SYNTHETIC TRAINING DATA GENERATION FOR SIDE-SCAN SONAR

STUDENTS: AMOGH BANTWAL, ANDERS PEARSON, ANKITH TUNUGUNTLA, IAN KINSELLA, KYLE ORTH

Introduction

- Side scan sonar (SSS) images the seafloor using sound pulses
- Autonomous Underwater Vehicle (AUV) emits pulses, records echoes
- SSS is widely used in search & rescue, pipeline inspections, and environmental monitoring

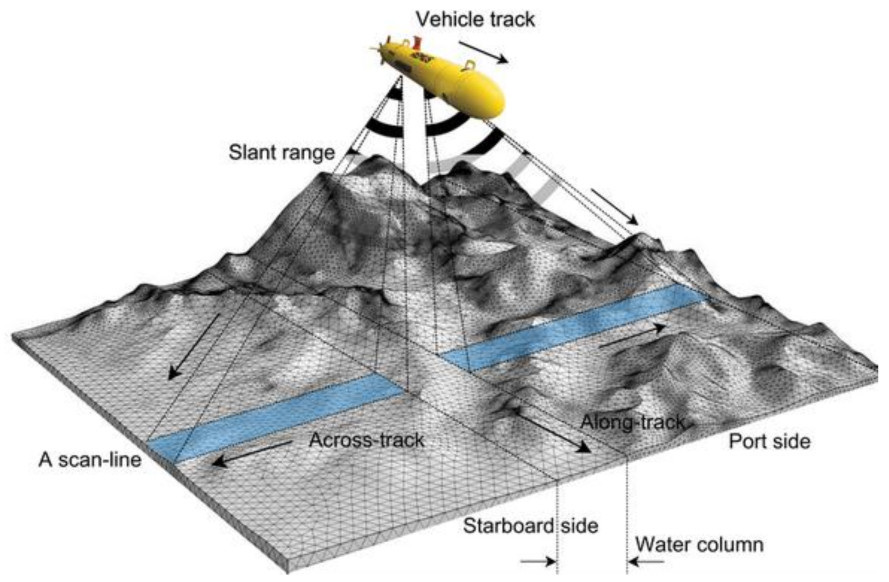

Figure 1: SSS system mounted on an AUV

Objective & Motivation

- **Objective**: Create a highly customizable system for mass-generating synthetic SSS images
- Model training requires large volume of diverse, labeled data [1]
- Easily generate images of underrepresented targets (wrecks, mines, etc.)
- Synthetic data reduces need for costly, time-consuming data collection
  - An AUV costs **~$300K**
  - Imaging for 8 hours costs **$900**

Figure 2: EdgeTech AUV used by Booz Allen Hamilton

Solution

- We present a scalable synthetic SSS data generation system
- Realistic, labeled SSS images using the Unity game engine
- Highly customizable (simulation settings, 3D assets)
- Post-processing follows real SSS processes
- Can generate 1000s of images ***in a few hours at low cost***

