NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

SENSOR PERFORMANCE ANALYSIS FOR MINE DETECTION WITH UNMANNED VEHICLES IN VERY SHALLOW WATER AND SURF ZONES

by

Alexander J. Fedorovich

June 2022

Thesis Advisor: Sean P. Kragelund
Co-Advisor: Joseph T. Klamo

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			*Form Approved OMB No. 0704-0188*

1. AGENCY USE ONLY *(Leave blank)*	**2. REPORT DATE** June 2022	**3. REPORT TYPE AND DATES COVERED** Master's thesis

4. TITLE AND SUBTITLE
SENSOR PERFORMANCE ANALYSIS FOR MINE DETECTION WITH UNMANNED VEHICLES IN VERY SHALLOW WATER AND SURF ZONES

5. FUNDING NUMBERS

6. AUTHOR(S) Alexander J. Fedorovich

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Postgraduate School
Monterey, CA 93943-5000

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
N/A

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

12a. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release. Distribution is unlimited.

12b. DISTRIBUTION CODE
A

13. ABSTRACT (maximum 200 words)

The very shallow water and surf zones present extraordinary challenges for classifying submerged objects such as mines or shoals. Accessing these areas with traditional unmanned underwater vehicles is difficult, and remotely operated vehicles often require putting operators in harm's way. This research explores the potential to perform object classification using only forward-looking sonar in the desired operating zones. Experiments were conducted in a controlled environment for two different target objects, a glass sphere and a rectangular cinder block. Next, forward-looking sonar images were analyzed to determine how the intensity and distribution of target returns changed as a function of distance and angle from the sonar. The ability to correlate experimentally measured intensity profiles with a target's physical size and shape is examined. Finally, recommendations for future research are proposed to further develop this approach for potential naval applications like mine countermeasures.

14. SUBJECT TERMS
mine, VSW, SZ, BZ, very shallow water, surf zone, beach zone, detection, sensing, standoff, unmanned, MCM, countermeasures, autonomous, seabed, warfare, UUV

15. NUMBER OF PAGES
87

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	**18. SECURITY CLASSIFICATION OF THIS PAGE**	**19. SECURITY CLASSIFICATION OF ABSTRACT**	**20. LIMITATION OF ABSTRACT**
Unclassified	Unclassified	Unclassified	UU

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

SENSOR PERFORMANCE ANALYSIS FOR MINE DETECTION WITH UNMANNED VEHICLES IN VERY SHALLOW WATER AND SURF ZONES

Alexander J. Fedorovich
Lieutenant, United States Navy
BS, United States Naval Academy, 2015

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN MECHANICAL ENGINEERING

from the

NAVAL POSTGRADUATE SCHOOL
June 2022

Approved by: Sean P. Kragelund
Advisor

Joseph T. Klamo
Co-Advisor

Garth V. Hobson
Chair, Department of Mechanical and Aerospace Engineering

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The very shallow water and surf zones present extraordinary challenges for classifying submerged objects such as mines or shoals. Accessing these areas with traditional unmanned underwater vehicles is difficult, and remotely operated vehicles often require putting operators in harm's way. This research explores the potential to perform object classification using only forward-looking sonar in the desired operating zones. Experiments were conducted in a controlled environment for two different target objects, a glass sphere and a rectangular cinder block. Next, forward-looking sonar images were analyzed to determine how the intensity and distribution of target returns changed as a function of distance and angle from the sonar. The ability to correlate experimentally measured intensity profiles with a target's physical size and shape is examined. Finally, recommendations for future research are proposed to further develop this approach for potential naval applications like mine countermeasures.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

List of Figures

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

List of Acronyms and Abbreviations

ATR Automatic Target Recognition

AUV Autonomous Underwater Vehicle

CAVR Center for Autonomous Vehicle Research

CNN Convolutional Neural Network

DOD Department of Defense

DOF Degrees of Freedom

EOD Explosive Ordnance Disposal

FLS Forward-Looking Sonar

FOV Field of View

GPS Global Positioning System

IMU Inertial Measurement Unit

INS Inertial Navigation System

MCM Mine Countermeasures

MIW Mine Warfare

NIWC Naval Information Warfare Command

NPS Naval Postgraduate School

RCNN Region-Based CNN

ROS Robot Operating System

ROV Remotely Operated Vehicle

SZ Surf Zone

USN U.S. Navy

VSW Very Shallow Water

UUV Unmanned Underwater Vehicle

CHAPTER 1:
Background and Motivation

Mine Warfare (MIW) is the strategic, operational, and tactical use of mines and mine countermeasures. Mines are typically emplaced to degrade the enemy's capabilities to wage land, air, and maritime warfare. MIW also involves countering enemy-emplaced mines to permit friendly maneuver or use of selected land or sea areas [1]. MIW can be traced back to the American Revolution, when a Yale student discovered gunpowder could be exploded underwater [2]. Mines have been featured in naval warfare ever since.

The use of mines offers an inexpensive and low-maintenance method of disrupting military operations. When placed in shallow water, mines can deny access to major waterways or ports. When placed in deep water, mines can cause the loss of expensive capital ships, such as destroyers and cruisers. The USS Samuel B. Roberts had a 15-foot hole blown into the port quarter and her keel broken when she struck a mine off the coast of Iran in April, 1988. A broken keel usually leads to sinking, but damage control efforts by the crew saved the ship. The ship was eventually heavy-lifted to a shipyard and repaired for a total cost of $90.8 million, or about $214 million in 2022 dollars. The mine the ship struck was estimated to cost less than $2,000 to manufacture and deploy [3]. Still relevant today, the USS Samuel B. Roberts example demonstrates the challenging and important nature of MIW.

The U.S. Navy (USN) process of defensive Mine Countermeasures (MCM) includes detection, classification, identification, and disposal [4]. Detection refers to the finding and localization of an object of interest. Classification involves determining if a detected target is either mine-like or benign. Possible attributes that could lead to an object's classification as mine-like include the presence of man-made features or having size and geometry that correlate to features of known mines. Identification requires visual confirmation by an operator and can be completed on-sight or with a camera. Disposal is excluded from the scope of this research.

Currently, the USN tactic to achieve detection of mines suspended in the water column is using the hull-mounted sonar of an Avenger-class MCM, which are small, wooden-hulled ships. To classify and identify mine-like objects, these platforms deploy an Explosive

Ordnance Disposal (EOD) diver or Remotely Operated Vehicle (ROV) system [5]. The Avenger-class is a legacy platform being phased out in favor of unmanned systems in combination with artificial intelligence and machine learning algorithms [6]. Unmanned MCM systems could be deployed at standoff ranges from small platforms, such as a littoral combat ship outfitted with a MCM mission module or a 10-meter RHIB from an amphibious warfare ship. Conducting MCM operations with unmanned systems significantly reduces the risk to human operators and increases the likelihood of operational mission success.

The USN must be able to perform MCM because missions such as amphibious assaults are highly vulnerable to adversary mines. These missions involve the movement of vehicles from the shallow water zone (12.2-61.0 m, 40-200 ft) of water depth through the Very Shallow Water (VSW) zone (3.28-12.2 m, 10-40 ft) and into the Surf Zone (SZ) (0-3.28 m, 0-10 ft) [7]. The SZ is not safely accessible by a submerged Unmanned Underwater Vehicle (UUV) at operating depth because of hazards like shoals, crashing waves, and sandbars. Crawling UUVs have the potential to successfully operate in the SZ, but are generally limited to only addressing mine-like objects on the bottom of the seafloor [8].

The VSW zone is shallow enough to force submerged UUVs to operate near the ocean surface. This causes them to experience wave-induced loads from passing waves that can result in unwanted vehicle motions. Also, the VSW zone sees wave breaking on the surface due to the decreasing water depth in this region. Because of this, the water column is usually well mixed with suspended sediment that reduces visibility. This region experiences a significant amount of underwater noise, making MCM operational planning within the VSW extraordinarily difficult. Finally, the shallow water zone is a transition region from Deep Water to the VSW and is usually orders of magnitude larger in volume than the VSW and SZ. These zones are depicted in Figure 1.1.

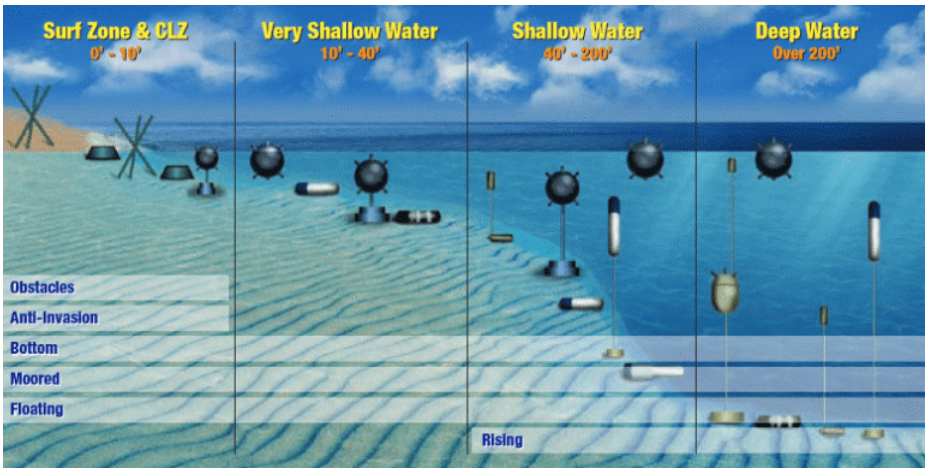

Figure 1.1. The zones of interest of this research.

Sending human operators into these zones is time-consuming, puts operator lives at risk, and reveals the operational force's intentions. As a result, missions like amphibious assaults must rely on brute force breaching techniques to destroy mines and barriers. As with other MCM tactics, breaching is difficult, inefficient, and puts lives at risk in addition to delaying the landing of vehicles [9].

Unlike an ROV, an Autonomous Underwater Vehicle (AUV) does not depend on a tether, which provides the advantage of increased stand-off range. Small AUVs are capable of operating in the zones that divers and ROVs cannot safely reach. These vehicles typically carry combinations of Global Positioning System (GPS), Inertial Navigation System (INS), camera, and sonar sensors. Unfortunately, mine detection in cluttered areas and/or the most shallow regions of the VSW is very challenging with side-looking sonars, because the environment does not allow vehicles to maintain or repeat the long, straight track-lines required by these sensors. Forward-looking cameras offer reliable, high-quality, real-time visual data when used in water with good clarity. The operational zones of interest in this research effort, however, typically contain poor water clarity. This amplifies the need for effective FLS utilization in the VSW. A FLS can be used to overcome these challenges because they deliver high-resolution, real-time imagery without specific maneuvering requirements. FLS excel with tasks such as object detection and avoidance, navigation, and hull inspections.

1.1 Previous Work

MCM research with AUVs has been explored for decades. In 1999, Smith et al. showed that AUV mine detection in water depths of 50-70 feet with calm conditions is possible using area search tactics and side-scan sonar [10]. That research also successfully used camera video to manually correlate data from the sonar. Topple and Fawcett attempted to apply MiNet, a deep learning model, for Automatic Target Recognition (ATR) with side scan sonars on AUVs [11]. Similarly, recent Naval Information Warfare Command (NIWC) research has implemented deep learning algorithms trained on a vast amount of historical data with successful object detection and classification trials [12]. However, that effort uses side-looking sonar on larger vehicles with precise navigation and maneuvering profiles. These vehicles and procedures cannot be utilized in the challenging environment of the VSW zone, which is the zone this research is focused on.

Hurtos et al. demonstrated an ability to detect underwater chains via cross-correlation of FLS data alone [13]. However, this was only accomplished if the sonar was 1-4 m from the links and moving at 0.25 m/s precisely in line with the chain. This motion profile would be difficult, if not impossible, to achieve in the VSW zone. Zhang et al. recently explored the use of a general regression neural network and Gaussian particle filter to use acoustic images from an FLS to attempt to track moving targets [14]. Their results show that tracking targets in underwater environments is feasible, but their work requires future validation in more complicated, realistic ocean environments.

Automated detection with FLS does have successful implementations. For example, Healey et al. implemented reactive obstacle avoidance for an AUV with a blazed-array FLS [15]. ATR algorithms can also "flag" potential targets for human operators to classify upon mission completion. This saves vast amounts of time that would be spent analyzing sonar data from the entire mission. However, automated classification with a standalone FLS in the VSW remains an area of active research.

1.2 Convolutional Neural Network (CNN) Approach

CNNs are artificial, multi-layer neural networks that are designed for use on two-dimensional data, such as camera images and videos [16]. An example application which uses a generic CNN for automated surface inspection techniques is described in [17]. In another applica-

tion, Girshick et al. adapted CNNs to create RCNNs for improved object detection with less training data [18].

Naval Postgraduate School (NPS) researchers therefore studied the ability of a standard, pre-built RCNN algorithm to detect targets with a forward-looking camera exposed to oscillatory motion. A 3D-printed mount was designed and built to restrict translation in the x, y, and z directions, but permit roll, pitch, and yaw rotations.

The first step in implementing the algorithm was to create training data. Camera data were trained using MATLAB's built-in Video Labeler application, which treated a video as a series of individual images and automatically labeled the target in each image. This series of labeled images served as the ground truth for the RCNN training algorithm. This algorithm was derived from a standard, pre-trained neural network provided by MATLAB [19]. After setting training options, the ground truth file was run through the training algorithm, which output a RCNN detector file. Applying this detector file to untrained video produced an annotated video with detected targets and confidence ratings (Figure 1.2).

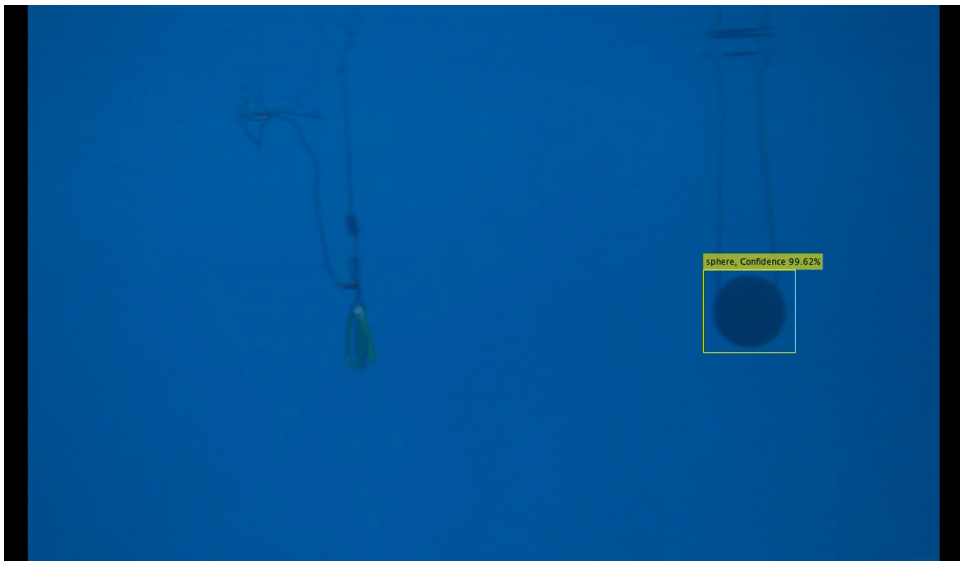

Figure 1.2. Example RCNN detector result with confidence rating.

Analysis of multiple camera trials using the RCNN algorithm determined that the standard algorithm worked exceptionally well. Even applying the maximum possible motion achiev-

able with the camera mount did not cause the detector to fail. Although the camera images were not in perfect focus and image blurring was observed, the algorithm still successfully detected the object. This indicates that the applied RCNN algorithm was robust to image distortion, even though it had only been trained with fewer than 1000 images.

Due to initial success with camera images, this research effort initially focused on applying the same training and detector approach to images from the FLS described in Chapter 2. This initial focus failed despite numerous trials. The failure to apply an RCNN designed for camera images to FLS images was most likely due to sensor modality. Cameras produce a two-dimensional signal that can be mapped to an array of pixels in an image. A sonar beam, however, produces a one-dimensional signal that can be mapped to its range from an object. An array of sonar beams can be used to construct a two-dimensional image based on the array's geometry and the range measurement from each beam. The resulting image, however, differs greatly from the two-dimensional signals produced by a camera. As a result, CNNs that were trained to detect color, contrast, texture, and/or spatial relationships in camera images fail for sonar images that lack this information [20]. Deep learning with FLS, however, is an ongoing field of study. Recent examples of positive contributions are described in [21], [22], and [23].

