CHALMERS
UNIVERSITY OF TECHNOLOGY

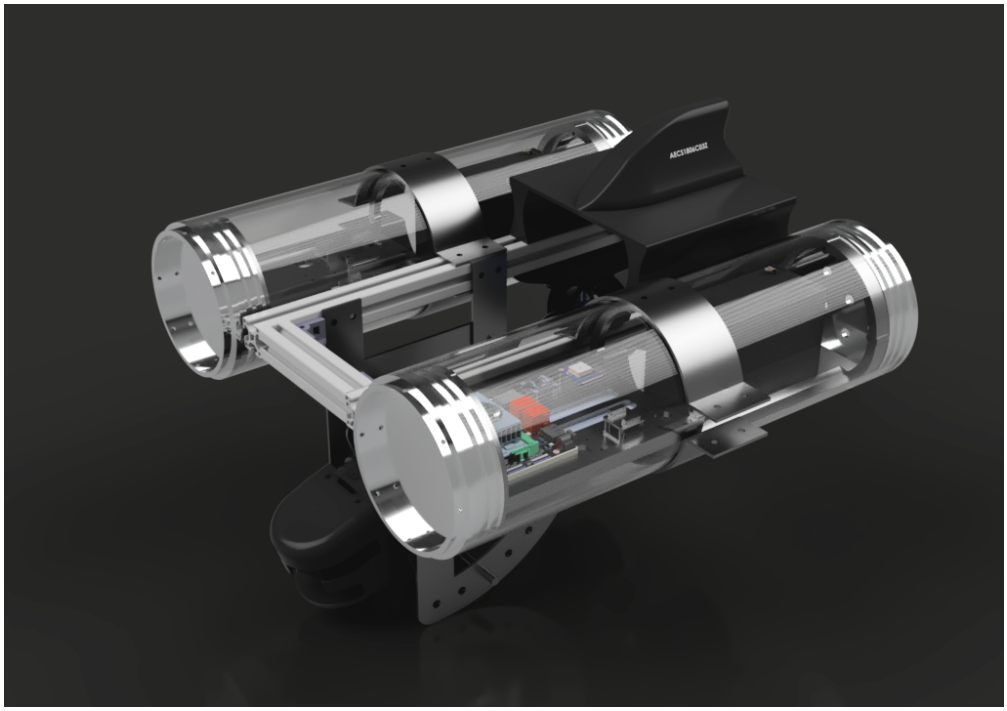

Sonar-equipped unmanned surface vehicle for search and rescue operations

Master's thesis in Complex adaptive systems

GABRIEL FORSBERG

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's Thesis in Complex Adaptive Systems

Sonar-equipped unmanned surface vehicle for search and rescue operations

GABRIEL FORSBERG

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology
Gothenburg, Sweden 2026

Sonar-equipped unmanned surface vehicle for search and rescue operations
GABRIEL FORSBERG

Supervisor: Tarun Kadri Sathiyan, Department of Mechanics and Maritime Sciences
Examiner: Ola Benderius, Department of Mechanics and Maritime Sciences

Master's Thesis 2026
Department of Mechanics and Maritime Sciences
Chalmers University of Technology
SE-412 96 Gothenburg
Sweden
Telephone +46 31 772 1000

Cover: Render of the revised Seadragon USV platform.

Typeset in LaTeX
Gothenburg, Sweden 2026

Sonar-equipped unmanned surface vehicle for search and rescue operations
GABRIEL FORSBERG
Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology

Abstract

Drowning remains a major public safety challenge and accounts for hundreds of thousands of fatalities globally, with survival rates heavily dependent on rapid response. Current underwater search and rescue (SAR) methods, including rescue divers, remotely operated vehicles, and side-scan sonar systems, are often resource intensive, slow to deploy, or otherwise operationally limited. This thesis investigates whether a lightweight, rapidly deployable and sonar-equipped unmanned surface vehicle (USV) can provide a faster and more efficient solution for underwater SAR operations. Building upon the ongoing *Seadragon* initiative within the Division of Vehicle Engineering and Autonomous Systems at Chalmers University of Technology, a redesigned USV platform (Seadragon 2.0) was developed with improved structural durability and sensing capability. The prototype integrates a Teledyne BlueView M900 Mk2 multibeam imaging sonar and an NVIDIA Jetson Orin Nano edge-computer within a compact catamaran-style hull architecture. To enable automatic underwater target detection, a custom sonar dataset was collected using recordings of submerged humans and simulated underwater debris. A two-stage transfer learning pipeline based on the YOLO26n architecture was trained both on the custom dataset and a larger public sonar dataset. While maintaining a real-time inference performance of 39.1 FPS on the embedded platform, the resulting model achieved a precision of 0.868, recall of 0.812, and mAP of 0.859. In addition, the redesigned Seadragon prototype achieved an estimated theoretical search coverage rate of 244.3 m^2/s, exceeding that of conventional rescue divers and comparable side-scan sonar vessel operations. Although full-scale field validation was limited by project time constraints, the results demonstrate the feasibility of the lightweight USV concept for enhancing underwater SAR operations and, critically, reducing time to rescue.

Keywords: unmanned surface vehicle (USV), search and rescue (SAR), imaging sonar, underwater target detection, machine learning, autonomous systems, acoustic imaging, human detection, edge computing.

Preface

This report presents the outcome of my master's thesis project conducted within the Complex adaptive systems programme at Chalmers University of Technology during the spring of 2026. The work is a continuation of an experimental USV search and rescue initiative at the Division of Vehicle Engineering and Autonomous Systems known as *Seadragon*.