Unity

System Architecture

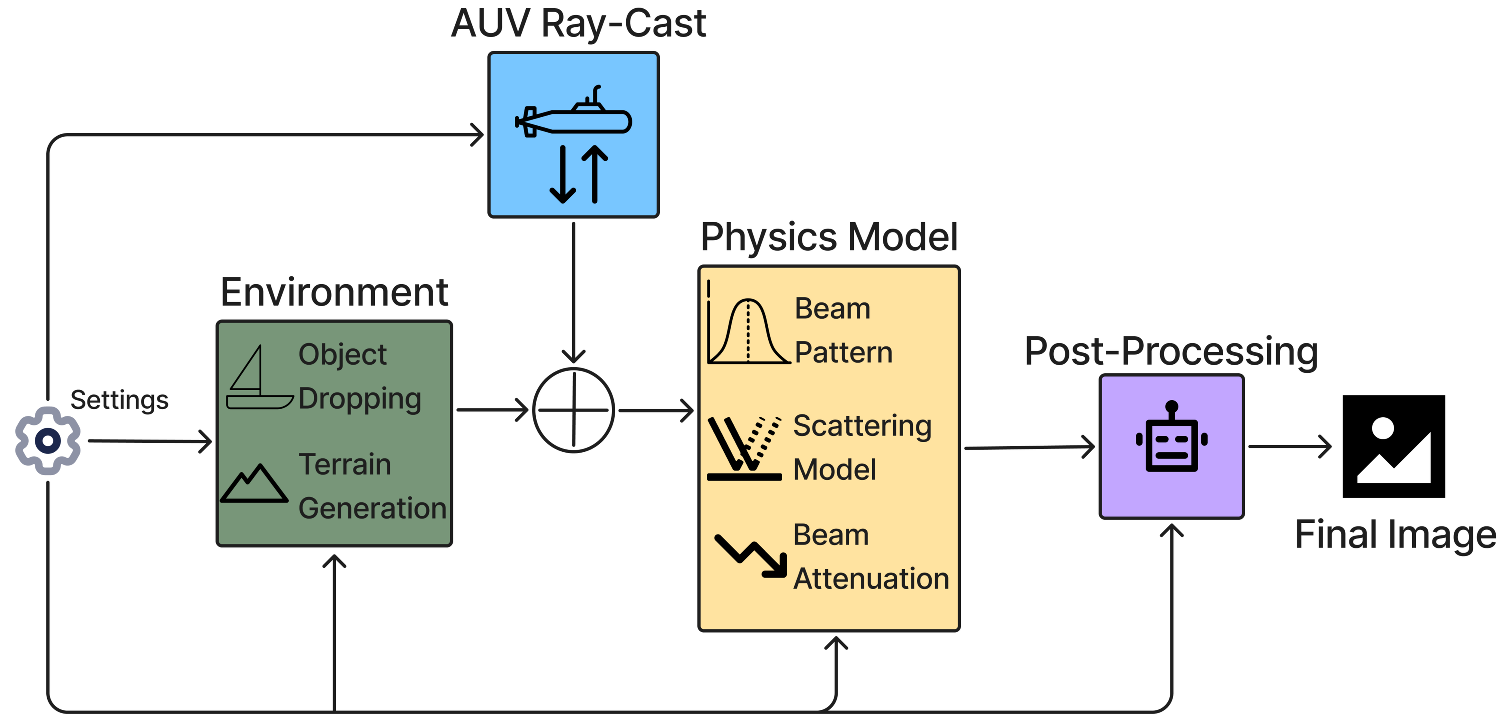

Figure 3: Flowchart of synthetic SSS data generation pipeline

Simulation Setup

STEP 1: Initialize simulation parameters (seabed, sonar physics, image rendering)
STEP 2: Generate seabed and drop objects
STEP 3: Fly AUV over seabed while casting rays, usually >2000 rays per ping
STEP 4: Collect backscatter value of each ray into large arrays
STEP 5: Apply post processing to raw value arrays and render into labeled images