Despite the inability to translate deep learning success from camera images to sonar images, it is important to explore the extended capabilities of a FLS (i.e., beyond detection). This will support AUV MCM missions in the VSW zone, where side-looking sonars are ineffective and forward-looking cameras can not be used because water clarity is often poor.

1.3 Contributions

AUVs which could autonomously detect, classify, and identify mines in challenging VSW environments at stand-off ranges would benefit Navy MCM operations. This capability will provide situational awareness to the force while reducing risk to human lives and maintaining a covert posture. FLS have proven their ability to successfully detect mine-like objects, and identification must be completed visually by an operator. This leaves classification as the remaining barrier to success.

Therefore, the objective of this research is to determine if a single AUV, equipped with

a standalone FLS and following a non-standard path around an object, can obtain FLS images with features that support the classification of that object. Utilizing lessons learned from the prior camera research, a controlled testing environment was constructed to record repeatable sonar data with targets in multiple positions and orientations. The targets varied by size, shape, and material. The recorded imagery was processed to analyze features of the returns, including temporal variation, peak intensity values, and intensity distribution relative to each target.

The experimental setup discussed in Chapter 2 provided a controlled, stable testing environment. The data collection process that generated the FLS images used for analysis is explained in Chapter 3. Chapter 4 describes the methods of extracting and processing recorded data in preparation for analysis. Chapter 5 details the analytical methods applied to FLS imagery and discusses key experimental findings. Finally, Chapter 6 summarizes the conclusions, evaluates the implemented methods, and recommends future work.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 2:
Experimental Setup

A controlled testing environment was constructed to generate and record sonar data from targets in multiple positions and orientations. This maximized the amount of information that can be learned from sonar images. A water basin was used to collect images from two forward-looking sensors, a camera and a sonar, of two different targets. Including a camera along with the sonar provides visual confirmation of the target's position in the sonar field of view. The particular sensors used are representative of sensors used on AUVs and ROVs that conduct MCM operations in the VSW. The Robot Operating System (ROS) middleware utilized in this research enabled the collection of simultaneous, time-stamped data from the camera, sonar, and an IMU. The experimental setup described in this chapter eliminates additional factors of the real world (e.g., variable currents, rapid depth changes, or water visibility) in order to examine FLS images of known targets under specified test conditions. The objective of this analysis is to identify image characteristics which correspond to known target geometry, characteristics that might support automatic target classification.

2.1 CAVR Tank

The CAVR tank, shown in Figure 2.1, is an indoor, concrete, freshwater tank with internal dimensions of 6.6 x 4.1 m and a notional water depth of 1.8 m (6'). A recirculating filtration system combined with ample lighting provides ideal water clarity for recording camera images without particulates or shadows. The concrete walls are 0.2 m (8") thick on all sides. The thickness and material of the walls offer a clearly visible outline of the tank in all sonar images.

A movable walkway staged on top of the walls provides the ability to place targets at any position within the tank by suspending the target down into the water basin from the walkway with ropes. Since the tank walls are clearly visible in all sonar images, target placement near the tank walls was avoided to minimize interference.

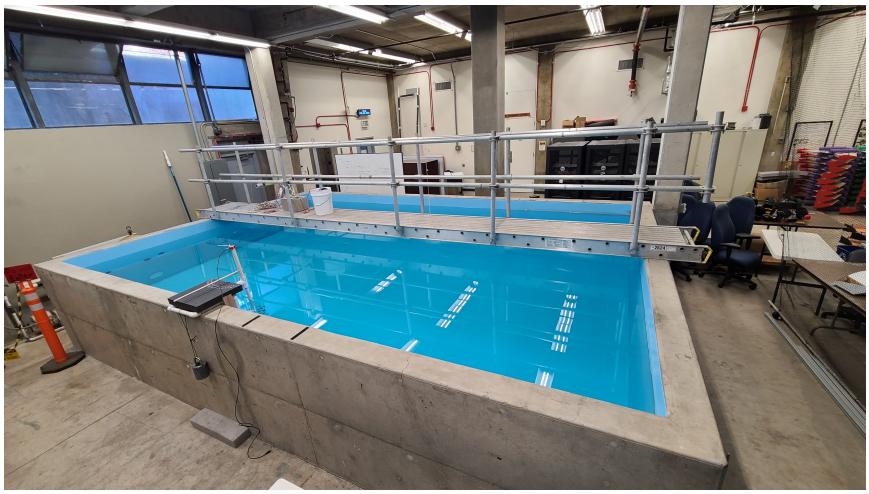

Figure 2.1. The CAVR tank with movable walkway.

2.2 Forward-Looking Sonar

The FLS used in this research was the BlueView M900-2250-130 Dual Frequency Sonar with a 1,000 meter depth rating. This sensor is commonly used in unmanned MCM vehicles, such as the REMUS 100 AUV [24] or vLBV300 ROV [25]. "Dual Frequency" refers to having both 2250 kHz and 900 kHz frequency sonar heads in a single unit. This sonar has 768 beams with a spacing of 0.18° between beams. Each beam has a horizontal width of 1° and a vertical width of 20°. The sonar unit is shown in Figure 2.2, and its nominal specifications are listed in Table 2.1. Part of the sonar's output includes metadata describing sonar settings and image characteristics like range resolution or bearing resolution. This metadata was used for image analysis. Due to the dimensions of the CAVR tank test, this research only utilized the 2250 kHz head; its operating range of 0.5-10 m (and stated optimal range of 0.5-7 m) was more appropriate than the 900 kHz head's operating range of 2-100 meters [26]. The entire CAVR tank occupies only about 5% of an image produced by the 900 kHz head, but occupies over 60% of an image produced by the 2250 kHz head.

Figure 2.2. The BlueView sonar used for this research.

Table 2.1. BlueView M900-2250-130 Sonar Specifications. Source: [27].

	900 kHz	2250 kHz
Field of View (FOV)	130°	
Max Range	100 m	10 m
Optimal Range	2-60 m	0.5-7 m
Beam Width	1° x 20°	
Beam Spacing	0.18°	
Number of Beams	768	
Range Resolution	1.3 cm	0.6 cm
Update Rate	Up to 25 Hz	
Supply Voltage	12-48 VDC	
Connectivity	Ethernet	
Weight in Water	0.43 kg (0.95 lbs)	
Depth Rating	1000m	
Size	21.8x9.1 cm (8.6"x5") (10.2 cm, 4" can)	

BlueView provides a downloadable Windows program, ProViewer, for viewing real-time images and adjusting sonar settings. While ProViewer was not used for data collection and

analysis during this thesis, it proved very useful for setting up the mount and learning about the sonar's operation.

2.3 Forward-Looking Camera

This research used a Basler acA1300-60gm GigE forward-looking camera. It has a 1.3 megapixel (1282x1026) resolution at 60 frames-per-second [28]. The camera weighs only 90 grams and has dimensions of 42x29x29 millimeters. The camera itself and its housing are shown in Figure 2.3. The housing is used on the REMUS 100 AUV, with a custom cable (seen in Figure 2.3) to provide a power and Ethernet connection. This camera utilized a 24 VDC power source and proved user-friendly and easy to operate.

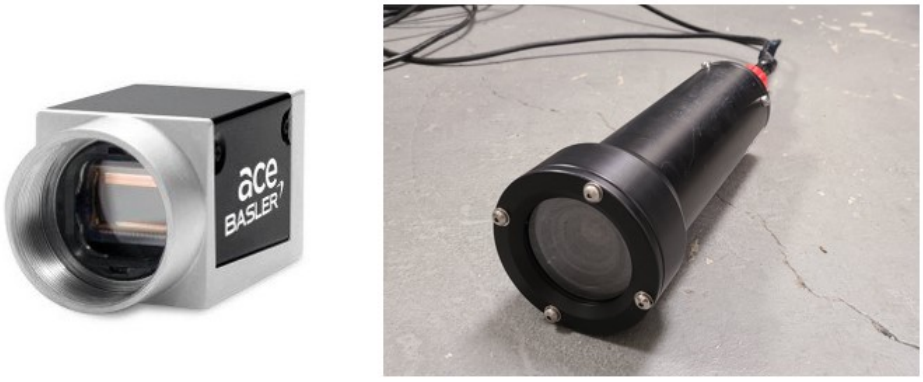

Figure 2.3. The Basler camera [29] (left) and the waterproof housing utilized during all camera testing. Note: images are not to scale.

2.4 Inertial Measurement Unit (IMU)

An inertial measurement unit was attached to the sensor mount to collect precise orientation data and support future experiments to identify relationships between platform dynamics and sensor performance.

The Xsens IMU used in this research is a plug and play, high-quality strapdown model [30] that collects and outputs 6-Degrees of Freedom (DOF) motion data. The specific unit used in this experimentation, the Xsens MTi-100, is shown in Figure 2.4. It measures orientation

and angular velocity, as well as linear velocity and acceleration at a sampling rate of up to 2 kHz.

Figure 2.4. This Xsens IMU recorded inertial motion data in every set of recorded data. Pencil included for scale.

2.5 Combined Camera and FLS Mount

On the original mount described in Chapter 1, the camera had to be manipulated by hand without precise, time-synchronized measurements of inertial data and camera images. The improved mount, shown in Figure 2.5, provided a rigid platform for both underwater sensors, a method for manual manipulation in one DOF, and precise time-synchronized measurement of sensor images and motions.

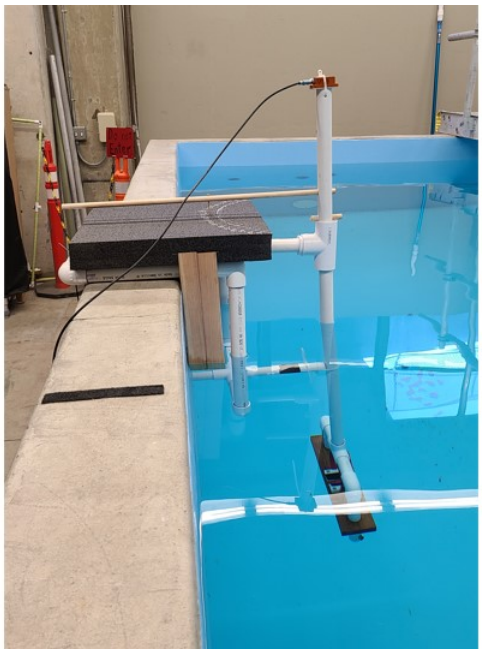

Figure 2.5. The side profile of the combined mount in the tank. Note: the camera and sonar are not installed in this image.

The new combined mount was restricted to one DOF, yaw, in order to maximize orientation control. Although only a single DOF is allowed, rotating the sensors 90° within the mount would simulate a pitching motion. Figure 2.6 illustrates the vertical FOV of the sonar when placed at half the nominal 1.8 m (6 ft) water depth on the tank's short wall (top sketch) or long wall (bottom sketch). Since the targets were always placed as close to the middle of the water depth as possible, the targets never left the vertical FOV of the sonar.

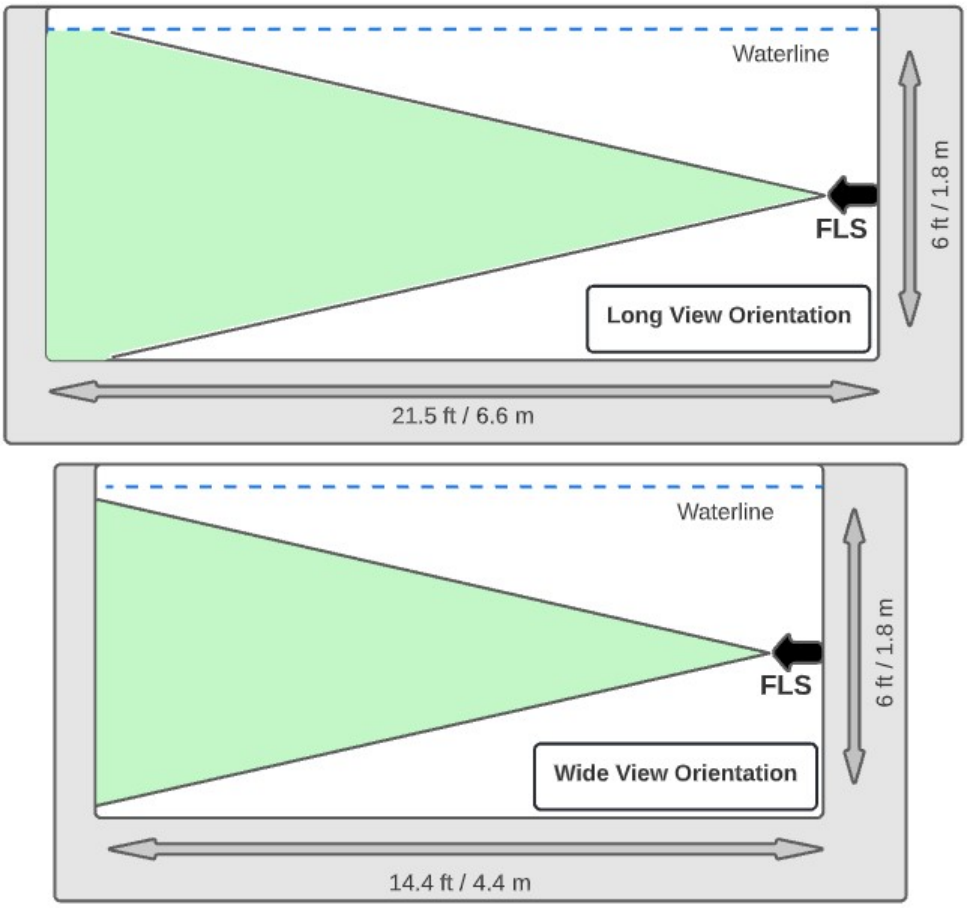

Figure 2.6. FLS coverage from side of tank. Note: diagrams are not to scale.

To simulate the relative change in range and bearing between the sensor and target that occurs as a results of vehicle motion, two different approaches were used during data collection. For the range tests, targets were re-positioned around the CAVR tank between data collection runs while the mount remained fixed in one position. For the offset tests, the mount was moved along the tank walls while the target was left in place.

The design concept, shown in Figure 2.7, of the combined mount was to slide the mount over the side of the CAVR tank with a hand-rotated vertical post supporting the sensors in the water. When in operation, the mount is in the same configuration as in the annotated figure, but placed on the other side of the CAVR tank wall.

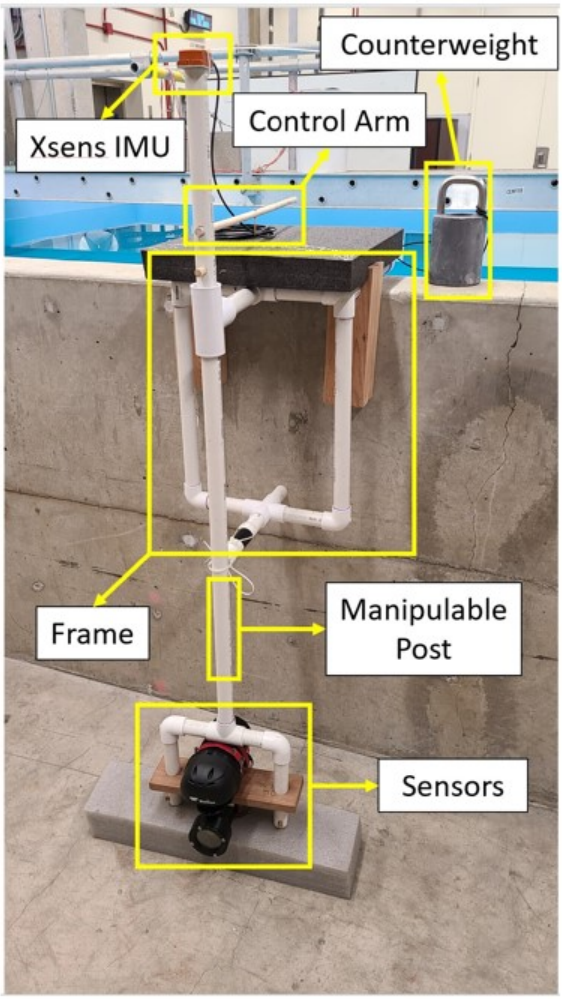

Figure 2.7. Annotated image of the combined mount from the front. Note: all sensors are disconnected in this image.