Acknowledgements

I would like to express my sincere gratitude to my supervisor, Tarun Kadri Sathiyan, for his active guidance and encouragement during the development of this project. I am also grateful to my examiner, Ola Benderius, for providing invaluable technical and academic input throughout the duration of the work. Finally, I wish to thank the excellent staff of Chalmers Prototype Laboratory—Anders Mattsson, Patrik Wåhlin, and Joakim Danielsson—who kindly offered their technical expertise during the construction of the prototype.

Gabriel Forsberg, Gothenburg, June 2026

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

ADVD	Acoustic drowning victim dataset
AI	Artificial intelligence
AUV	Autonomous unmanned vehicle
CNN	Convolutional neural network
COCO	Common objects in context
DoF	Degrees of freedom
EKF-SLAM	Extended kalman filter simultaneous localization and mapping
ESC	Electric speed controller
FHSS	Frequency hopping spread spectrum
FOV	Field of view
GNSS	Global navigation satellite systems
GPS	Global positioning system
GUI	Graphical user interface
IMU	Inertial measurement unit
IPC	Inter-process communication
ISC	Inter-service communication
ML	Machine learning
MVP	Minimum viable product
PWM	Pulse width modulation
R-CNN	Region-based convolutional network
ResNet	Residual network
ROV	Remotely operated vehicle
SAR	Search and rescue
SLS	Swedish life saving society
TOPS	Trillion operations per second
TCP	Transmission control protocol
UATD	Underwater acoustic target detection
UDP	User datagram protocol
UI	User interface
USV	Unmanned surface vehicle
WHO	World Health Organization
YOLO	You only look once

Nomenclature

Below is the nomenclature of constants, parameters, and variables used throughout this report.

Constants

g	Acceleration due to gravity ($9.81\,\mathrm{m/s^2}$)
ρ_{Al}	Density of aluminium ($2700\,\mathrm{kg/m^3}$)
ρ_{water}	Density of fresh water ($1000\,\mathrm{kg/m^3}$)

Parameters

A	Projected frontal area of the craft
C_d	Hydrodynamic drag coefficient
ID	Inner diameter of acrylic tube
L	Length of acrylic tube
OD	Outer diameter of acrylic tube
r	Radius of acrylic tube
W_{swath}	Effective swath width of the sonar

Variables

B_{one}	Buoyant force of a single sealed tube
C_{rate}	Theoretical search coverage rate
F_d	Hydrodynamic drag force
m_{air}	Mass of a component in air (dry weight)
m_{wdisp}	Mass of water displaced by an object
v_{max}	Operational velocity of the USV
V_{one}	Total volume of one sealed tube
V_{wall}	Material volume of the acrylic tube wall

Contents

List of Acronyms **ix**

Nomenclature **xi**

List of Figures **xv**

List of Tables **xvii**

1 Introduction **1**
1.1 Problem statement . . . 2
1.2 Research questions . . . 2
1.3 Limitations . . . 2

2 Background **3**
2.1 Current state of maritime SAR . . . 3
2.2 Existing portable ROV and USV platforms . . . 4
2.3 Imaging sonar technology . . . 5
2.3.1 Challenges in underwater acoustic imaging . . . 6
2.4 Machine learning for target detection . . . 6
2.5 Emerging sonar SAR capabilities . . . 7

3 Methods **9**
3.1 System architecture and design . . . 9
3.1.1 Hardware redesign and prototyping . . . 10
3.1.2 Propulsion and power system . . . 13
3.1.3 Teledyne multibeam imaging sonar . . . 14
3.1.4 Computer and sensor integration . . . 15
3.1.5 Software architecture and development environment . . . 16
3.1.6 Remote control . . . 17
3.2 Machine learning pipeline . . . 18
3.2.1 Dataset acquisition and augmentation . . . 19
3.2.2 Model architecture and training . . . 21

4 Results **23**
4.1 Final assembled prototype . . . 23
4.1.1 Mass and buoyancy analysis . . . 24
4.2 Search coverage and operational time . . . 25

4.2.1 Operational speed . . . 25
4.2.2 Operational endurance . . . 25
4.2.3 Theoretical search coverage . . . 26
4.2.4 Comparison to conventional methods . . . 27
4.3 Detection and localization accuracy . . . 27
4.4 Computational efficiency . . . 30

5 Discussion **31**
5.1 Effectiveness of the lightweight USV concept . . . 31
5.2 Machine learning performance . . . 32
5.2.1 Limits and improvements of the deployed model . . . 33
5.3 Future work . . . 34
5.3.1 Enhanced autonomy and localization . . . 34
5.3.2 Connectivity and communication stability . . . 34
5.3.3 Hardware and user interface refinement . . . 35