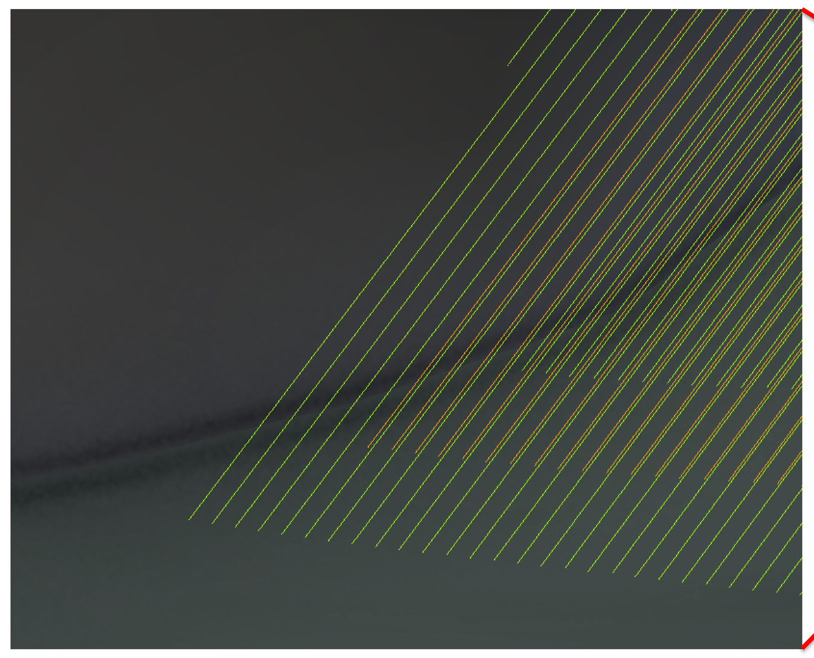

Figure 4: Closeup of rays hitting an object

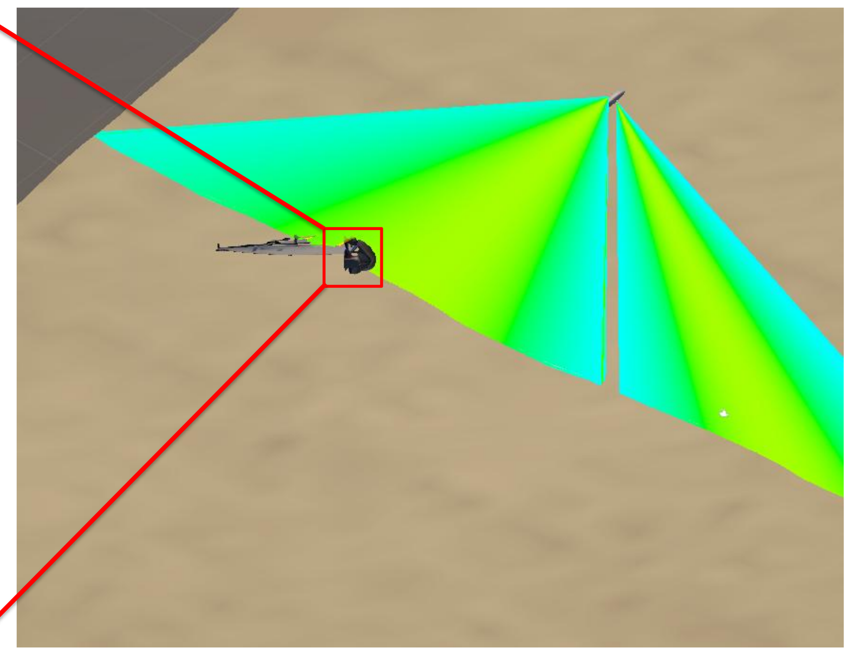

Figure 5: The AUV casting rays in Unity

Acoustic Modeling

APL's Jackson Backscatter Model [2] predicts seabed acoustic reflectivity; grazing angle and sediment properties determine how much energy is reflected vs absorbed.
Francois-Garrison Model accounts for frequency-dependent signal attenuation through the water column.
AUV Beam Pattern models transducer direction-dependent sensitivity.

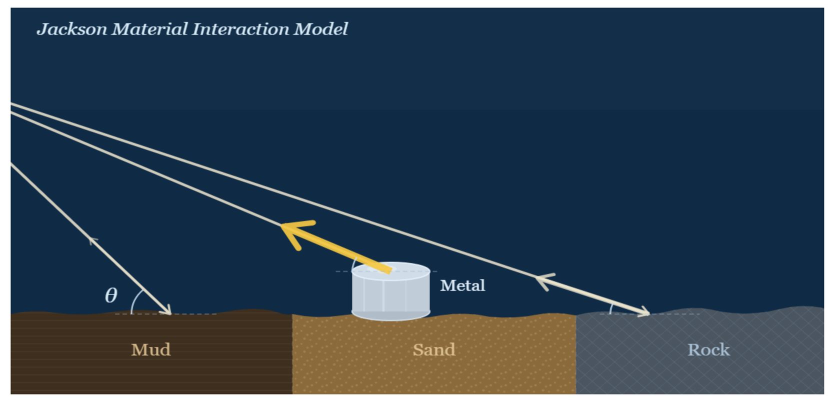

Figure 6: Ray intensities from backscatter model

Results

- YOLO object detection model was trained on a combination of real and synthetic data
- Overall recognition accuracy saw minor improvement when trained with synthetic data
- Synthetic data taught model to recognize objects with little to no real data
- System created 1000s of accurate labeled SSS images