Frame

The combined mount was constructed of glued PVC piping and fittings with wood dowels and two-by-fours integrated for alignment. A counter-weight outside of the tank ensured the mount stayed still during data collection. The manipulable post, however, still allows a user to rotate the sensors in the yaw direction.

Manipulable Post

A vertical post hangs from the frame via a collar and wood dowel peg. The sensor platform was attached to the bottom of this post in order to place the sensors at half the tank's nominal water depth. The post was installed with a long wood dowel control arm. A screw in the control arm (similar to the needle of a vinyl record player) is used to indicate the approximate yaw angle on a scale marked every 5° on a large, square piece of foam. This portion of the mount is shown in Figure 2.8. The control arm was calibrated with live sonar imagery via the BlueView ProViewer software for a target hung precisely in the center of the tank with the mount located exactly in the center of the tank wall. Calibration was accurate to one tenth of a degree, although a small source of error exists from having to approximate the location of the exact center of the target return. The IMU was fixed to the top of the manipulable post with a zip-tie.

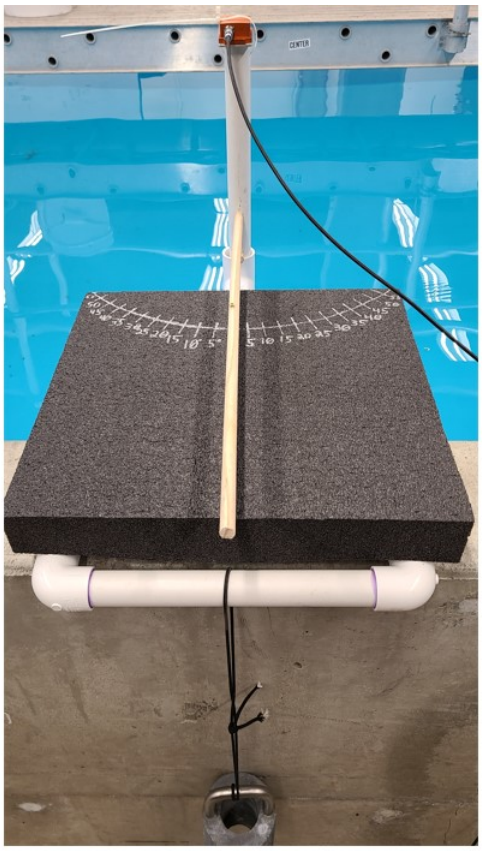

Figure 2.8. A view from above the control arm showing the arm and protractor. The Xsens IMU is visible on top of the manipulable post and counterweight at the bottom of the image.

Sensors Platform

A small section of one-by-four wood planking served as the mounting platform for the FLS and camera, as shown in Figure 2.9. The sonar and camera were fixed to the platform with adjustable metal ring clamps. This setup provided stable and secure mounting for the sensors while offering simple installation and removal. One person with a flat head screwdriver can easily loosen the clamps and remove each sensor. Secure and consistent mounting of the sensors is the key aspect to the effectiveness of the entire mount: data from both the camera and sonar is taken simultaneously from a precisely controlled and measured

orientation. This eliminates potential error from repeating experiments for the camera and sonar individually. This also simplifies data synchronization because camera, sonar, and IMU data are time-stamped using a common time reference.

Figure 2.9. A close-up of the FLS above the camera.

Data and Power transmission

Ethernet cables from the sonar and camera were connected to the Linux laptop used to execute ROS commands via a network switch, while the IMU utilized a USB connection for data transfer and power. The sonar was powered by a 3-prong 120 VAC plug, but the camera required a 24 V power supply.

2.6 Targets

The two targets chosen for this research vary in size, shape, and material, and are shown in Figure 2.10. The first is a 21 cm (8.25") diameter, thin-walled, water-filled, glass sphere. The second is a standard builder's 20x20x41 cm (8"x8"x16") cinder block. These were selected to vary the shape and cross-section of targets for the sonar. Ideally, material would not vary between the targets to eliminate material reflection properties as a parameter, but this was not possible for this research. Figure 2.10 also shows the attached ropes that were used in suspending the targets in the CAVR tank.

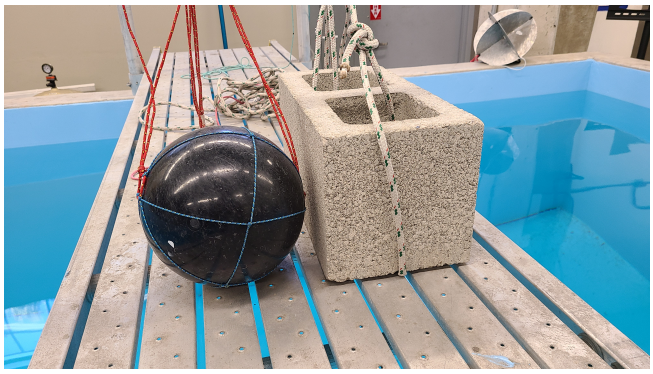

Figure 2.10. The glass sphere (left) next to the cinder block (right).

Each target was hung from the walkway's railing with ropes in a manner that ensured minimal movement in any direction. The ropes were secured using a self-tightening clove hitch knot. The depth of the targets was approximated at half the water depth, well within the vertical FOV of the sonar beams. Targets were easily suspended from different points along the walkway to adjust their distance from the sensors when aligned with the long dimension of the CAVR tank. The walkway itself was moved to adjust distance when aligned with the short dimension of the tank. Re-positioning the sphere was simple, as its own weight prevented movement and its cross-sectional profile does not change when the sphere rotates due to its inherent rotational symmetry.

The cinder block, on the other hand, was often difficult to properly orient. For these experiments, the cinder block needed to be level with a desired face as close to parallel with the tank wall as possible. The cinder block was supported by two ropes, one at each end. While this made leveling more difficult, it was still easier to work with than using four or more support ropes to manipulate each face independently. Additionally, the cinder block is heavy for one person to work with as it weighs about 15.9 kilograms (35 pounds).

In summary, the experimental setup described in this chapter provided control of one dependent variable at a time while minimizing uncontrollable variables for each test run. Testing designs, outlined in Chapter 3, leveraged this capability to incorporate multiple geometries and test setups in order to explore specific aspects of FLS imagery.

CHAPTER 3:
Data Collection

Fulfilling the objectives of this research called for the collection of a large amount of data. This meant that efficient, repeatable processes were necessary to setup the equipment and collect, store, and access recorded data. The experimental setup outlined in Chapter 2 was used to conduct two different types of tests. First, FLS images of each target were collected at different ranges. The objectives of these tests were to analyze a) the impact of range on the intensity of target returns, and b) the temporal variation of sonar images under static test conditions. Second, FLS images of a rectangular cinder block and a glass sphere were collected from different offset positions to investigate whether lateral offset can support target classification with a standalone FLS. This chapter describes the processes used for collecting and recording data, as well as the two unique testing structures used to support target return analysis. For the purposes of this research, a "return" refers to a group of pixels in a sonar image with increased intensity values attributable solely to target reflections. In these images, each pixel corresponds to a specific range and angle from the sonar.

3.1 Data Recording

All experimental data collection utilized the open-source ROS to log IMU, camera, and sonar data in a single, time-stamped rosbag (*.bag) file. ROS middleware provides a publish/subscribe architecture for inter-process communication [31]. Each process publishes data in one or more well-defined messages as topics that other processes can subscribe to. Individual sensor processes publish time-stamped messages at unique sampling rates, but every message in a given experiment was saved to a single bag file with a common time reference. MATLAB tools allow for selection of individual topics by name and/or time range to extract messages of interest. NPS CAVR developed a ROS driver using the BlueView software development kit to generate and record images from each sonar ping. Similarly, vendor-provided ROS drivers were used to record camera images and IMU data in real-time.

Additionally, all sonar data was separately recorded in the BlueView SON (*.son) file

format. This format is unique to BlueView and typically supports post-mission hydroacoustic visualization and analysis [32]. SON files were saved for every test as a backup for the bag files and to support sonar playback with BlueView's ProViewer software.

3.2 Test Procedures

Each of the two test types were repeated twice: once with the glass sphere, and once with the cinder block. Approximately 100 images were recorded during each test run to have a statistically significant sample size. These tests maximize control and minimize variability.

3.2.1 Range Test

The range tests explored how different factors affect the recorded images from the FLS. The first goal was to determine how varying range affected the size, intensity, and distribution of the target returns. Size refers to the amount of beams that received a return. Intensity is the strength of the return, on a standard scale of 0-255. Distribution is defined as the spread of intensities within a return. This testing also explored the temporal variation of the FLS images, i.e., how target returns changed from image to image for a specified static test condition.

The range test setup depicted in Figure 3.1 involved taking a series of images for each target at ranges of 2, 3, 4, and 5 meters from the sensors. The mount was placed in the center of the short wall (facing the length of the tank) with each target hung from the walkway spanning the tank. Distances were verified by measuring the distance between the manipulable sensor post and the rope suspending each target. The cinder block was oriented with its short side facing the sensors. One set of approximately 100 images was collected. For each data set, a few images were checked at random to verify proper sensor operation and data collection.

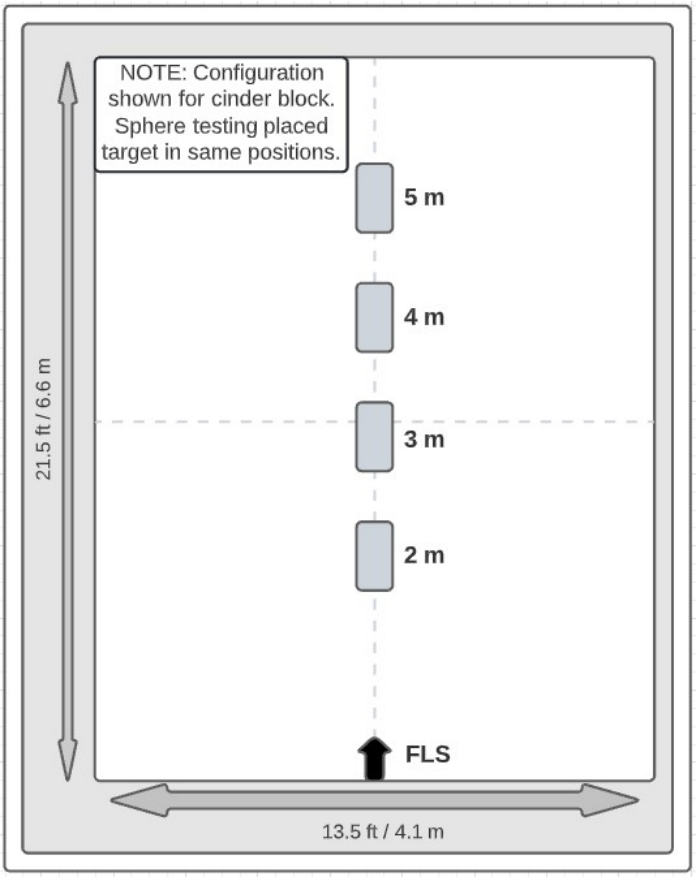

Figure 3.1. This depicts the setup of the Range Test. Note: diagram not to scale.

One-meter distances were not included in the range test analysis, because images in these data sets contained too much scattering to accurately isolate target returns. Although this distance is within the sonar's stated optimal range of 0.5-7 m, experiments suggested that targets need to be much smaller, or have lower target strength, to produce a usable return at this close range.

3.2.2 Offset Test

The goal of the offset testing was to generate images for qualitative analysis (via visual inspection) to investigate how each target's sonar return changed with lateral separation

from the FLS. This explored whether offset images could provide additional information about target size and geometry in support of classification. For these tests, the target was fixed at a distance 1.5 m (5') away from the initial sonar position, 1 m (3.3') from the nearest tank wall, and suspended in the middle of the water column. The sensor mount was moved laterally along the CAVR tank's long wall for each new test position. From there, the sensors remained normal to the long wall (i.e., at a 0-degree yaw angle). All other data collection procedures mirrored those of the range test.

Table 3.1 describes the locations and associated slant ranges to the target from each test position sketched in Figure 3.2. In this figure, the first and last test positions are denoted by a black sonar icon. For offset testing, the cinder block was placed with its long side facing the sensors (the short side faced the sensors during range testing). The offset values were measured relative to the center point of the target. This distance was then normalized by the target's body length. Observing changes in terms of body length could unveil a relationship between the physical size of the object and how far the sonar needs to be offset in order to support classification.

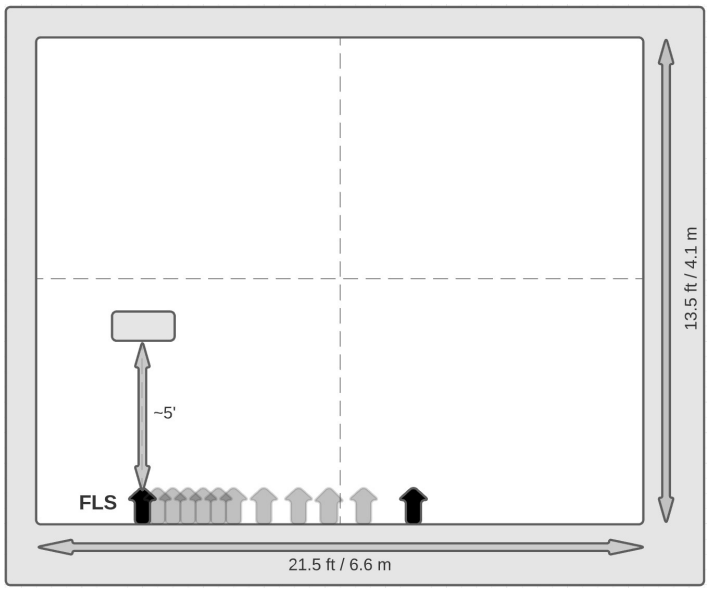

Figure 3.2. This depicts the setup of offset testing. Note: diagram not to scale.

Table 3.1. Offset Test Values.

Offset, d [cm]	Angular Offset [deg]	Slant Range [m]	Block d/L	Sphere d/L
0	0°	1.52	0	0
10.16	3.8°	1.53	0.25	0.5
20.33	7.6°	1.54	0.5	1
30.49	11.3°	1.55	0.75	1.5
40.65	14.9°	1.58	1	2
50.81	18.4°	1.61	1.25	2.5
60.98	21.8°	1.64	1.5	3
81.30	28.1°	1.73	2	4
101.63	33.7°	1.83	2.5	5
121.95	38.7°	1.95	3	6
142.28	43.0°	2.09	3.5	7
182.93	50.2°	2.38	4.5	9

The range and offset tests utilize controlled, repeatable procedures for consistent data collection. The consistency of the recorded data is crucial to the successful analysis of sonar returns in support of object classification efforts. Chapter 4 details actions required to process the data in preparation for analysis. These steps include bag file data extraction and compilation as well as target sub-region isolation.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 4:
Data Processing

The data processing methods applied to this research were necessary to properly isolate target data from background interference within the CAVR tank or target reflection scattering. Experimental data was first extracted from bag files to allow for comparative analysis. IMU data was plotted to verify stationary test conditions. Finally, sonar returns for both the range and offset tests were isolated and extracted from their original image formats through multiple extraction techniques.

4.1 ROS Bag Extraction

Upon completion of each test, bag files were saved to cloud storage for easy access. All data processing was then completed in MATLAB R2021a. MATLAB has built-in commands specifically designed for extracting and processing data from bag files [33]. A timeseries format was used for IMU data extracted from a bag file because it guaranteed time synchronization with the sensor images. Additionally, MATLAB timeseries can easily be manipulated to account for different sampling frequencies [34]. The camera and FLS data were extracted as images from the bag files. Once all the data was extracted, the IMU and sonar data could be processed and prepared for analysis separately.

4.2 IMU Data Processing

The IMU state variables analyzed for this research were angular position (orientation), angular velocity, and linear acceleration. For this research, linear acceleration was verified to be zero for the stationary test conditions. Future work that subjects the sensors to motion will be able to utilize the acceleration data to analyze these motions. Using built-in MATLAB functions, angular position values were first converted from the standard ROS units [35] of quaternions to Euler angles in radians, and then from radians to degrees. Angular velocity only required conversion from radians/sec to degrees/sec. Timestamps were converted from Unix time to elapsed seconds by subtracting the first timestamp from each data point in a timeseries.

Once extracted, the angular positions and velocities were simply graphed against time to verify that angular and linear accelerations were within the noise level of the sensors. This provided confidence that the sensor remained stationary during testing. IMU data from an early trial is displayed in Figure 4.1. This example trial, completed to gather data for RCNN applications, involved rotating the sensors back and forth by hand to explore the effects of yaw rate on sonar and camera images.