6 Conclusion **37**

Bibliography **39**

A Appendix 1 **I**

List of Figures

2.1 Examples of existing ROV platforms. 4
2.2 Examples of existing USV platforms. 5

3.1 Previous Seadragon prototype. 9
3.2 Sealed acrylic tube with internal electronics. 10
3.3 Custom machined aluminium end caps. 10
3.4 3D-printed electronic trays. 11
3.5 Structural aluminium frame, constructed from 3030 aluminium profiles. 11
3.6 Tube clamp and 3D-printed stabilizers. 12
3.7 Sonar mount. 12
3.8 3D-printed antenna mount. 13
3.9 BlueRobotics T200 thrusters. 13
3.10 Teledyne BlueView M900 Mk2 multibeam imaging sonar. 14
3.11 Simplified system architecture diagram. The layout illustrates the physical distribution of the respective electrical components in the port (command) and starboard (propulsion) hulls, as well as their connection through the shared umbilical and to the externally mounted components. No LTE module was included in the final prototype. . . 16
3.12 Overview of the employed microservice architecture. The IMU and GPS services were planned but never realized. 17
3.13 Control map of the PlayStation 5 DualSense bluetooth controller during remote control operation. 18
3.14 Experimental setup for sonar data collection. 19
3.15 Objects added to scan area to improve model specificity. 19
3.16 Collection of sample images. Contrast enhanced and colour adjusted for easier viewing. 20
3.17 Object distribution in the ADVD dataset. 21
3.18 Dual-stage fine-tuning pipeline for YOLO26n. In stage 1, the model is trained for 20 epochs on the UATD dataset to learn general sonar features. In stage 2, the neck and head are fine-tuned for 100 additional epochs on the custom ADVD dataset to specialize in target-specific textures and shadows. 21

4.1 Final assembled prototype. 23
4.2 Illustration of the Seadragon 2.0 search swath. 26
4.3 F_1-score vs confidence across all object classes. 28
4.4 Consfusion matrix. y-axis: Predicted objects, x-axis: Ground truth. . 29

4.5 Ground truth vs. prediction. 29

5.1 Example of labelling mistake (left) and correct prediction (right). . . 33

A.1 Complete schematic of the electrical system of the Seadragon 2.0 prototype. I

List of Tables

3.1 BlueRobotics T200 thruster performance specifications at 12 V. . . . 14
3.2 Technical specifications for Teledyne BlueView M900 Mk2. 15

4.1 Seadragon 2.0 mass analysis. 24
4.2 Comparison of search coverage and operational metrics. 27
4.3 Validation results per class. P = Precision, R = Recall. 28

1
Introduction

Globally, the World Health Organization (WHO) conservatively estimates that 236,000 people die by drowning annually [1], and attribute more than 2.5 million preventable deaths to drowning in the past decade alone [2]. Sweden and the Nordic region have one of the fewest reported rates of drowning in the world [3], and since the Swedish Life Saving Society (SLS) began monitoring fatal drowning incidents in 1898, the number of drowning victims in Sweden has declined from approximately 1,100 cases annually, to about 220 today [4]. Despite this encouraging trend, drowning remains a critical public safety concern, and over the past decade the downward trajectory has stagnated [5]. The SLS's goal of halving drowning incidents by 2030, coupled with its long-term vision of eliminating all fatalities, places increased pressure on existing strategies and highlights the need for new innovative approaches.

One promising avenue for improving drowning prevention and response is the application of advanced technological solutions in underwater *search and rescue* (SAR) operations [6]. While aerial tools such as thermal cameras and drones have become essential for rapid deployment and effective detection above the water surface, underwater search capabilities remain comparatively limited and often labour-intensive. Current methods, including side-scan sonar mounted on commercial vessels, traditional divers, and remotely operated vehicles (ROVs), typically require significant deployment time and may reduce overall operational efficiency. As survival rates are highly dependent on rapid response [7], addressing these inefficiencies is of great importance.

This thesis builds on an ongoing experimental initiative at Chalmers University of Technology known as the *Seadragon*, where a rapidly deployable sonar-equipped unmanned surface vehicle (USV) prototype is being developed. The previous prototypic design is too bulky for the stated mission requirement and is primarily made from 3D-printed components. This work aims to provide a redesign of the existing platform, hereafter referred to as *Seadragon 2.0*, and achieve a more durable construction, with enhanced software and machine learning (ML) capabilities.

1.1 Problem statement

Existing underwater search methods are too slow for life-critical windows. This work investigates whether a lightweight USV with ML-enhanced sonar can provide a faster, more accurate alternative for detecting and localizing submerged victims.

1.2 Research questions

Based on the stated goals of the project, the following research questions were formulated.

1. How effectively can a lightweight, rapidly deployable USV prototype support time-critical open-water search operations compared to conventional deployment methods?
2. To what extent can a USV-mounted sonar system, combined with machine-learning-based signal processing, detect and localize submerged human-sized targets in controlled aquatic environments?

1.3 Limitations

To narrow the scope of the project and adhere to national laws and regulations, the following limitations are set.

- Testing is restricted to controlled environments like pools and lakes due to national regulations on underwater surveying.
- The prototype is not designed for volatile sea states nor high-velocity currents.
- The project prioritizes rapid prototyping and technical proof-of-concept over a commercial-grade, fully durable product or a polished user interface.
- The USV is designed for remote-controlled operation rather than full autonomy, with software evaluation focusing specifically on detection accuracy of human-sized targets.

2
Background

In a 2014 case control study of drowning victims (n = 1377) by Quan *et al.* [7], they concluded that more than 90% of drowning victims suffer a bad outcome (death or severe disability) if submerged for 11 minutes or longer. A Swedish study (n = 6965) by Claesson *et al.* [8], further showed that the victims are never found in 18% of the reported incidents, highlighting the need for not only quicker response times, but also more advanced SAR target detection solutions.

2.1 Current state of maritime SAR

When a suspected drowning is reported, an initial fast first response team is typically dispatched to the area, and depending on the situation, followed by rescue divers or other more advanced SAR resources [9–11]. However, as time is of the essence, the amount of equipment the first responders can carry with them is limited, necessitating careful equipment and resource allocation.