Real SSS

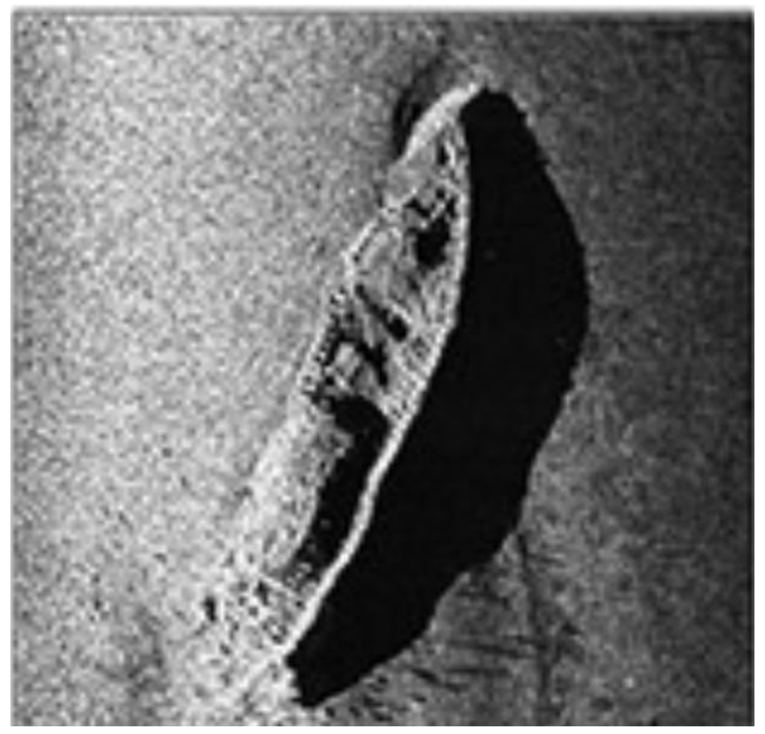

Synthetic SSS

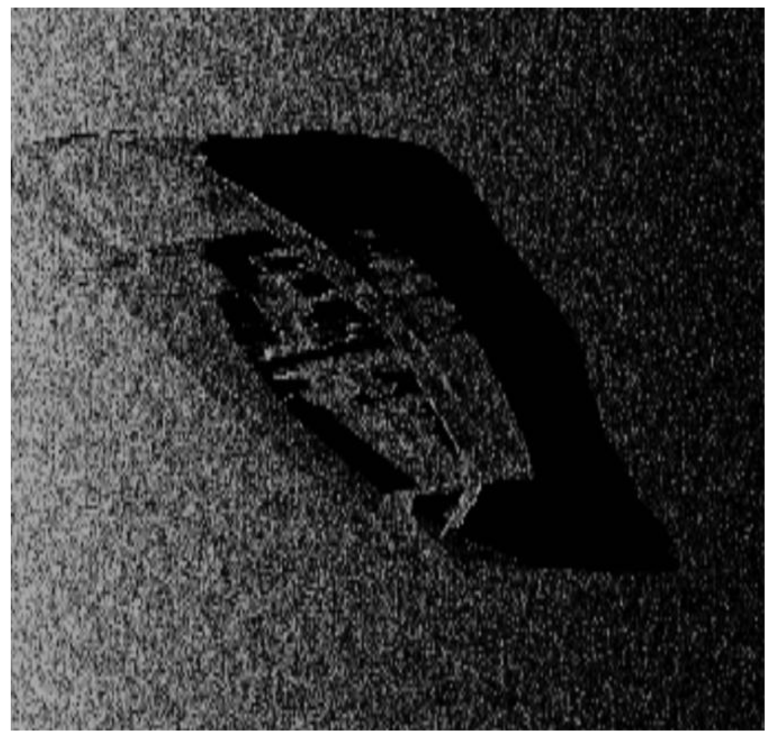

Object	Images (#)	Real mAP50	Real + Synthetic mAP50
People	2,457	0.988	0.98
Grenade	780	0.327	0.349
Ships	2,458	0.242	0.223
All Classes	25,986*	0.452	0.453

*Not every object that we trained on was inside test set due to insufficient real SSS with objects
mAP = Mean Average Precision

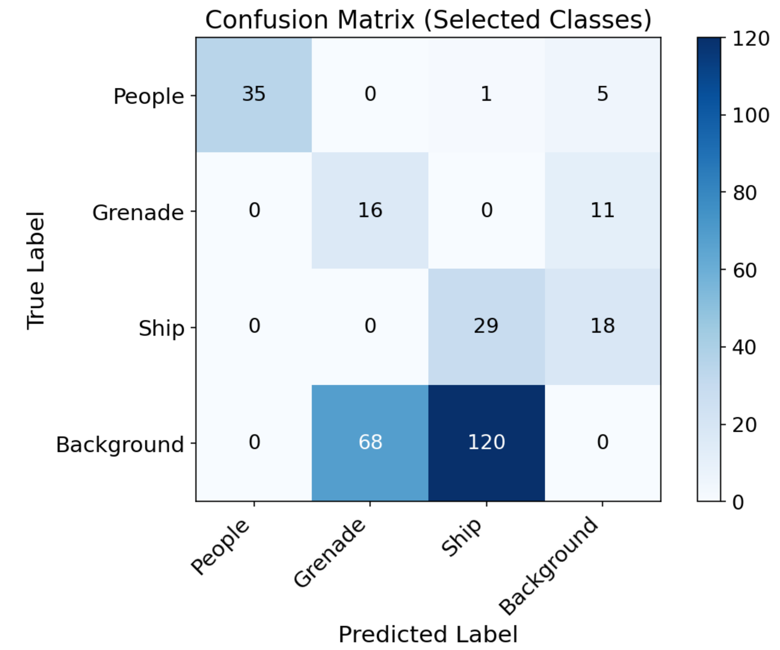

Scan for BAH Side Scan Sonar Footage!

[1] I. Apeināns, "Optimal size of agricultural dataset for YOLOv8 training," in *Environment. Technologies. Resources. Proceedings of the International Scientific and Practical Conference*, vol. 2, 2024, pp. 21–24
[2] Jackson, Darrell. (2000). High-Frequency Bistatic Scattering Model for Elastic Seafloors. 45.

ELECTRICAL & COMPUTER ENGINEERING
UNIVERSITY of WASHINGTON

ADVISERS: NOAH PARKER, VASILEIOS CHARISOPOULOS, VICTOR MARCENAC
SPONSOR: BOOZ ALLEN HAMILTON