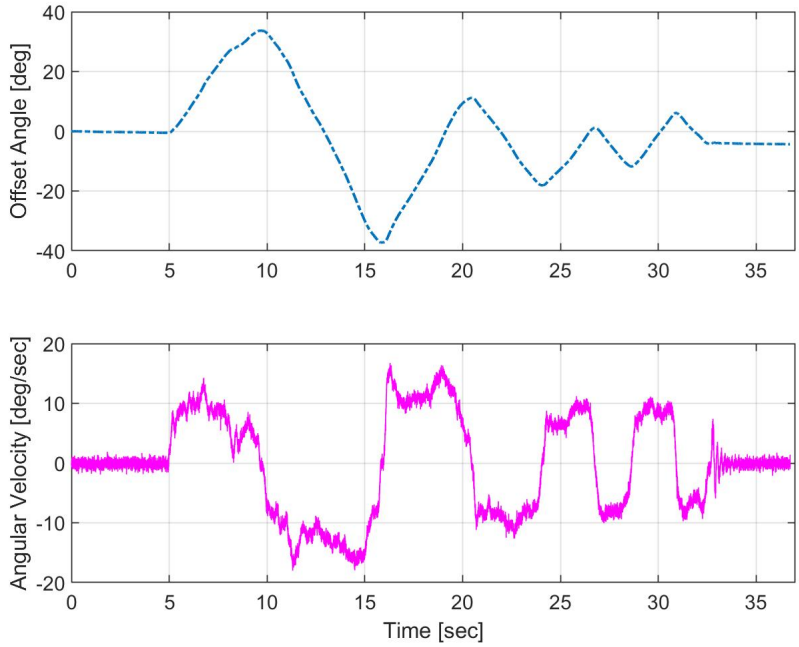

Figure 4.1. Graphs of IMU yaw data during an early trial. Note: positive numbers indicate panning the sensors to the right.

4.3 Sonar Data Processing

Sonar data was extracted from the bag files via the same method used to extract IMU data, except as MATLAB cell arrays instead of timeseries. This process produced a one-dimensional column vector of two-dimensional image data. Therefore, each image had to be reconstructed by reshaping this vector into the appropriate image size. BlueView

sonars output two different image arrays, using either polar "r-theta" or Cartesian "X-Y" coordinates. The r-theta images are 768x1223 pixels. Each row corresponds to one of the sonar's 768 unique beams and each column corresponds to a unique range bin. Sonar beams can be mapped to physical angles from the sonar centerline using the sonar's angular resolution, while sonar range can be mapped to physical range measurements using the sonar's range resolution. Both of these parameters are output as metadata along with the images themselves. X-Y images are 2298x1242 pixels, and map the same sonar intensity data into a Cartesian coordinate system whose origin is centered on the sonar, with rows and columns corresponding to lateral and longitudinal distance from this origin, respectively. Although X-Y images are useful for presenting sonar data to a human operator, r-theta images are a more natural representation of sonar measurements. The r-theta formulation also makes it easier to compute and apply geometric relationships. Therefore, this research used r-theta images for analysis, an example of which is depicted in Figure 4.2.

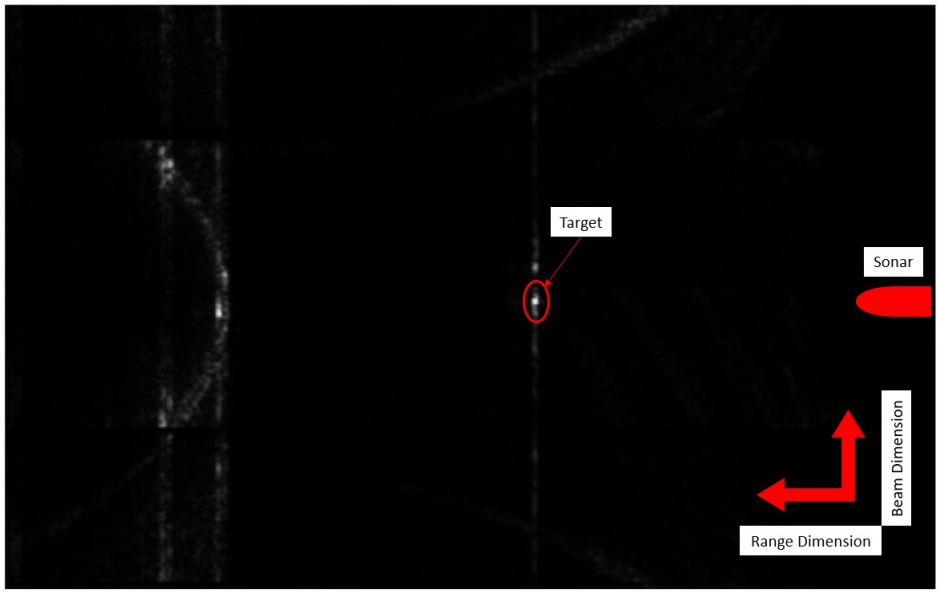

Figure 4.2. Annotated image from range test with the cinder block target at 4 meters.

Figure 4.3 compares r-theta and X-Y images for the 50th ping of a range test with the cinder block placed 3 meters away from the sonar. Intensity data is presented as a gray-scale

image, where pixels range from black (intensity = 0) to white (intensity = 255). Light pixels indicate higher intensity returns, while dark pixels indicate low-to-no returns. The target in Figure 4.3 is denoted by the red box.

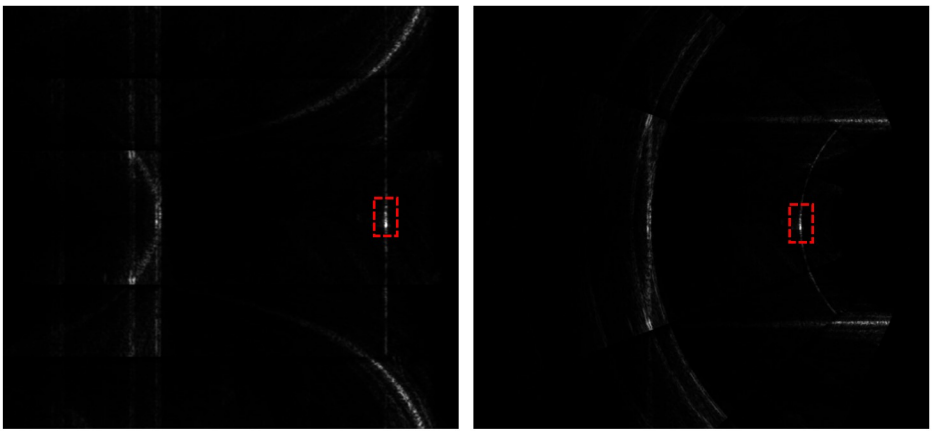

Figure 4.3. Sample of r-theta (left) and X-Y (right) FLS images from the same sonar ping.

The sonar driver produces images for the entire range of the sonar due to the assumption that a vehicle operates in an open, unbounded environment. As a result, each sonar image contains range data out to 10 meters, nearly twice the size of the tank's longest dimension. Figure 4.3 shows clear returns from the CAVR tank walls, as well as scattering from the target. Note that straight walls appear as curves in the r-theta image. In order to conduct proper quantitative assessment of the target, these extraneous returns required elimination by creating sub-regions containing the target and a small amount of empty space around the target.

4.3.1 Range Test Target Sub-Regions

New images of smaller sub-regions were created to exclude spurious returns and returns from the walls of the CAVR tank. These sub-regions were centered on the highest observed intensities and sized to include both the target and a very small amount of empty space around the target. Figure 4.4 depicts an example of an extracted sub-region.

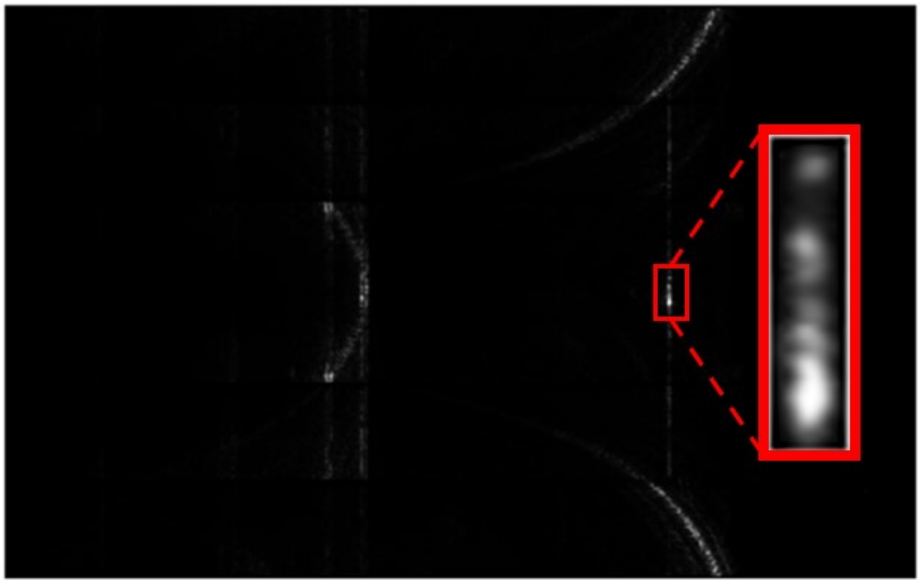

Figure 4.4. The original cinder block 3-meter range test image. Inset: the selected sub-region.

Expected Number of Beams Spanned by Target

The trigonometric relationship

$$\text{Expected Beams Spanned by Target} = 2\tan^{-1}\left(\frac{\text{Target Width}/2}{\text{Target Range}}\right)$$

was used to determine the expected number of beams spanned by the target in a return. The limitations of the CAVR tank and combined mount only allow for hand-placement of the mount and targets. This introduces the potential for up to a few inches of alignment error to either side of the sonar's center line. To account for this error, the sub-regions of each image were selected to include the expected number of beams for the largest target, the cinder block at 2 meters, plus 5 beams on either side. A perfectly controlled experimental setup would make it easier to obtain a sub-region that minimizes the number of pixels corresponding to spurious returns or empty water. This ideal sub-region could then be automatically correlated with the center of a sonar image to produce a new image containing only the target. As previously stated, however, this accuracy was not possible with the setup utilized by this research.

4.3.2 Offset Testing Target Sub-Regions

Offset testing utilized a different isolation strategy than range testing because the sub-region needed to be large enough to reveal other information about the return (e.g., other faces). This strategy also provided a cleaner presentation for comparison of each offset location. Therefore, the size of the cropped image was the same for each offset data set and was large enough to include every test's target return. This facilitated side-by-side visual comparison of offset images without changing image size or scale.

4.3.3 Final Image Processing

Once each data set was appropriately isolated, the full set of images for each range test trial (e.g., all 111 images) were "stacked" into one three-dimensional array. This facilitated analysis of the intensity changes of each pixel from image to image (i.e., its temporal variation). Temporal variation results can provide insights about features of a sonar return, like noise distributions, and how returns vary with aspect changes for a given target.

At this point, a controlled experimental setup was constructed, data was recorded, and the recorded data was processed. In Chapter 5, analysis of the range and offset tests is conducted. For the range test, this involves examination of temporal variation before analyzing sonar returns in the range and beam dimensions. Range test results motivate offset test analysis, which are analyzed to investigate the potential for classifying targets in FLS images taken during a laterally offset approach.

CHAPTER 5:
Results and Analysis

In this chapter, results from the range and offset tests are presented using plots of intensity values and comparisons of sonar return images. First, temporal variation is explored to determine the variability of a sonar return in consecutive images under static test conditions. Next, range test imagery is analyzed in the range and beam dimensions. The range tests suggest that targets can not be definitively classified by a standalone Forward-Looking Sonar (FLS) using a head-on approach. Therefore, offset test data are analyzed to investigate the potential for classifying targets in FLS images taken from an offset approach.

5.1 Range Test

Figure 5.1 shows the time progression of intensities within the isolated sub-region of the glass sphere while Figure 5.2 illustrates the same for the cinder block returns. Both figures show results from the 2-meter range test. The red boxes within the images indicate the estimated location of the target. The number of beams spanned by the boxes was derived from the geometric relationship between range to the target and its known width. The boxes were centered on the location of the peak intensity of the sub-region.

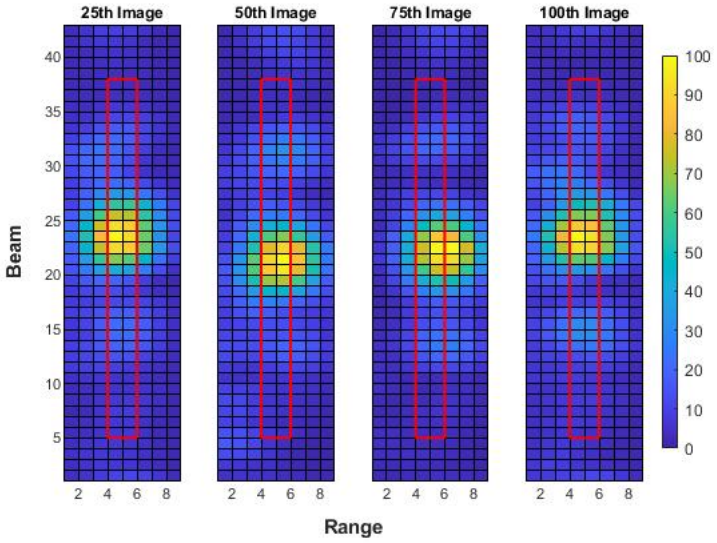

Figure 5.1. Representative progression of images from the 2-meter range test with the glass sphere target.

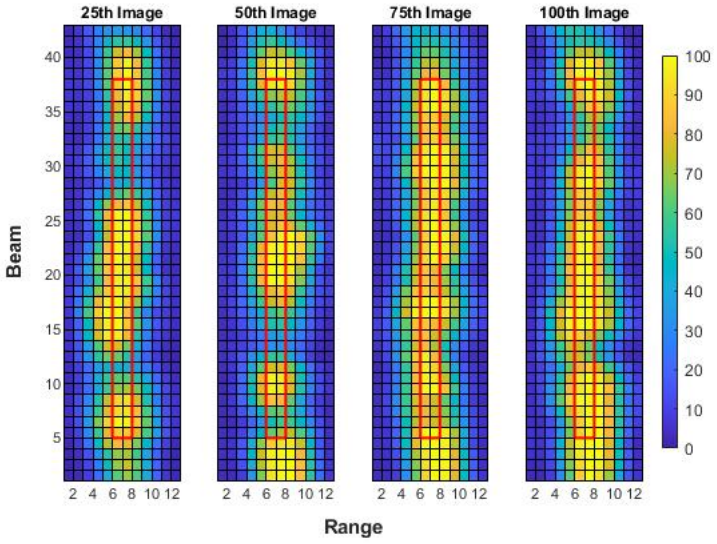

Figure 5.2. Representative progression of images from the 2-meter range test with the cinder block target.

The representative results of Figure 5.1 indicate that three areas of analysis are required for the range test. First, the images show that intensity variability over time is present within

the data set. This is most clear in the 75^{th} image of Figure 5.1 due to noticeably higher intensities at the peak compared to the other images. This temporal variability is even more pronounced for the cinder block target. Next, the peak returns in all images span multiple range bins rather than a single bin. Therefore, an examination of the range dimension will be conducted. Finally, the peak intensity values do not span the full amount of beams expected within the return for the sphere. While for the cinder block test in Figure 5.2 the intensity values extend outside the expected span of beams. These results necessitate an examination of the beam dimension.

5.1.1 Temporal Variation

Temporal variation analysis utilized the three-dimensional arrays of 100 "stacked" sonar images from each data set of the range test. In addition to beam and range, these arrays allowed for the examination of time, the third dimension. The standard deviation of the intensity values over time for each pixel of every image of the range test was computed. Histograms of intensity values over time as a percentage of the maximum value, 255, were created for six pixels within each image sub-region. The histograms, along with the standard deviation plot, for the 2-meter test with the glass sphere target are shown in Figure 5.3. Pixels were selected from the following approximate locations:

1. Peak intensity, which correlates to the assumed center of the target.
2. In the beam dimension: halfway between the peak intensity pixel and the edge of the visible return. This does not reflect the true geometric width of the target.
3. In the beam dimension: the edge of the visible target return. This also does not reflect the target's actual geometric width.
4. In the beam dimension: a pixel far outside the visible and geometrically-expected target edge. This should reflect empty water in the CAVR tank.
5. In the range dimension: halfway between the peak intensity pixel and the edge of the visible return.
6. In the range dimension: the edge of the visible target return. This occurred near or at the edge of the sub-region in all data sets.