Above the surface, tools such as night vision goggles, drones, and handheld thermal cameras are commonly utilized [6]. These are generally easy to carry and require little to no training. Conversely, sub-surface tools are generally more limited and often require extensive resources such as specialized sonar-equipped vessels and professional divers. While standard goggles and water binoculars are frequently used, they are severely limited by water visibility. Although handheld sonars have successfully been adopted by some SAR organizations in Sweden in recent years [12], they remain limited to small area scans and often require the operator to be in or on the water, exposing them to the same environmental hazards as the victim.

In deeper or more challenging water conditions, SAR operations currently rely heavily on experienced rescue divers [8]. Despite this, only 15 out of 290 municipalities in Sweden currently have SAR divers stationed, with the number of total divers reportedly declining [13]. Apart from making the rescue operations more difficult to perform in large parts of the country, it also risks increasing the time of rescue beyond the critical survival limit. However, increasing the number of divers alone is not likely to remedy the situation, as they too frequently face challenges. In an interview by Johansson (2021) [6], a Swedish rescue diver said: “Sometimes when we dive you only see 20 cm, sometimes nothing, sometimes 3 meters. If you are going to search an area big as a football field it takes a very long time.” Consequently, the study identifies a significant demand among rescue services for new and innovative cost-effective approaches that could enhance maritime SAR operations beyond their current capabilities.

2.2 Existing portable ROV and USV platforms

To bridge the gap between the availability of divers and the often vast and opaque search areas, autonomous and remotely operated technologies have emerged as a potential solution [14]. In particular, Johansson emphasizes that portable solutions hold the most promise for time-critical operations [6]. As a response to this, a wide range of commercial options have begun to emerge in this category [15–17].

To aid in a search and rescue mission in a large water reservoir, a Colorado based search and rescue team acquired a commercial VideoRay Pro 4 ROV (see Fig. 2.1a), leading to four successful drowning victim recoveries within the first three months of the purchase [18]. In a similar incident, the Houston County Rescue Unit successfully utilized a Deep Trekker DTG3 (see Fig. 2.1b) underwater ROV to locate two submerged bodies with the help of the onboard camera, after initially spending weeks searching using conventional divers [19].

(a) VideoRay Pro 4 [20]. **(b)** Deep Trekker DTG3 [16].

Figure 2.1: Examples of existing ROV platforms.

The amount of sensors and utility these vehicles provides varies greatly, and the cost can range from as low as few thousand dollars and go as high as $300.000 for state-of-the art sonar equipped platforms [21]. Commonly used systems today, such as the VideoRay PRO 4 and Deep Trekker DTG3 ROVs, typically have onboard optical cameras mounted as default, with more advanced features like multibeam imaging sonars being optional accessories that can be added based on the mission requirements [22].

In contrast to the more common tethered ROV platforms, there are also increasingly remote controlled and autonomous solutions coming to market [23–25]. Since 2010, Hydronalix, one of the leading suppliers of man-portable USVs, have sold their EMILY (emergency integrated lifesaving lanyard) platform (Fig. 2.2a), to both law enforcement and SAR organizations worldwide [23]. The vehicle can be carried by two persons and quickly be deployed to serve many time-critical SAR operations. In the case of an active drowning incident, the platform can be controlled remotely to the area of interest and immediately provide buoyancy assistance to up to four people in distress while awaiting further assistance. If bodies are suspected to be submerged in the water, the platform can also be fitted with a side-scan sonar (see

Section 2.3), providing a detailed overview of the area immediately below the vehicle.

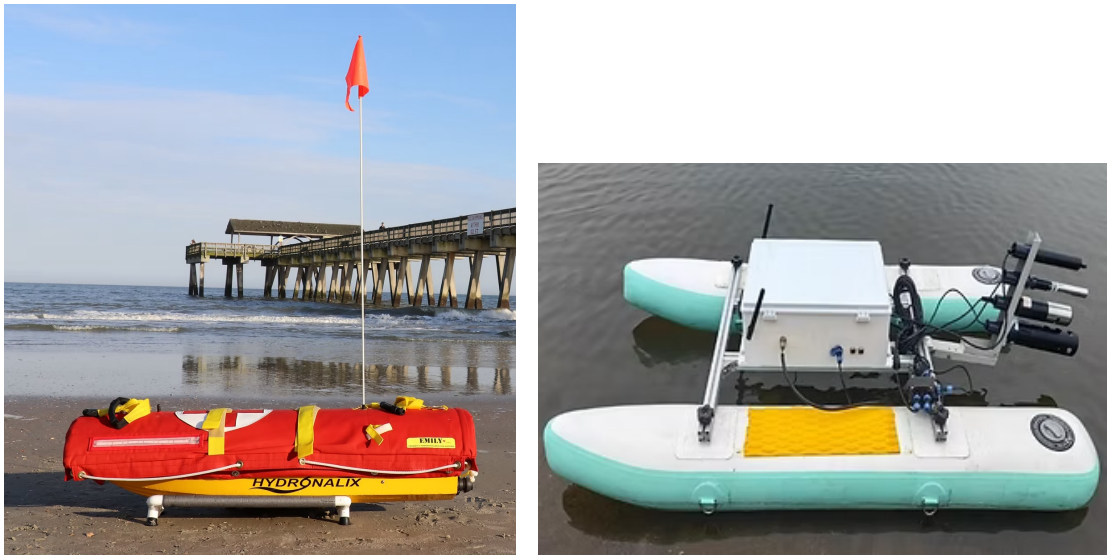

(a) EMILY [23]. **(b)** Agistar BX-USV II [24].

