. Folkesson, "Bathymetric surveying with imaging sonar using neural volume rendering," *IEEE Robotics and Automation Letters*, vol. 9, no. 9, pp. 8146–8153, 2024. **1, 2, 5**
- [13] Y. Feng, W. Lu, H. Gao, B. Nie, K. Lin, and L. Hu, "Differentiable space carving for 3D reconstruction using imaging sonar," *IEEE Robotics and Automation Letters*, vol. 9, no. 11, pp. 10065–10072, 2024. **1, 2, 3, 5, 7**
- [14] Y. Xie, N. Bore, and J. Folkesson, "Sidescan only neural bathymetry from large-scale survey," *Sensors*, vol. 22, no. 14, 2022. [Online]. Available: <https://www.mdpi.com/1424-8220/22/14/5092> **1**
- [15] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3D gaussian splatting for real-time radiance field rendering," *ACM Transactions on Graphics*, vol. 42, no. 4, July 2023. [Online]. Available: <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/> **1, 2, 3, 4, 5**
- [16] D. Yang, J. J. Leonard, and Y. Girdhar, "Seasplat: Representing underwater scenes with 3d gaussian splatting and a physically grounded image formation model," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025. **1**
- [17] H. Li, W. Song, T. Xu, A. Elsig, and J. Kulhanek, "Watersplating: Fast underwater 3D scene reconstruction using gaussian splatting," in *2025 International Conference on 3D Vision (3DV)*, 2025. **1**
- [18] N. Mualem, R. Amoyal, O. Freifeld, and D. Akkaynak, "Gaussian splashing: Direct volumetric rendering underwater," 2024, *arXiv:2411.19588*. **1**
- [19] Z. Qu, O. Vengurlekar, M. Qadri, K. Zhang, M. Kaess, C. Metzler, S. Jayasuriya, and A. Pediredla, "Z-splat: Z-axis gaussian splatting for camera-sonar fusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–12, 2024. **1, 4, 5, 7**
- [20] Y. Ge, H. Behl, J. Xu, S. Gunasekar, N. Joshi, Y. Song, X. Wang, L. Itti, and V. Vineet, "Neural-sim: Learning to generate training data with nerf," in *Computer Vision – ECCV 2022*, 2022, pp. 477–493. **2**
- [21] S. Isaacson, P.-C. Kung, M. Ramanagopal, R. Vasudevan, and K. A. Skinner, "Loner: Lidar only neural representations for real-time slam," *IEEE Robotics and Automation Letters*, vol. 8, no. 12, pp. 8042–8049, 2023. **2**
- [22] A. Rosinol, J. J. Leonard, and L. Carlone, "Nerf-slam: Real-time dense monocular slam with neural radiance fields," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 3437–3444. **2**
- [23] A. Simeonov, Y. Du, A. Tagliasacchi, J. B. Tenenbaum, A. Rodriguez, P. Agrawal, and V. Sitzmann, "Neural descriptor fields: Se(3)-equivariant object representations for manipulation," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 6394–6400. **2**
- [24] H. Matsuki, R. Murai, P. H. Kelly, and A. J. Davison, "Gaussian splatting slam," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024, pp. 18 039–18 048. **2, 7**
- [25] J. Luiten, G. Kopanas, B. Leibe, and D. Ramanan, "Dynamic 3D Gaussians: Tracking by Persistent Dynamic View Synthesis," in *2024 International Conference on 3D Vision (3DV)*. Los Alamitos, CA, USA: IEEE Computer Society, Mar. 2024, pp. 800–809. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/3DV62453.2024.00044> **2**
- [26] D. Borts, E. Liang, T. Broedermann, A. Ramazzina, S. Walz, E. Palladin, J. Sun, D. Brueggemann, C. Sakaridis, L. Van Gool, M. Bijelic, and F. Heide, "Radar fields: Frequency-space neural scene representations for fmcw radar," in *ACM SIGGRAPH 2024 Conference Papers*, ser. SIGGRAPH '24. New York, NY, USA: Association for Computing Machinery, 2024. [Online]. Available: <https://doi.org/10.1145/3641519.3657510> **2**
- [27] A. Carlson, M. S. Ramanagopal, N. Tseng, M. Johnson-Roberson, R. Vasudevan, and K. A. Skinner, "Cloner: Camera-lidar fusion for occupancy grid-aided neural representations," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2812–2819, 2023. **2**
- [28] T. Huang, J. Miller, A. Prabhakara, T. Jin, T. Laroia, Z. Kolter, and A. Rowe, "Dart: Implicit doppler tomography for radar novel view synthesis," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024, pp. 24 118–24 129. **2, 3**