Figure 5.3 shows that pixel intensities near the assumed target center range from approximately 35% to 65%. Adjacent pixels in the beam or range dimensions range from approx-

imately 10% to 50%. The points outside the return (histograms 3, 4, and 6) show a very narrow peak of weak intensities, between 0% and about 15%. The intensity distributions in histograms 1, 2, and 5 are approximately uniform or normally-distributed. These distributions indicate that a mean value is suitable because most of the data points occur near the mean. As a result, the time-averaged intensity value for each pixel will be used to represent the intensity values of the sphere target data.

Intensities from the sphere 3-, 4-, and 5-meter tests are weak in all locations, leading to low variance. Since these data sets only include a sharp spike at very low intensities and no other occurrences, the only information they provide is that low intensity returns also have low variability. Histograms and standard deviation plots for the entirety of the range test can be found in Figures A.9-A.16 in the Appendix.

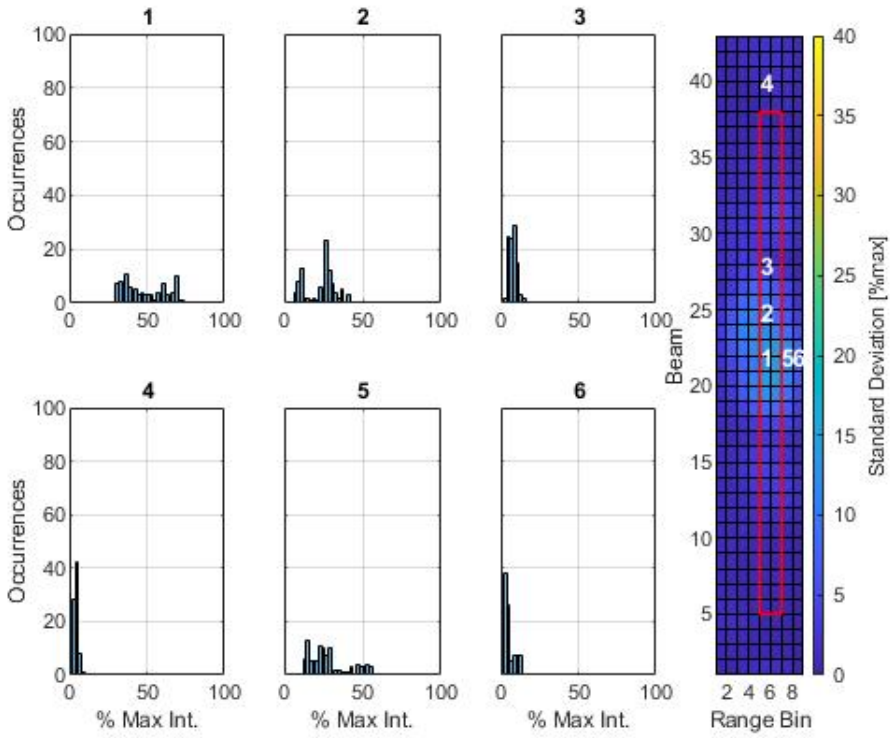

Figure 5.3. Sphere 2-meter range test standard deviation plot with histograms from six points within the image.

Figure 5.4 shows data from the cinder block 3-meter range test. The peak (histogram 1) intensity value occurred more than twice as frequently as any other value, but this only occurred in the 2- and 3-meter range tests. The peak returns in both of these tests also displayed very high intensity values. The 4- and 5-meter tests show intensity distributions that are similar to Figure 5.3. The other regions of the return (histograms 2-5) display either uniform or somewhat-normal distributions. This trend holds true in almost all of the range test data sets. Like the sphere data sets, taking the time average intensities of cinder block data sets is a reasonable representation of the target intensity through time. Therefore, time-averaged data will be utilized in the analysis below.

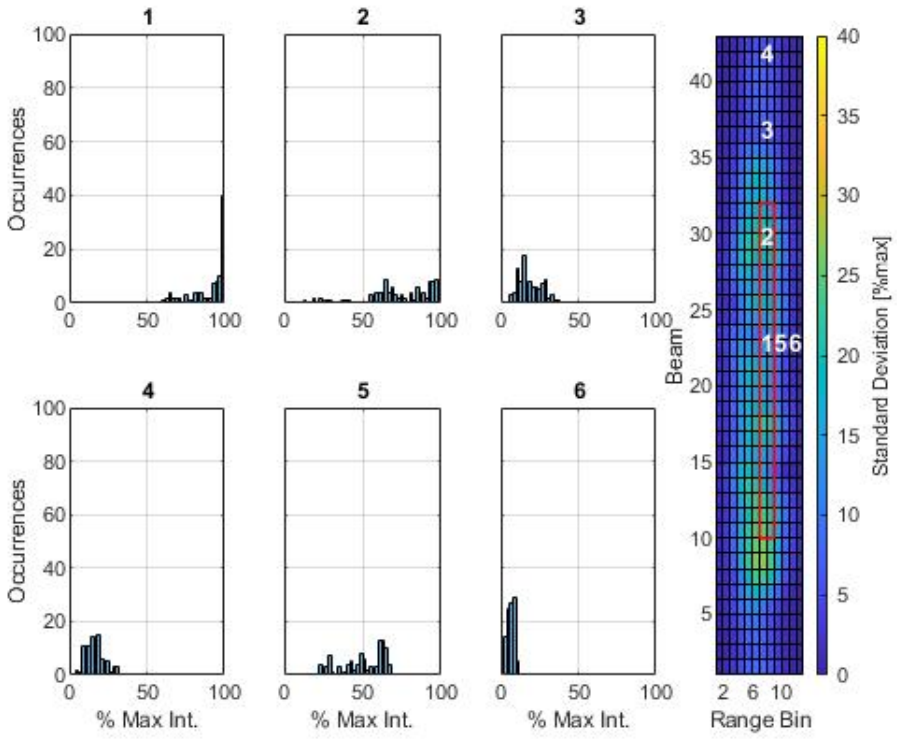

Figure 5.4. Cinder block 3-meter range test standard deviation plot with histograms from six points within the image.

5.1.2 Range Dimension

Range dimension analysis used the time-averaged images for each sub-region and individually plotted the intensities for each beam in every data set. Figure 5.5 shows the resulting plots, with each line representing the intensities measured by an individual sonar beam for each range bin in the target sub-region. These plots clearly show that intensities were approximately normally-distributed in the range dimension. This trend was consistent for both targets at all ranges (glass sphere target plots are shown in Figure A.18 of the Appendix). While the shape of each plot is consistent across all beams, the peak intensity measured by each beam decreases with angular separation from the beam with the target's peak intensity value.

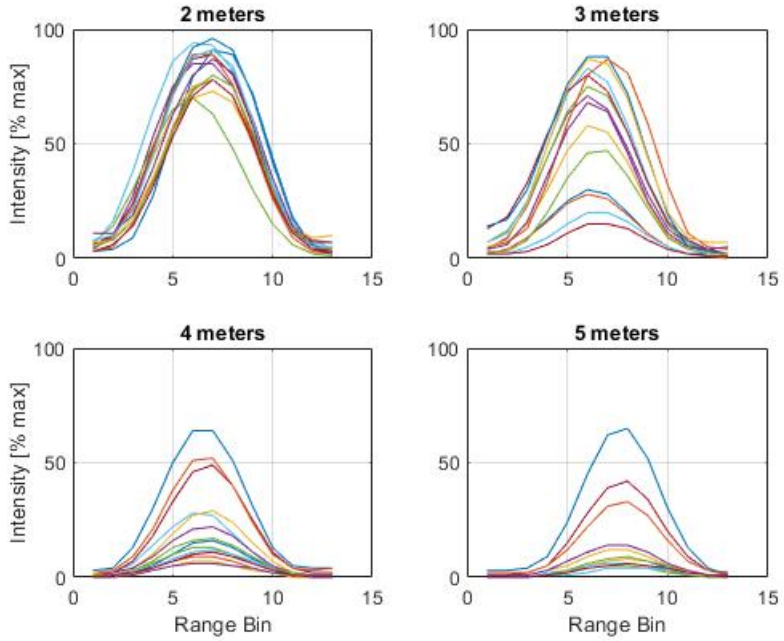

Figure 5.5. Cinder block range dimension plots for range location.

Normal distributions in the range dimension, with means centered on the true value, are expected because each beam provides independent range data with some amount of random measurement uncertainty in each measurement. Averaging over all bins in the range dimension for each beam would inaccurately underestimate the intensity value corresponding to

the target's expected range for that beam. The normally-distributed intensity values in the range dimension suggest that using the range bin containing the peak intensity value will more accurately represent the target's true range estimate for each beam.

This allowed the target to be represented by an array of intensity values that appear as a line in the beam-range plane. Note that the location of the peak value in the range dimension may vary by one or two successive range bins, as the range measurement can be influenced by noise or multi-path effects. However, the sonar's range resolution of 1.6 cm is well below the target placement uncertainty of the experimental setup; therefore, the peak value in the range dimension will be used in the remaining analysis.

5.1.3 Beam Dimension

Based on the temporal and range-dimension findings above, beam-dimension analysis utilized the peak intensity values for each beam in a time-averaged image of the target sub-region. This process reduced a range of intensity values in a two-dimensional group of "target" pixels to a one-dimensional array of intensity values corresponding to each sonar beam. The peak return intensities were used to align this vector with the expected number of beams spanned by the target, as computed from test setup geometry. Table 5.1 lists the difference between the beam with highest observed target intensity and the beam at the center of the image sub-region. This difference amounts to 0.2 to 2.8 degrees of angular uncertainty in placing the target relative to the sonar, which can be observed by the positions of the red boxes in Figures 5.1 - 5.4. While this uncertainty does not negatively affect the beam-dimension analysis described below, a more precise target placement method could be devised for future experiments.

Table 5.1. Measured sub-region target location errors for the range test.

Target	Range [m]	Error [beams]	Error [deg]
Cinder Block	2	5	0.87°
	3	10	1.73°
	4	9	1.56°
	5	3	0.52°
Sphere	2	16	2.77°
	3	14	2.42°
	4	1	0.17°
	5	2	0.35°

Figure 5.6 displays the time-averaged peak intensity across a number of beams centered on the target. The left subplot shows the cinder block results while the right subplot contains the sphere results. Each subplot contains four curves corresponding to the four range test distances. The solid vertical lines correlate to the expected number of beams spanned by the target, computed as described in Chapter 4, at the specified range from the sonar. Each plot depicts the intensity as a percentage of the maximum possible value, 255. The beam range shown (+/- 21 beams) was selected to include the expected target return at a range of 2 meters, plus 5 beams on either side.

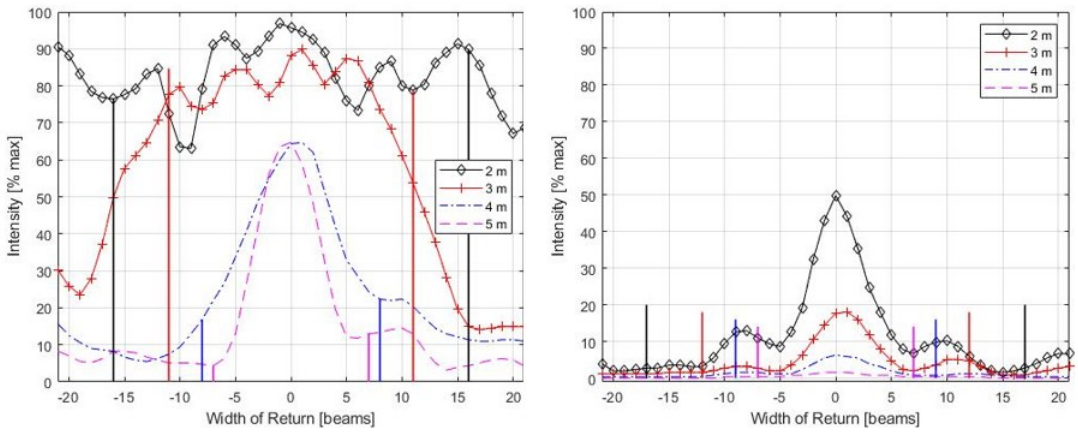

Figure 5.6. Cinder block (left) and glass sphere (right) beam dimension plots for each range test location.

The 2- and 3-meter cinder block tests both show high intensity values that span more beams than expected. The 2-meter test, in particular, does not have a significant decrease in intensity within the sub-region. Figure 5.7 shows that the strong return intensities do not decline until noticeably outside of the sub-region, depicted by the dashed red box, giving the false impression that the target is extremely large relative to its actual size. Even after the intensities drop off from the near-maximum values, returns of up to approximately 50% remain well beyond the edges of the actual target. Additionally, intensities that exceed the "empty water" intensity values observed in other data sets can be found across the entire beam dimension shown.

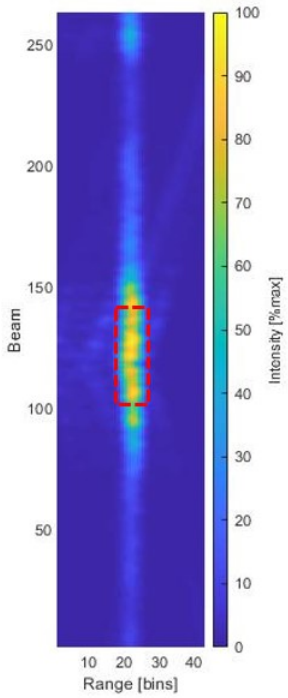

Figure 5.7. Extended beam dimension view of the 2-meter cinder block range test data.

The remaining data from both targets in Figure 5.6 show intensity values that drop off significantly before reaching the expected number of beams for each target. Surprisingly, their intensity profiles are all about 10 beams wide, although the 4-meter cinder block return does show a slightly wider profile. The exact reason for this is currently unknown. All sphere tests produced returns spanning significantly fewer beams than geometrically expected. This suggests that the number of beams spanned by a sonar return cannot be correlated with actual target width for flat-faced targets that are close to the sonar or targets with rotational symmetry. The cinder block 3- and 4-meter range returns, however, indicate their returns may correlate to their actual geometric widths. This finding requires further investigation to confirm. Also, both targets produce observable, low-intensity secondary (and possibly tertiary) peaks in their returns. Although the root cause of these sub-peaks is unknown, they are produced by both target types and are not geometry-dependent. Therefore,

these features cannot be used to aid in target classification.

Finally, Figure 5.8 shows that targets with different shapes and different ranges can have nearly indistinguishable intensity profiles. This suggests that it is possible to switch the size and shape of an object at a given range and receive the same intensity profile across the sonar beams. In this research, for example, it is impossible to distinguish whether the return of an unknown target was produced by a small cinder block or a large glass sphere. Object material can also affect the intensity profile of an unknown target, further complicating target classification.

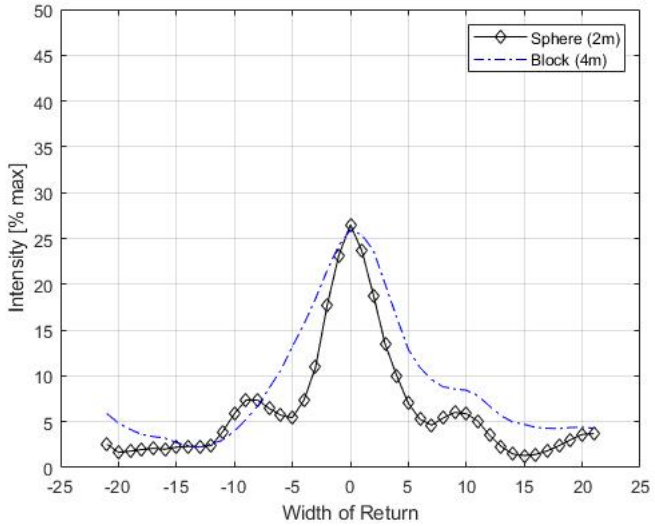

Figure 5.8. Comparison of sphere and block returns from the range test

Range and beam dimension analysis of the range tests showed that a head-on sonar orientation does not provide enough information to make a size or shape determination for spherical or flat-faced targets. The remaining analysis will investigate if laterally offsetting the sonar from the target will provide additional information needed to support object classification.