Figure 2.2: Examples of existing USV platforms.

One example of an autonomous platform (AUV), requiring little to no human intervention, is the Agistar BX-USV II vehicle [24], offering a highly customizable set of sensors and capabilities (Fig. 2.2b). Today, these platforms are however typically only used for general hydrographic surveying and water quality testing.

2.3 Imaging sonar technology

General sonar-based object detection has long been employed in both commercial and military maritime operations, with use ranging from object avoidance and fish detection, to seafloor mapping and submarine hunting [26]. Historically, sonar-based detection has been limited to simple pinging, with "blips" appearing at specific bearings and distances, or slightly more advanced side-scan sonars, which sends out a narrow fan of beams and relies on movement of the sonar to map the seafloor in thin, high-resolution slices.

With more recent advancements in high frequency multibeam imaging sonars and advanced phased arrays, so-called "acoustic cameras" which can provide live 2D images have been made possible, reaching frame rates of up to 40 frames per second [27]. These sonars offer a new way to image deep subsurface environments that previously have been difficult to capture. As the computational cost of *beamforming* (the mathematical processing of directional signal transmission) continues to drop, these high-resolution systems are now also beginning to be commercially viable and a practical choice for use in cost-effective SAR-platforms [28, 29].

2.3.1 Challenges in underwater acoustic imaging

Conventional optical underwater imaging commonly suffers from a combination of physical phenomena such as light absorption, refraction and particulate scattering caused by microscopic organisms and sediment [30, 31]. In contrast, as sound travels approximately 4.5 times faster in water than in air, acoustic imaging typically thrives in water. While sonars are not affected by the same light absorption and particulate scattering issues as optical cameras, they also face their own set of challenges. Sharp changes in water temperature can act like a mirror and reflect signals before they actually hit any target, and surface disturbances such as heavy rain or waves can create acoustic noise, cluttering the image [32, 33]. Additionally, although modern multi-beam sonars are getting close [34], they still generally lack the fine, granular detail of high-definition optical systems. For sonar ML-applications, the lack of diverse and high-quality training data also remains a significant bottleneck [35]. Efficient and innovative ML-architectures and training pipelines are therefore critical to achieve reliable and high-accuracy results [36].

2.4 Machine learning for target detection

Since the inception of AlexNet in 2012 [37], convolutional neural networks (CNNs) have been the dominant paradigm in modern computer vision and image processing tasks, being used as the backbone in a wide range of applications, from state of the art medical diagnostics, to autonomous vehicles [38, 39]. More recent advancements include the introduction of residual network (ResNet) connections [40], which in part solved the vanishing gradient issue commonly plaguing regular CNNs, and the transformer architecture with its now ubiquitous self-attention mechanism [41].

While the transformer architecture has been wildly successful in a wide range of artificial intelligence (AI) applications and the main driver behind the recent developments of large language models (LLMs), it requires a significant amount of data to overcome its weak *inductive bias*, meaning it has no or very little prior knowledge of the problem at hand. In contrast, CNNs have built-in knowledge about how images work (e.g., locality and translation invariance), which means they can perform exceptionally well even with relatively small datasets. For object detection specifically, region-based convolutional networks (R-CNNs) have proven to be very accurate and offer a good trade-off between performance and the amount of training data required [42]. The multiple stage architecture of the common R-CNN networks however also inherently comes with a delay in the processing time required to parse the images, making real-time object detection challenging.

As a response to the albeit powerful, but slow, multi-stage nature of the R-CNN architecture, the first you only look once (YOLO) architecture was developed in 2015 [43]. As the name implies, it replaces the multi-stage pipelines (as found in R-CNN networks), with a single neural network that predicts bounding boxes and class probabilities directly from full images in one evaluation. This drastically improves the speed at which images can be processed. The original model proposed in the 2015 paper processed images at 45 frames per second, enabling real-time detection, with limited effects on the overall performance.

2.5 Emerging sonar SAR capabilities

In a 2019 study by Nguyen *et al.* [29], a Teledyne BlueView M900-90 multibeam imaging sonar was used to train two different CNN based neural networks (AlexNet/-GooLeNet), on small dataset of submerged body shaped objects. To reduce the effects of the inherent noise in the real-world sonar images, they employed a number of data augmenting techniques to their clean pool dataset. The authors added salt and pepper noise to better represent the ubiquitous background noise from the surrounding environment, and a varying polarization filter to better capture the effects of sharp and directionally defined reflections sporadically caused by the sea floor. To account for varying levels of wave absorption at different target distances, they also included inverted images and artificially altered pixel intensities. Finally, they performed various image translation and rotation transformations to better capture different angles of attack not fully represented in the recorded dataset. With these, the best performing GoogLeNet model achieved a final accuracy of 91.6%, up from only 31.6% with no additional augmentations.

In a more recent (2024) study by Hu and Liu [36], the authors successfully increased the accuracy of the their models by utilizing novel transfer learning techniques, incorporating both synthetic aperture radar- and optical images into the training pipeline, effectively increasing the size of the dataset. In addition, they modified the YOLOv5s architecture, which served as the backbone in their model, to incorporate a ShuffleNetv2 channel splitting technique. This reduced the computational complexity (GFLOPs) from 60.4, down to just 5.9, with no discernible impact on the mean average precision (0.812) of the model.

3
Methods

This chapter presents the design improvements implemented in Seadragon 2.0, the revised version of the original Seadragon USV platform. It also covers the machine learning pipeline, from sonar data acquisition to target detection, as well as the software and system architecture employed.