- [29] Y. Y. Lin, X.-Y. Pan, S. Fridovich-Keil, and G. Wetzstein, "ThermalNeRF: Thermal Radiance Fields," in *2024 IEEE International Conference on Computational Photography (ICCP)*. Los Alamitos, CA, USA: IEEE Computer Society, July 2024, pp. 1–12. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/ICCP61108.2024.10644336> **2**
- [30] T. Guerneve, K. Subr, and Y. Petillot, "Three-dimensional reconstruction of underwater objects using wide-aperture imaging sonar," *Journal of Field Robotics*, vol. 35, no. 6, pp. 890–905, 2018. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.21783> **2**
- [31] M. Valdenegro-Toro, "Submerged marine debris detection with autonomous underwater vehicles," in *2016 International Conference on Robotics and Automation for Humanitarian Applications (RAHA)*. IEEE, 2016, pp. 1–7. **2**
- [32] S. Negahdaripour, "On 3-D reconstruction from stereo FS sonar imaging," in *OCEANS 2010 MTS/IEEE SEATTLE*, 2010, pp. 1–6. **2**
- [33] E. Potokar, K. Lay, K. Norman, D. Benham, T. B. Neilsen, M. Kaess, and J. G. Mangelson, "Holocean: Realistic sonar simulation," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 8450–8456. **2, 4**
- [34] Y. Wang, C. Wu, Y. Ji, H. Tsuchiya, H. Asama, and A. Yamashita, "2D forward looking sonar simulation with ground echo modeling," in *2023 20th International Conference on Ubiquitous Robots (UR)*, 2023, pp. 724–729. **2**
- [35] T. A. Huang and M. Kaess, "Towards acoustic structure from motion for imaging sonar," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 758–765. **2**
- [36] E. Westman, I. Gkioulekas, and M. Kaess, "A volumetric albedo framework for 3D imaging sonar reconstruction," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 9645–9651. **2**
- [37] P.-C. Kung, X. Zhang, K. A. Skinner, and N. Jaipuria, "LiHigs: Lidar-supervised gaussian splatting for highway driving scene reconstruction," 2024, *arXiv:2412.15447*. [Online]. Available: <https://arxiv.org/abs/2412.15447> **3**
- [38] J. W. Lee, H. Lim, S. Yang, and J. Choi, "Micro-splatting: Maximizing isotropic constraints for refined optimization in 3d gaussian splatting," 2025. [Online]. Available: <https://arxiv.org/abs/2504.05740> **4**
- [39] V. Ye, R. Li, J. Kerr, M. Turkulainen, B. Yi, Z. Pan, O. Seiskari, J. Ye, J. Hu, M. Tancik, and A. Kanazawa, "gsplat: An open-source library for Gaussian splatting," *arXiv preprint arXiv:2409.06765*, 2024. [Online]. Available: <https://arxiv.org/abs/2409.06765> **5, 6**
- [40] "Blueprint Subsea | Oculus M-Series." [Online]. Available: <https://www.blueprintsubsea.com/> **5**
- [41] "DVL A50." [Online]. Available: <https://waterlinked.com/shop/dvl-a50-114> **5**
- [42] V. Leroy, Y. Cabon, and J. Revaud, "Grounding image matching in 3d with mast3r," 2024. [Online]. Available: <https://arxiv.org/abs/2406.09756> **5**
- [43] H. Shen and L. Zhang, "A map-based algorithm for destriping and inpainting of remotely sensed images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 5, pp. 1492–1502, 2008. **6**
- [44] P. Rakwatin, W. Takeuchi, and Y. Yasuoka, "Stripe noise reduction in modis data by combining histogram matching with facet filter," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 6, pp. 1844–1856, 2007. **6**
- [45] J. E. Nichol and V. Vohora, "Noise over water surfaces in landsat tm images," *International Journal of Remote Sensing*, vol. 25, no. 11, pp. 2087–2093, 2004. **6**
- [46] A. V. Sethuraman and K. A. Skinner, "Stars: Zero-shot sim-to-real transfer for segmentation of shipwrecks in sonar imagery," in *34th British Machine Vision Conference 2023, BMVC 2023, Aberdeen, UK, November 20-24, 2023*. BMVA, 2023. [Online]. Available: <https://papers.bmvc2023.org/0606.pdf> **7**
- [47] A. V. Sethuraman, P. Baldoni, K. A. Skinner, and J. McMahon, "Learning which side to scan: Multi-view informed active perception with side scan sonar for autonomous underwater vehicles," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 8348–8354. **7**