5.2 Offset Test

Unlike the range test, the offset test utilized a larger, constant sub-region for each image in the data set, as explained in Chapter 4. Offset test analysis adopted the range test approach of time-averaging each data set to produce one image for each sub-region. Figure 5.9 is a reproduction of Figure 3.2 as a reminder of the offset test setup.

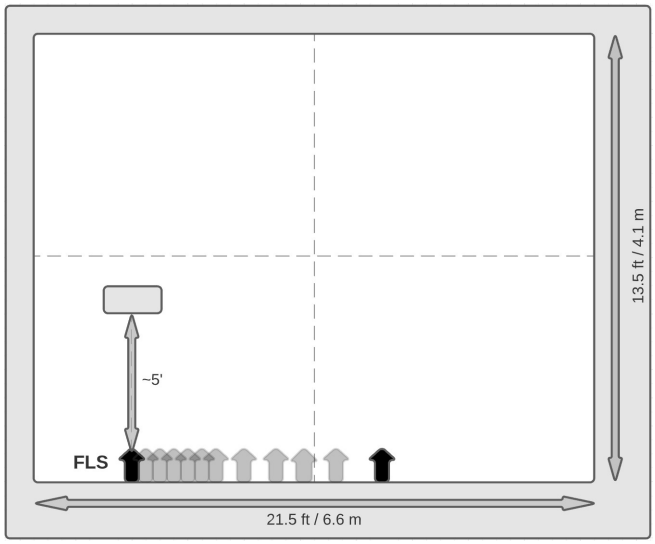

Figure 5.9. Offset test setup. Note: diagram not to scale.

A sequence of time-averaged images of the target sub-region, taken from each offset position, is presented in Figure 5.10. Each image in this sequence has the same orientation as the range test figures, with the sonar to the right of the image. However, because target returns span different beam and range bins within the sub-region as lateral offset distance increases, the offset test analysis was not performed along range and beam dimensions. The range to the target in the first image (0 L offset) is approximately 1.5 meters, and the range in the final image (4.5 L offset) is about 2.4 meters.

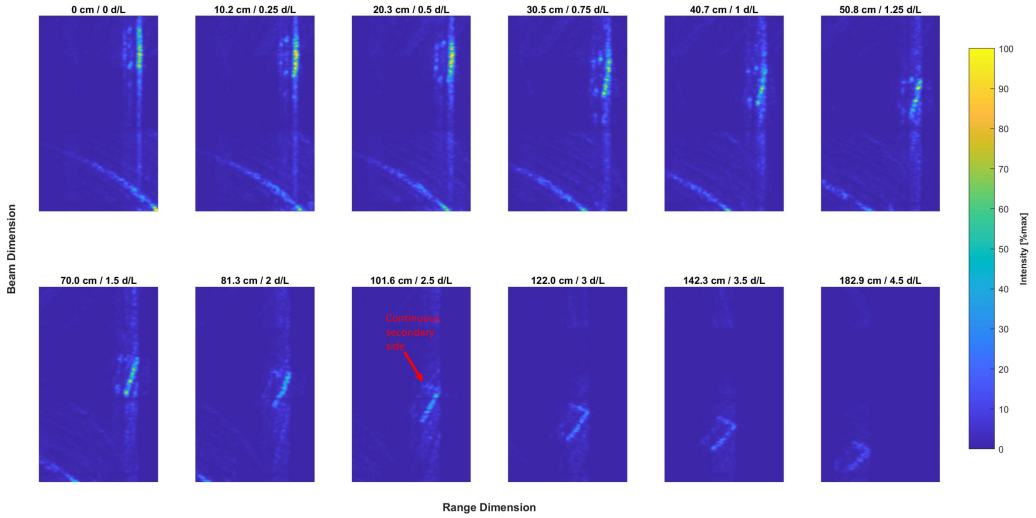

Figure 5.10. Collage of offset test sub-region images with the cinder block target.

Visual examination of images from offset positions 0 L through 1.25 L indicate the presence of a pronounced primary face. Each of these sub-regions contain 100% intensity values at several points along the face, but these peak intensity values are not distributed uniformly across the face. Intensity values range between 50% and 100% of maximum intensity along the primary face in each of these first six data sets. After a significant offset of 2.5 L, a complete second face is visually observable. Roy, et al. describe image processing techniques that could be applied to this image to "fill in" the gaps between the primary and secondary faces [36]. These techniques could enable a corner-detection algorithm such as the ones described in [37] to classify this target as a rectangular object. The information provided in offset FLS images might therefore allow a geometry determination to be made prior to imaging the entire target, potentially enabling target classification.

The fully connected secondary side is indicated by the red arrow in Figure 5.10. Even though it is fully connected and clearly visible, the secondary face only contains intensity values around 25%, while the corner of the block contains intensity values near 60%. By comparison, most of the pixels in the primary face have intensities around 30%, although some pixels in one portion of this face have intensities just over 50%.

As offset distance increases, overall intensities reduce significantly. Therefore, care must be taken during path planning to avoid large arbitrary offsets. By the final image (4.5 L), visual observation is difficult and the primary face, secondary face, and corner pixels have approximately equal intensity values around 15%, less than half as strong as the 5-meter range test and only a small fraction of the 2-meter range test. This intensity reduction occurs even though the target is only about 2 meters away from the sonar. While offset testing did not identify the root cause of this, one factor could have been the increasingly off-normal relative geometry of the cinder block as offset increased. Despite its relatively weak intensity values, a complete secondary face remains clearly visible in the remaining three data sets, as shown in Figure 5.11.

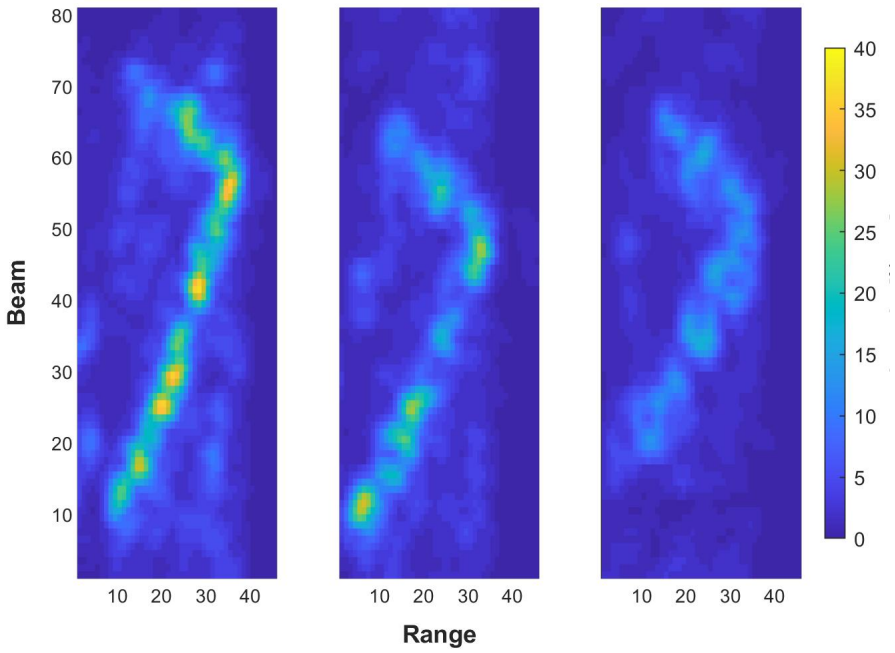

Figure 5.11. From left to right, the 3.0 L, 3.5 L, and 4.5 L offset test sub-region images with the cinder block target.

The offset test results for the glass sphere are shown in Figure 5.12. An obvious takeaway from visual examination of this image sequence is that the shape of the target return does not

change with offset position. This result was expected because of the rotational symmetry of the sphere. It also made it possible to analyze the beam dimension profiles of the returns from all twelve offset tests using the same approach utilized for the range test analysis. Figure 5.13 plots this analysis for each offset test.

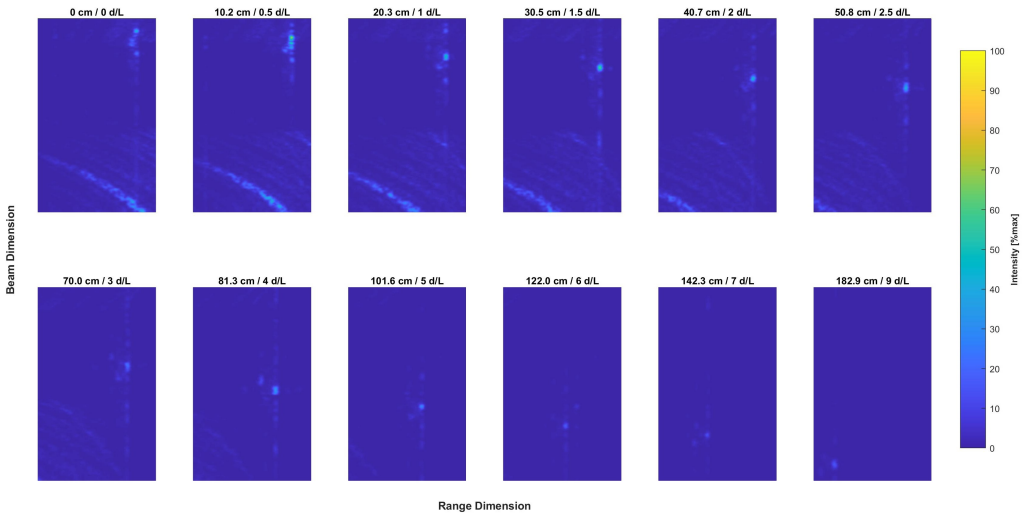

Figure 5.12. Collage of offset test sub-region images with the glass sphere target.

Similar to the cinder block offset tests, the sphere's offset test intensities decrease significantly as offset increases; peak intensity values in the last few images of Figure 5.12 barely exceed 10%, despite the target being only about 2 meters from the sonar. Unlike the cinder block returns, however, these final intensity values are still higher than both the 4- and 5-meter range tests for the sphere.

Figure 5.13 shows that the beam dimension profiles of most of the sphere's data sets mimic the bell curve shapes from the range tests in Figure 5.6. Additionally, their returns span a similar number of beams.

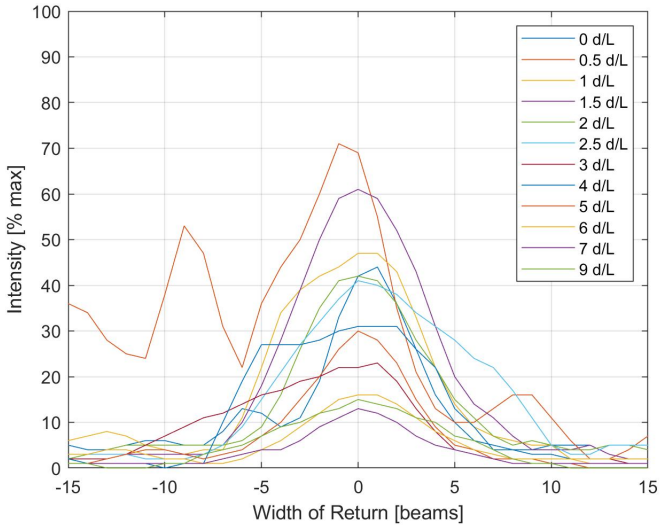

Figure 5.13. Beam dimension plot of the glass sphere at each offset test location.

This plot depicts another expected result: the peak intensities for head-on or nearly head-on data sets (about 60-70%) were higher than those of the 2-meter range test (50%). This higher value is expected because the sphere is closer than 2 m, about 1.5 m. The head-on (0 L offset) peak intensity values for both targets can be predicted by extrapolating the corresponding range test results to 1.5 meters. Table 5.2 compares these predicted values to the actual measurements. The extrapolated plots are shown in Figure A.19 of the Appendix. The consistency of these head-on results suggests that the experimental setup and data collection procedures were repeatable for both the range and offset tests.

Table 5.2. Comparison of measured intensities from the head-on (0 L) offset test with predicted intensities from the range test.

	Intensity [% Max]	
	Predicted	Measured
Cinder Block	71.3	70.0
Glass Sphere	100.0	98.0

Interestingly, offset test positions that were not head-on returned intensities that were significantly lower than the corresponding range test intensities for both targets. Although determining the exact cause of this phenomenon requires further investigation, at least two potential causes exist. First, target placement could have unintentionally introduced a vertical offset. Second, the off-normal aspect of the cinder block may cause sonar signals to deflect away from the sensor instead of reflecting back towards it. Granted, this would not apply to the sphere due to its rotational symmetry. This suggests that the low intensities measured during offset testing were not due to target aspect alone.

Offset Test Summary

Offsetting a FLS from the target allows for successful visual identification of important geometric aspects, such as sides or corners. This suggests that classification is possible from a standalone FLS, with lateral offset, for objects with these features. The shape of FLS returns from objects with rotational symmetry do not change with increased offset. However, few natural objects have this characteristic. Therefore, this result can provide grounds for classification as "mine-like." This would also be consistent with current Navy methods which classify some naturally-occurring objects like spherical coral as mine-like objects.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 6:
Conclusions and Future Work

This research set out to determine the possibility of successfully classifying mine-like objects using only a FLS in support of the Navy's MCM needs in the VSW. The combined mount and CAVR tank proved to be a successful setup in that they provided the ability to collect data in a controlled, methodical manner that supported the range and offset experiments. Data collection, extraction, and processing techniques allowed for qualitative and quantitative analysis to derive conclusions from the range and offset tests.

6.1 Conclusions

The head-on range test outcomes support three different, but related conclusions. First, temporal variation can be neutralized for a static test by averaging through time. This is important for future static tests in a controlled environment but is not necessarily applicable to real-world data sets, where the sensor is usually in motion and range to a target is constantly changing. The second conclusion is that target locations can be precisely determined by a large peak in intensity that correlates with the center of the target, especially for a target with a convex surface (such as a sphere). Lastly, the geometrically expected number of beams spanned by a target does not generally correlate with the measured width of a sonar return. The one exception observed during range testing was for a flat-faced target placed perpendicular to, and sufficiently far away from, the sonar. Therefore, assuming a detected target had a flat face and known target strength, a reasonable size estimate could be made in support of classification. In all other cases, however, target size and shape determination in support of classification cannot be determined from an FLS alone.

Additionally, the 2-meter head-on range test included too many spurious (i.e., false) returns to be able to discern the actual target within the sonar image. This result indicates that although getting closer to a target will provide a stronger return, AUV path planning must account for avoiding false positives from spurious returns when too close to the target (e.g., about 3 meters for a cinder block).

The offset test revealed that target geometry and apparent shape changes are observable for

objects without rotational symmetry. However, quantifying the apparent shape change is difficult and this research was unable to correlate this change to appreciable intensity differences between different sides of a rectangular target. Additionally, detection of rotational symmetry could be grounds for mine-like classification. Since naturally-occurring objects rarely have the same apparent shape from all perspectives, rotational symmetry is highly indicative of a man-made object.

6.2 Future Work

Multiple aspects of this research were identified as either areas to improve or expand upon through future work.

6.2.1 Vehicle Dynamics

A study should be conducted to evaluate the impact that dynamic disturbances in the operating environment have on sensor performance. The VSW zone subjects vehicles to inconsistent waves, currents, and water depths, all of which can affect AUV sensor performance. The study should first characterize the dynamics of disturbances in the operating environment before analyzing the sensor performance of vehicles subjected to those disturbances.

6.2.2 Water Clarity

Camera testing in the CAVR tank could only be conducted in ideal, clear water conditions. Studies should be conducted to determine the degradation of trained object detection algorithm performance as water clarity decreases and relate those findings to expected conditions in the MIW zones of interest.

6.2.3 Experimental Setup Improvement

The setup utilized in this research allowed for the successful execution of experimentation as designed. However, as with any research, aspects of the setup can be improved upon to reduce error and uncertainty. The primary sources of uncertainty in this experimentation were the unknown acoustic properties of the CAVR tank and the precision of sensor and target placement.

CAVR Tank Limitations

Conducting similar range and offset tests for the same targets and FLS in a larger tank, whose walls and bottom are not within range of the sonar head, would be beneficial. Alternatively, an acoustic test facility designed to eliminate echoes could be utilized. These tests could determine which findings of this research directly resulted from the test environment. This would quantify the viability of future sonar-related research in the CAVR tank and potentially establish a noise baseline for future data sets.