3.1 System architecture and design

The previous prototypic Seadragon design consisted of two large 3D-printed pontoons for buoyancy, a metal sheet frame connecting the bodies, and a central waterproof box to house the electronics (Fig. 3.1).

Figure 3.1: Previous Seadragon prototype.

This design provided considerably more buoyancy than needed for normal operation, had water ingress problems at the joints of the 3D-printed pontoon sections, and was not equipped with the sensors needed for accurate positioning and efficient wireless communication. To address these issues, several changes to the design were made.

3.1.1 Hardware redesign and prototyping

To improve the structural integrity and simplify the assembly process of the prototype, the original 3D-printed pontoons were replaced with two acrylic tubes, each with an outer diameter of 150 mm and length of 480 mm. To prevent water from entering, two sets of custom aluminium (AW-5083) end caps were machined to seal the tubes. The two resulting watertight volumes were then used to house the internal electronics (Fig. 3.2), removing the need for the centrally placed watertight box.

Figure 3.2: Sealed acrylic tube with internal electronics.

The aluminium end caps were designed with two radial o-rings, providing a robust double-layered seal against the inner walls of the acrylic tubes (Fig. 3.3). On the face of the rear mounted end caps, holes were drilled to allow for cable penetration using BlueRobotics M10 WetLink penetrators. A lip extruding from the outer face of the end caps were also added for easier removal and as a potential attachment point for future hydrodynamically improved nose cones.

(a) Front face of rear end cap.

(b) Front view of bow end cap.

(c) Rear view of bow end cap.

Figure 3.3: Custom machined aluminium end caps.

Inside the tubes, 3D-printed trays were designed to hold all of the electronics securely in place (Fig. 3.4). They were printed in two sections and held together by pressure-fitted dowel joints.

(a) Logic tray (port/logic side).

(b) Battery tray (identical on both sides).

(c) Combined tray (starboard/propulsion side).

Figure 3.4: 3D-printed electronic trays.

To connect two tubes together and provide a structural frame for the thrusters and sonar to mount to, 3030 aluminium profiles were used (Fig. 3.5). This simplified the construction of the prototype and provided an easy way to shift the balance of the entire craft. By simply translating the mounts of the thrusters, antenna and sonar along the length of the profiles, the centre of gravity of the platform could be adjusted as needed.

Figure 3.5: Structural aluminium frame, constructed from 3030 aluminium profiles.

The side-mounted tubes were secured to the frame by two custom made aluminium clamps (Fig. 3.6a). They were made from 3 mm aluminium sheet and split into two sections, one secured to the frame from above, and one from below, clamping the tube tight (Fig. 3.6b). To stabilize the tubes and prevent side-to-side wiggle at the ends, 3D-printed stabilizers were inserted into the T-slots of the profiles at the bow and stern of the frame (Fig. 3.6c).

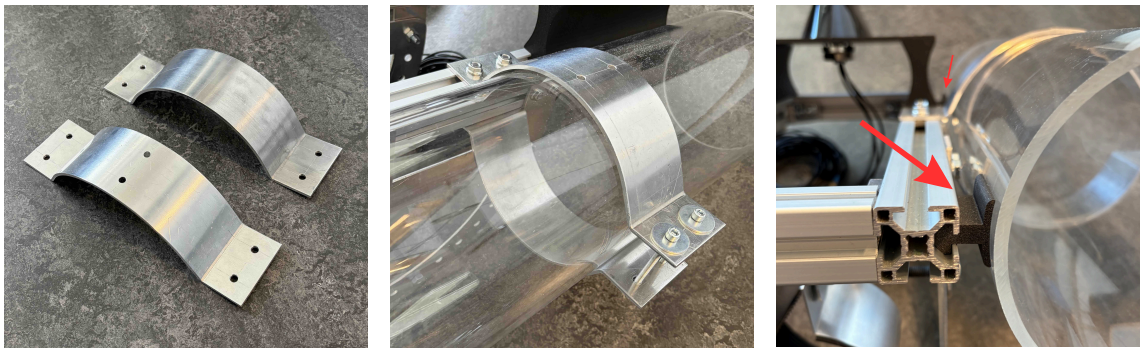

(a) Tube clamp. **(b)** Mounted clamp. **(c)** Tube stabilizers.

Figure 3.6: Tube clamp and 3D-printed stabilizers.

Similarly, the mount for the sonar clamp was made from 3 mm aluminium sheet, consisting of two distinct parts. The vertical mounts (Fig. 3.7a), were designed to lower the placement of the sonar such that its vertical field of view was neither obstructed by the body of the craft, nor interfering with surface of the water. The lower adjustable mounts were bolted to both the vertical mount connected to the frame (Fig. 3.7b), and with the clamp holding the sonar (Fig. 3.7c). To allow for a wide range of search configurations, spanning from looking directly down at the seafloor, to the body of water immediately in front of the craft, mounting holes were drilled 10° apart radially along the mount.