Target Selection and Placement

The targets used in this experimentation proved to be good choices, as they showed how shape and size affected sonar returns. Future research should utilize targets with comparable size and materials to expected mine-like objects. These targets would produce returns that more closely resemble objects likely to be encountered during real-world missions.

Future research with the experimental setup described in this thesis would benefit from improved target placement. Instead of suspending targets with ropes, an apparatus should be designed to accurately control the position, depth, and aspect of targets placed within the tank. The apparatus should be easily maneuverable and should have minimal to no sonar return. These improvements would offer a more controlled and precise testing environment. The offset test, in particular, would greatly benefit from increased target control. The sonar (with a "head-on" aspect) and targets could maintain the same positions within the tank while only the aspect of the target is changed. This would ensure intensity changes are due to target aspect alone.

Data Processing and Analysis Automation

Future effort should be committed to automating the data processing and analysis of sonar test data. The manual methods used in this thesis to "hard-code" image isolation and intensity analysis is time-consuming, repetitive, and susceptible to minor errors that skew data. Despite the challenges outlined in Chapter 4, modern computer vision techniques can be leveraged to accomplish this goal. The ability to automatically collect and process data, without human operator intervention, would greatly accelerate future research in this area. It would also provide ample data for development of machine learning and/or artificial intelligence algorithms.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX: Supplementary Figures

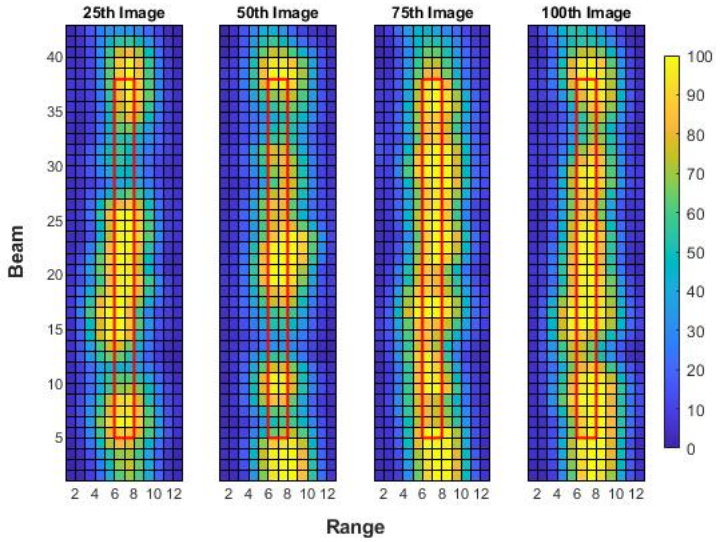

Figure A.1. Cinder block 2-meter range test image progression.

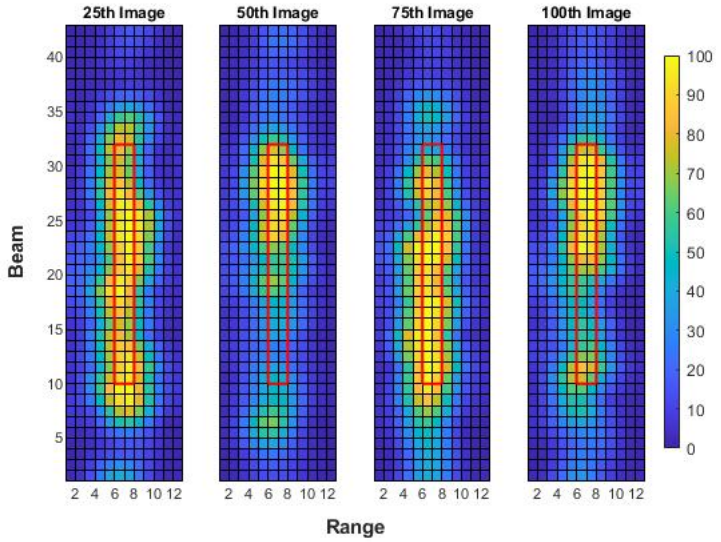

Figure A.2. Cinder block 3-meter range test image progression.

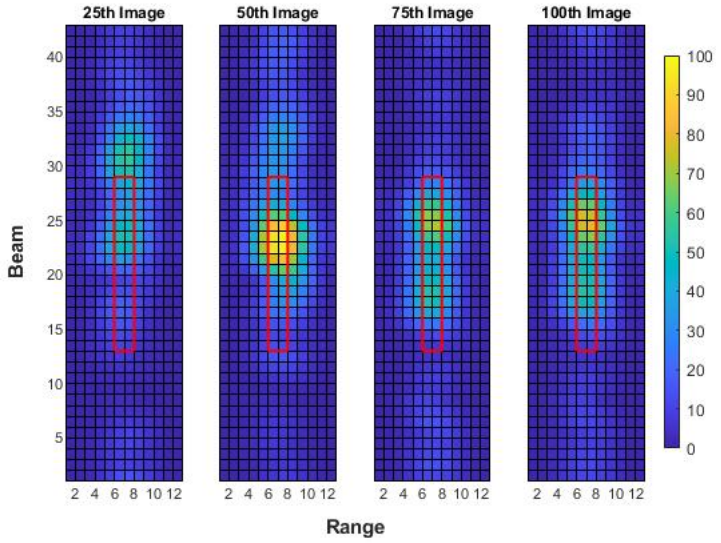

Figure A.3. Cinder block 4-meter range test image progression.

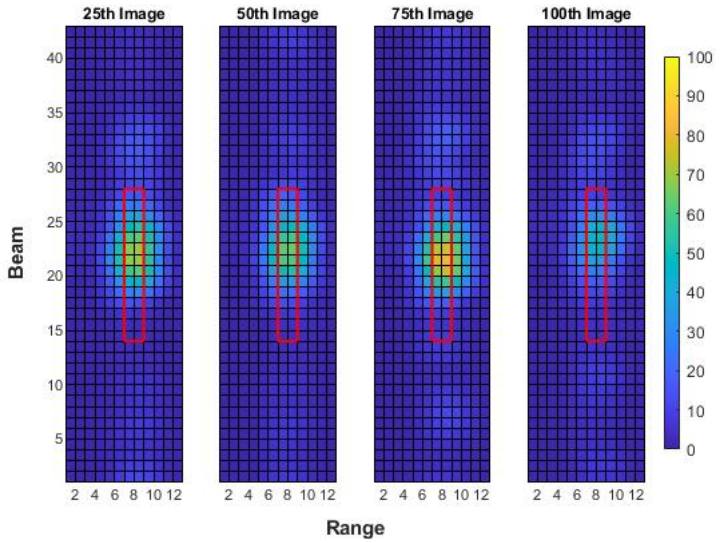

Figure A.4. Cinder block 5-meter range test image progression.

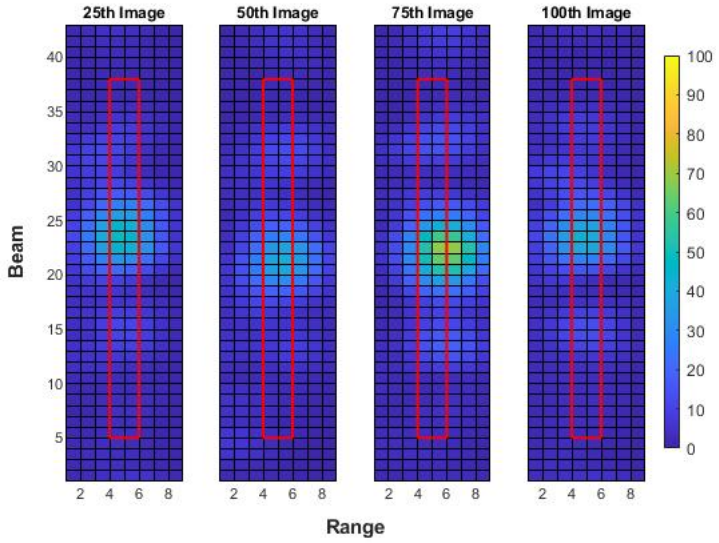

Figure A.5. Sphere 2-meter range test image progression.

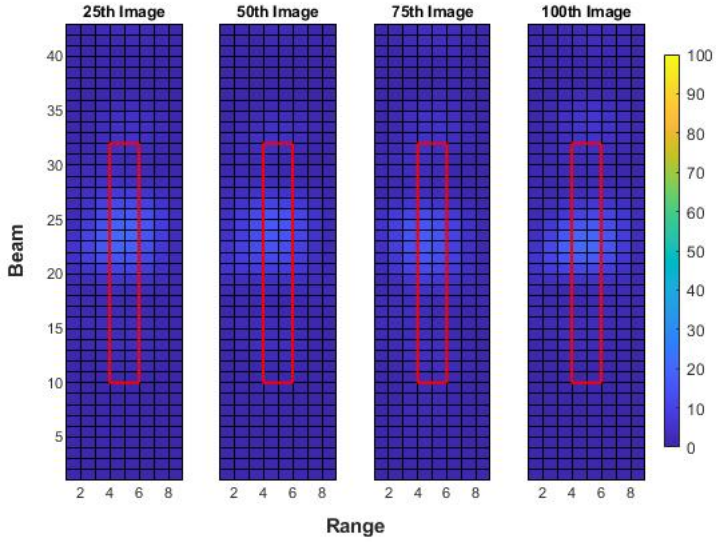

Figure A.6. Sphere 3-meter range test image progression.

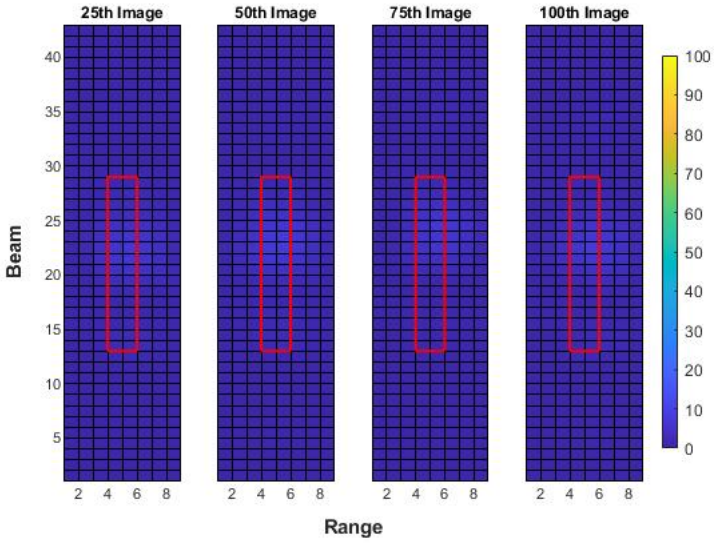

Figure A.7. Sphere 4-meter range test image progression.

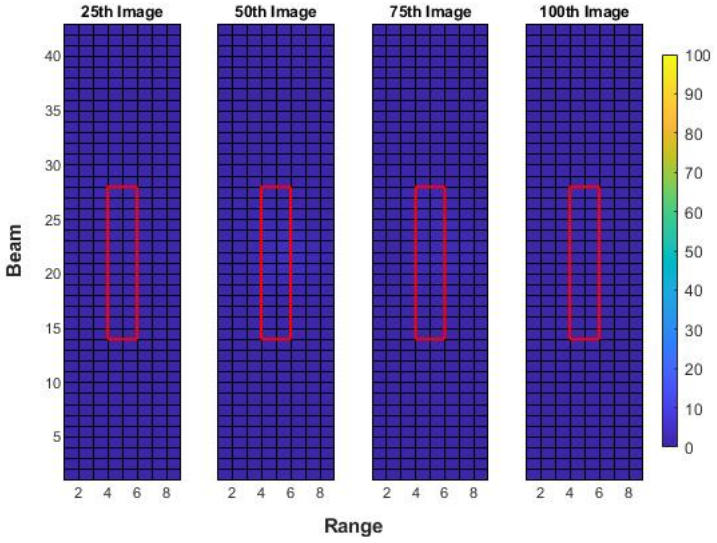

Figure A.8. Sphere 5-meter range test image progression.

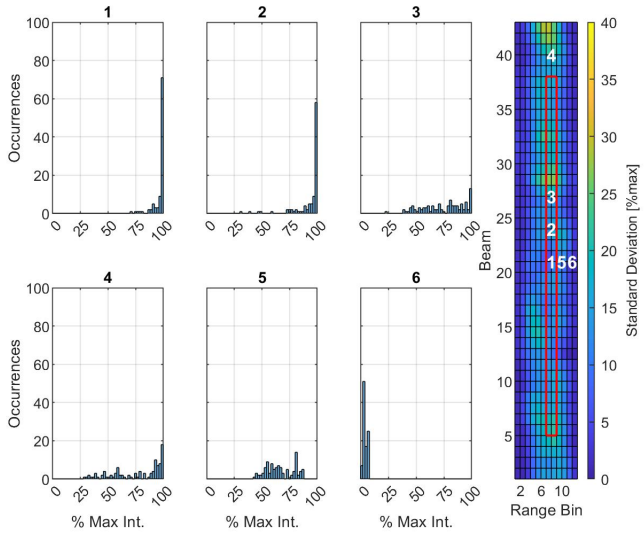

Figure A.9. Cinder block 2-meter range test standard deviation plot with histograms from six points within the image.

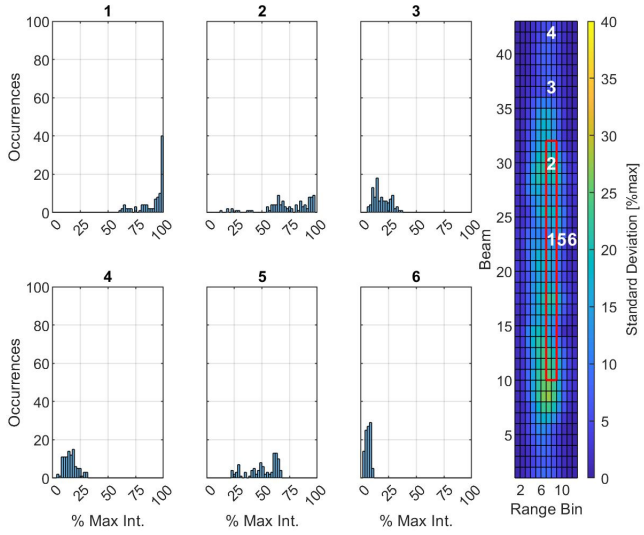

Figure A.10. Cinder block 3-meter range test standard deviation plot with histograms from six points within the image.

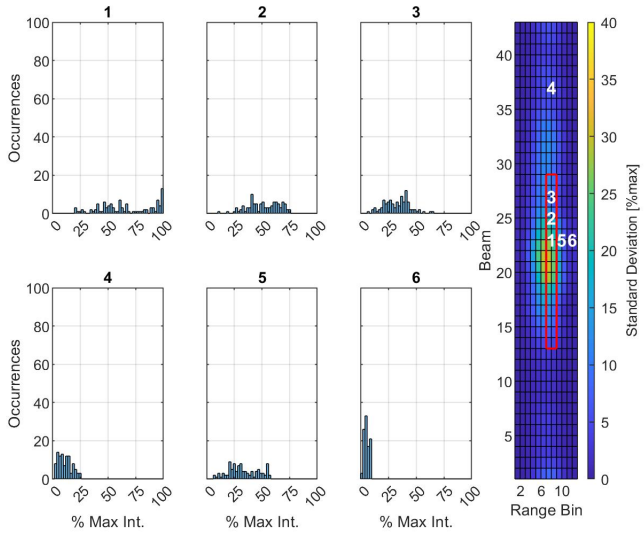

Figure A.11. Cinder block 4-meter range test standard deviation plot with histograms from six points within the image.

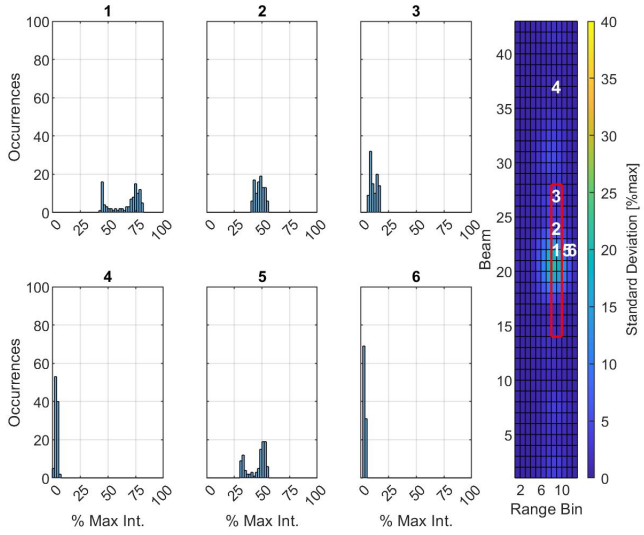

Figure A.12. Cinder block 5-meter range test standard deviation plot with histograms from six points within the image.