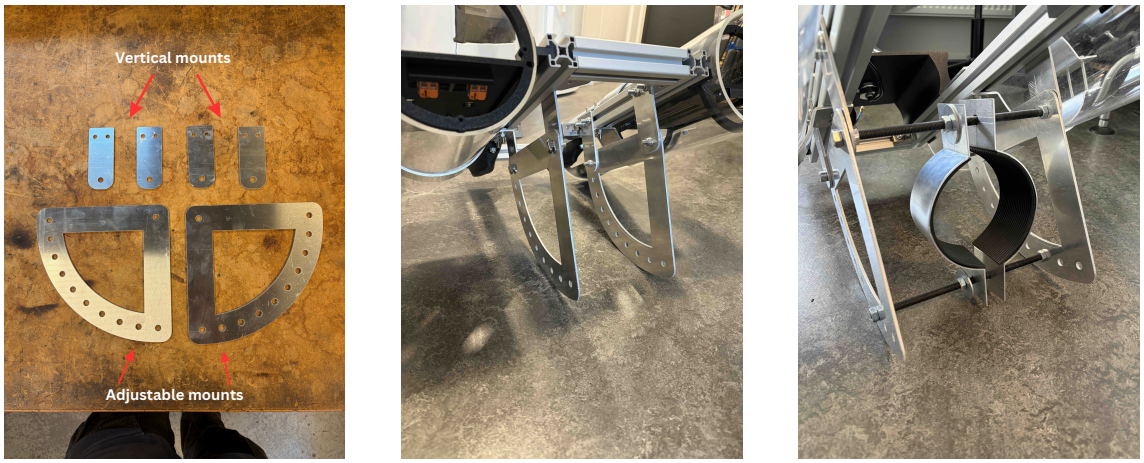

(a) Vertical and adjustable mounts. **(b)** Assembled mount secured to the frame. **(c)** Sonar clamp secured to adjustable mount.

Figure 3.7: Sonar mount.

The clamp holding the sonar was made from a thinner 1.5 mm aluminium sheet and dressed with protective rubber sheet. Threaded rods spanning horizontally across of the adjustable mount were used to position and hold the clamp secure in the desired angular position.

On the rear of the craft, a raised 3D-printed antenna mount was inserted and pressure fitted into the top-facing t-slots of the frame (Fig. 3.8), giving the external antenna a clear and unobstructed view from all sides.

Figure 3.8: 3D-printed antenna mount.

3.1.2 Propulsion and power system

For forward propulsion, two BlueRobotics T200 thrusters were used, along with two electronic speed controllers (ESCs) and a 12V lead-acid battery powering the system. Using the two thrusters, steering was achieved through differential thrust vectoring, enabling the USV to make both precise course corrections as well as stationary manoeuvrers. The complete electrical system was housed entirely inside the starboard tube of the USV, with the two T200 thrusters being mounted in the rear and attached under the longitudinal running profiles of the frame (Fig. 3.9).

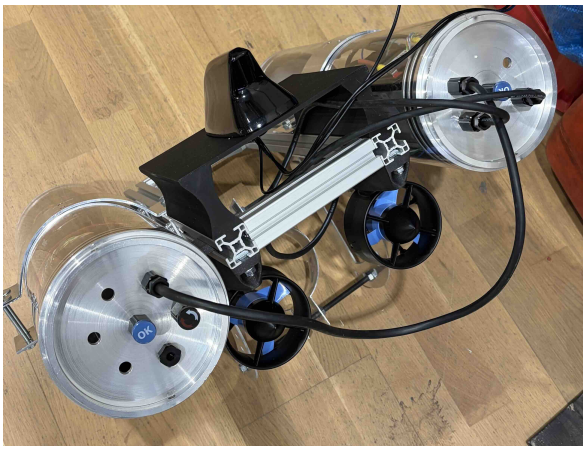

Figure 3.9: BlueRobotics T200 thrusters.

The T200 thrusters were chosen due to their reasonable cost to performance ratio and their extensively proven track record in other ROV and AUV platforms [44]. At 12V, they provide 3.71 kgf each of forward thrust at full throttle, giving a combined maximum thrust of 7.42 kgf. The ESCs which control the 3-phase brushless motors of the T200s, are themselves controlled by a pulse width modulation (PWM) signal. A 1100 µs long pulse tells the controllers to set the motors to full speed reverse, a 1900 µs pulse indicate full speed forward, and a neutral state is defined as a signal length of 1500 µs, with a small deadband (±25 µs) around the centre to prevent the thrusters from jittering. The complete performance specifications of the T200 thrusters are listed in Table 3.1 below.

Table 3.1: BlueRobotics T200 thruster performance specifications at 12 V.

Parameter	Value	Unit
Full throttle forward thrust	3.71	kgf
Full throttle reverse thrust	2.92	kgf
Minimum throttle	0.02	kgf
Full throttle current	17.0	A
Full throttle power	205.0	W

3.1.3 Teledyne multibeam imaging sonar

To capture the required high-resolution 2D sonar images required for the Seadragon project, a Teledyne BlueView M900 Mk2 (Fig. 3.10) multibeam imaging sonar was used. It's a widely employed sonar in ROV and AUV applications [45], operating at 900 kHz, striking a good balance between high-resolution detail and a suitable range for SAR operations.

Figure 3.10: Teledyne BlueView M900 Mk2 multibeam imaging sonar.

With a 130° horizontal and 12° vertical field of view, it can efficiently perform wide-area searches crucial for SAR operations. Additionally, unlike traditional side-scan sonars that require constant forward motion to reconstruct an image, the M900 Mk2's update rate of up to 25 Hz allows for continuous streaming of 2D acoustic imagery, even when completely stationary. The complete specifications of the sonar are listed in Table 3.2 below.

Parameter	**Specification**
Operating frequency	900 kHz
Field of view	130° Horizontal × 12° Vertical
Max range	100 m
Optimum range	2–60 m
Number of beams	768
Beam spacing	0.18°
Range resolution	1.3 cm
Update rate	Up to 25 Hz
Power consumption	20 Watts (typical)
Depth rating	1,000 m

Table 3.2: Technical specifications for Teledyne BlueView M900 Mk2.