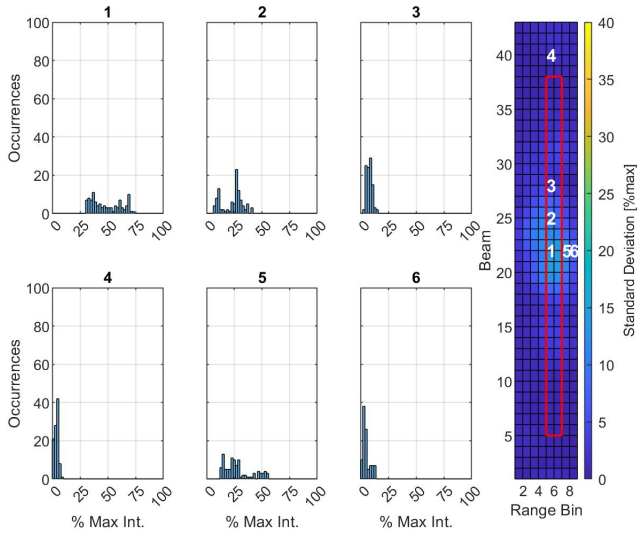

Figure A.13. Sphere 2-meter range test standard deviation plot with histograms from six points within the image.

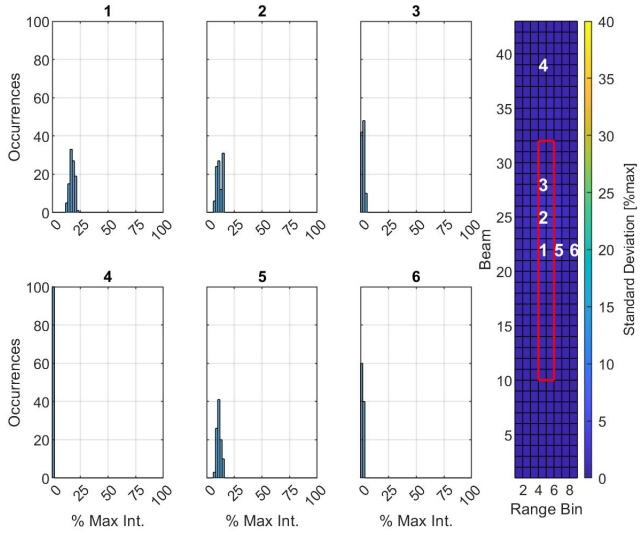

Figure A.14. Sphere 3-meter range test standard deviation plot with histograms from six points within the image.

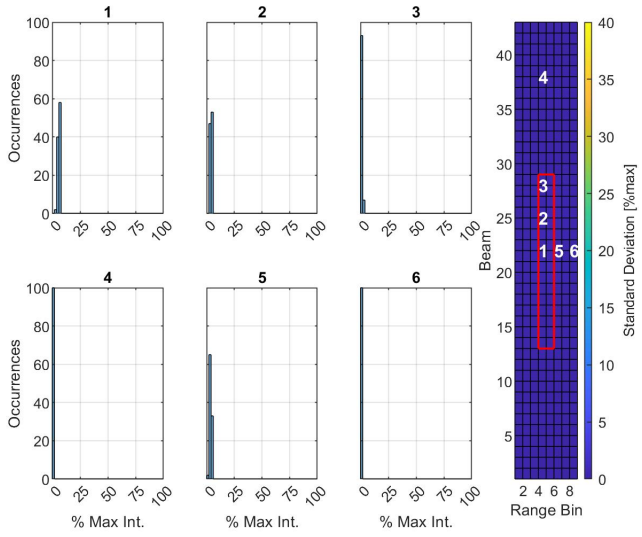

Figure A.15. Sphere 4-meter range test standard deviation plot with histograms from six points within the image.

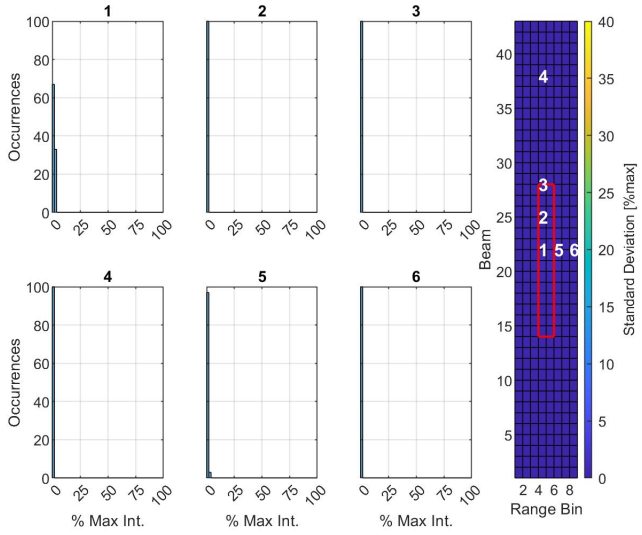

Figure A.16. Sphere 5-meter range test standard deviation plot with histograms from six points within the image.

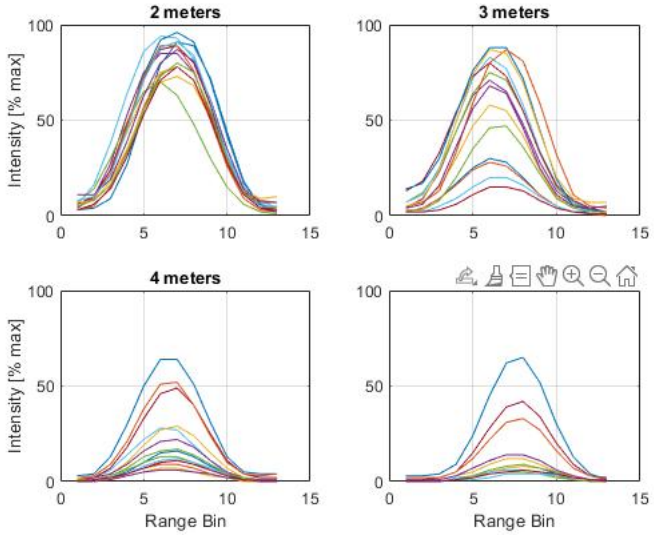

Figure A.17. Cinder block range dimension plots for range location.

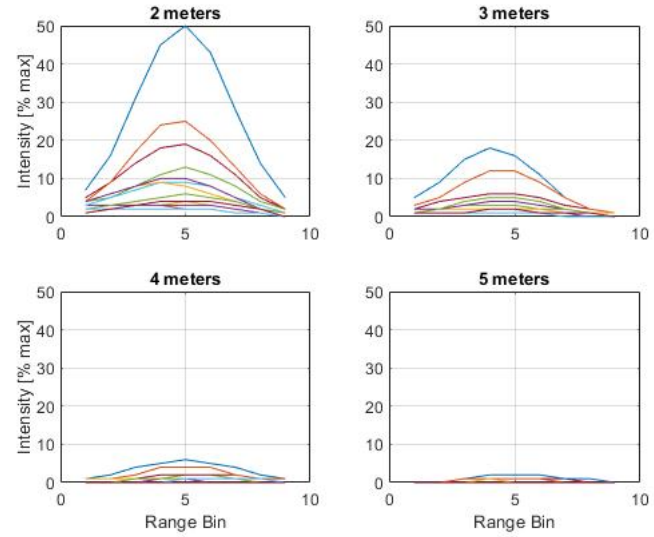

Figure A.18. Sphere range dimension plots for range location.

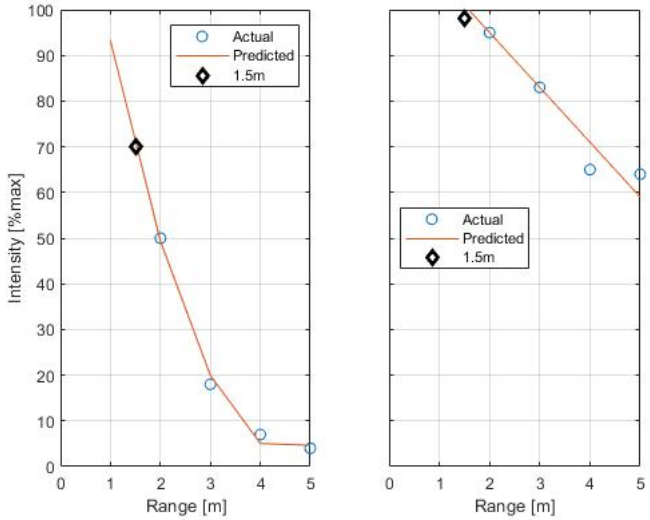

Figure A.19. Predicted sphere (left) and cinder block (right) intensities at 1.5 meters.

List of References

[1] *Department of Defense Dictionary of Military and Associated Terms*, Joint Publication 1-02, Department of Defense, Washington, DC, 2010. [Online], pp. 11, 12, 38, 67, 155.

[2] N. History and H. Command, "Naval Mine Warfare," U.S. Navy, Jul 2021. [Online]. Available: https://www.history.navy.mil/browse-by-topic/exploration-and-innovation/naval-mine-warfare.html

[3] "*USS Samuel B. Roberts (FFG-58)*," *Wikipedia*. Accessed Jan 18, 2022. [Online]. Available: https://en.wikipedia.org/wiki/USS_Samuel_B._Roberts_(FFG-58)

[4] *Barriers, Obstacles, and Mine Warfare for Joint Operations*, Joint Publication 3-15, Department of Defense, Washington, DC, 2018. [Online], pp. B–14, B–24.

[5] Naval Sea Systems Command. "Mine Countermeasures Ship - MCM." Accessed Jan 20, 2022. [Online]. Available: https://www.navy.mil/Resources/Fact-Files/Display-FactFiles/Article/2171622/mine-countermeasures-ships-mcm/

[6] A. Magnuson, "Navy Minesweepers look to A.I. to boost speed, reduce risk," Feb 2022. [Online]. Available: https://www.nationaldefensemagazine.org/articles/2022/2/9/navy-minesweepers-look-to-ai-to-boost-speed-reduce-risk

[7] N. R. Council, E. C. on Geosciences, O. S. Board, and N. A. of Sciences, *Oceanography and Mine Warfare*, 7th ed. Washington, D.C., USA: National Academies Press, 2000. [Online].

[8] M. Lawlor, "Navy dives deeper into mine countermeasures," Signal, Jul 2001. [Online]. Available: https://www.afcea.org/content/navy-dives-deeper-mine-countermeasures

[9] M. Aylward, "Mines in the Surf Zone: A proposed Breaching Concept," M.S. thesis, Naval Postgraduate School, Monterey, CA, USA, 1994. [Online]. Available: https://apps.dtic.mil/sti/citations/ADA285184

[10] S. Smith, E. An, R. Christensen, J. Kloske, S. Snowden, D. Kronen, and L. Marquis, "Results of an Experiment using AUVs for Shallow-Water Mine Reconnaissance"," *Proceedings of SPIE*, vol. 3711, no. 1, pp. 132–172, Jul 1999. [Online].

[11] J. Topple and J. Fawcett, "MiNet: Efficient Deep Learning Automatic Target Recognition for Small Autonomous Vehicles," *IEEE Geoscience and Remote Sensing Letters*, vol. 18, no. 6, pp. 1014–1018, Jun 2021. [Online].

[12] A. Houck, Private communication, Mar 2021.

[13] N. Hurtos, N. Palomeras, S. Nagappa, and J. Salvi, "Automatic detection of underwater chain links using a forward-looking sonar," *2013 MTS/IEEE OCEANS - Bergen*, Jun.

[14] T. Zhang, L. Shuwei, X. He, H. Huang, and K. Hao, "Underwater Target Tracking Using Forward-Looking Sonar for Autonomous Underwater Vehicles," *Sensors*, vol. 1, no. 102, Dec 2019. [Online].

[15] A. Healey, D. Horner, and S. Kragelund, "AUV experiments in obstacle avoidance," in *Proceedings of the IEEE Oceans 2005 Conference*, 2005. [Online]. Available: https://ieeexplore.ieee.org/document/1639962

[16] C. Mishra and D. L. Gupta, "Deep machine learning and neural networks: an overview," *IAES international journal of artificial intelligence*, vol. 6, no. 2, June 2017. [Online]. doi: 10.11591/ijai.v6.i2.

[17] R. Ren, R. Hung, and K. C. Tan, "A Generic Deep-Learning-Based Approach for Automated Surface Inspection," *IEEE Transactions on Cybernetics*, vol. 48, no. 3, pp. 929–940, Mar 2018. [Online].

[18] R. Girshick, J. Donae, T. Darrell, and J. Malik, "Region-Based Convolutional Networks for Accurate Object Detection and Segmentation," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 1, Jan.

[19] MathWorks. "Convolutional Neural Network" Accessed Mar. 7, 2022. [Online]. Available: https://www.mathworks.com/discovery/convolutional-neural-network-matlab.html

[20] M. Kolsch, Private communication, Dec 2021.

[21] P. Ribeiro, M. dos Santos, P. Drews, and S. Botelho, "Forward looking sonar scene matching using deep learning," in *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2017. [Online]. Available: https://ieeexplore.ieee.org/document/8260693

[22] T. Xiao, Z. Cai, C. Lin, and Q. Chen, "A shadow capture deep neural network for underwater forward-looking sonar image detection," *Mobile Information Systems*, vol. 2021, p. 10, Dec 2021. [Online]. doi: 10.1155.2021.3168464.

[23] Z. Fan, X. Liu, W. Xia, and H. Li, "Detection and segmentation of underwater objects from forward-looking sonar based on a modified mask RCNN," *Signal, Image, and Video Processing*, vol. 15, Sep.

[24] “REMUS 100,” Woods hole Oceanographic Institution, 2022. [Online]. Available: https://www2.whoi.edu/site/osl/vehicles/remus-100/

[25] “vLBC300,” Seascape Subsea BV, 2022. [Online]. Available: https://www.seascapesubsea.com/product/vlbv300/

[26] Teledyne BlueView. “M900-2250 Dual Frequency Series” Accessed Jan. 10, 2022. [Online]. Available: http://www.teledynemarine.com/M900-2250/

[27] Teledyne BlueView, *Dual-Frequency Sonar*, M900-2250-130 Series, 2019. [Online]. Available: http://www.teledynemarine.com/M900-2250/

[28] Basler AG. “2021 Basler ace specifications”. [Online]. Available: https://www.baslerweb.com/en/products/cameras/area-scan-cameras/ace/aca1300-60gm/. Accessed Sep. 24, 2021.

[29] Basler. “cA1300-60gm Basler ace” Accessed Jan. 10, 2022. [Online]. Available: https://www.baslerweb.com/en/products/cameras/area-scan-cameras/ace/aca1300-60gm/

[30] XSENS. “MTi-100 IMU”. [Online]. Available: https://www.xsens.com/mti-100. Accessed Nov. 23, 2021.

[31] MathWorks. “Get Started with ROS” Accessed Dec. 10, 2021. [Online]. Available: https://www.mathworks.com/help/ros/ug/get-started-with-ros.html

[32] Echoview Software. “About Us” Accessed Feb. 7, 2022. [Online]. Available: https://echoview.com/about-us/

[33] MathWorks. "ROS Log Files" Accessed Nov. 23, 2021. [Online]. Available: https://www.mathworks.com/help/ros/ug/ros-log-files-rosbags.html

[34] MathWorks. “Timeseries” Accessed Nov. 23, 2021. [Online]. Available: https://www.mathworks.com/help/matlab/ref/timeseries.html?s_tid=doc_ta

[35] T. Foote and M. Purvis. “Standard units of measure and coordinate conventions” Accessed Mar. 1, 2022. [Online]. Available: https://www.ros.org/reps/rep-0103.html

[36] SPIE. “Filling the gaps” Accessed Mar. 7, 2022. [Online]. Available: https://spie.org/news/filling-the-gaps?SSO=1

[37] University of Southhampton. “Computer Vision Demonstration Website” Accessed Mar. 7, 2022. [Online]. Available: https://www.southampton.ac.uk/~msn/book/new_demo/corners/

THIS PAGE INTENTIONALLY LEFT BLANK

Initial Distribution List

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California