On Seadragon 2.0, the BlueView M900 Mk2 sonar is powered by the same 12V lead-acid battery that powers the onboard Jetson Orin Nano (Fig. 3.11), with an ethernet connection directly passed to the computer for live data transfer.

3.1.4 Computer and sensor integration

An NVIDIA Jetson Orin Nano was chosen to be the main computing device on the Seadragon 2.0. Due to its compact form factor and high computational throughput, it is particularly well suited for edge computing applications with significant spatial constraints, such as the internal volume of the Seadragon 2.0. It is built around the NVIDIA Ampere architecture with 1,024 CUDA cores, 32 Tensor cores and a 6-core Arm Cortex-A78AE CPU, delivering up to 67 TOPS (trillion operations per second). This configuration provides both the processing power needed to facilitate the real-time target detection algorithm and to timely transmit the necessary motor controls to the ESCs.

Although never utilized, two Adafruit breakout modules were connected to the Jetson Orin Nano: one Ultimate GPS module (PA1616D), and one 9 DoF (degree of freedom) orientation IMU (inertial measurement unit) module (BNO085). This to allow for more precise and reliable positioning in future iterations of the project. Moreover, to allow for a reliable signal reception within the watertight enclosure at large distances, an external multi-band Abracon shark fin antenna (AECS1806C03ZS-3000S) was mounted to the exterior hull. The antenna features three distinct antenna elements, one tuned for cellular (LTE), one tuned for satellite positioning (GNSS/GPS), and one for Wi-Fi.

To safely send the control signals from the *command* tube, over the umbilical to the ESCs in the *propulsion* tube, an opto-isolator circuit was added between the two. This ensured that no stray current or electrical noise could travel from the high-current circuit required to power the thrusters, into the more sensitive and lower-current circuit of the command tube. In Fig. 3.11, a simplified system architecture and wiring block diagram of the setup is presented. A complete wiring diagram can be found in Appendix A.1.

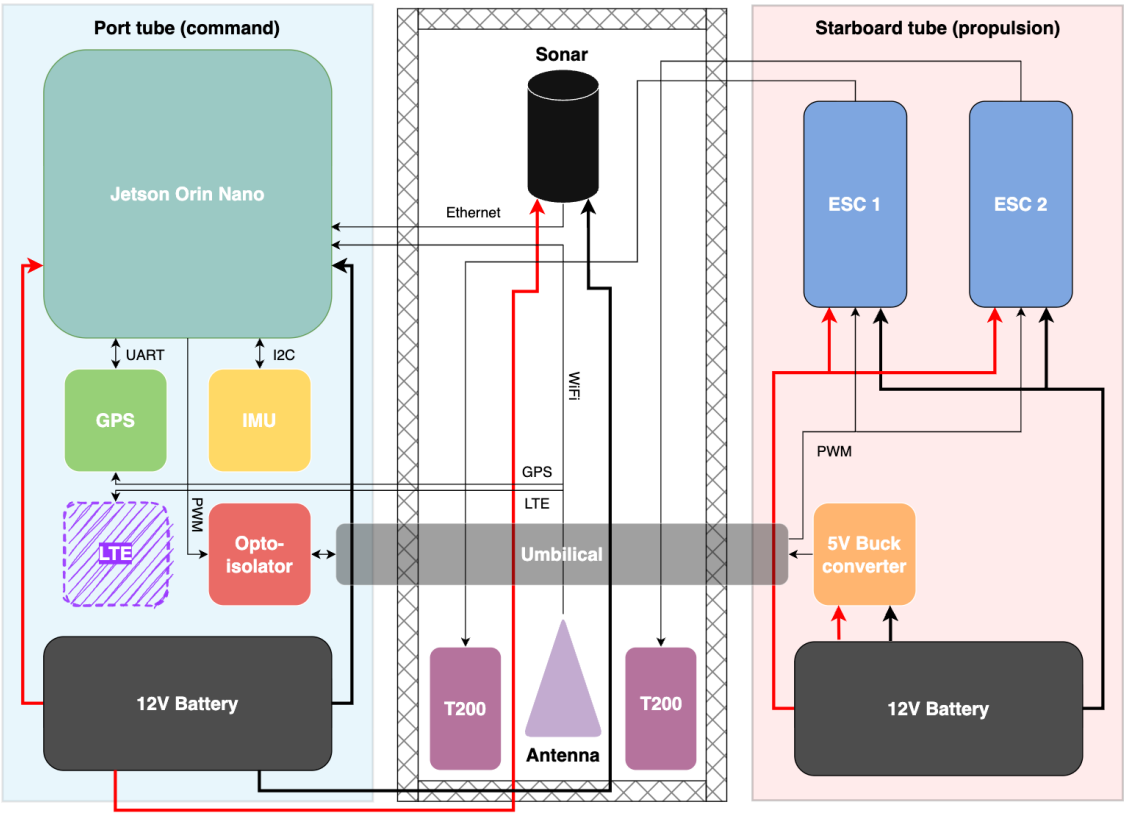

Figure 3.11: Simplified system architecture diagram. The layout illustrates the physical distribution of the respective electrical components in the port (command) and starboard (propulsion) hulls, as well as their connection through the shared umbilical and to the externally mounted components. No LTE module was included in the final prototype.

3.1.5 Software architecture and development environment

The software environment deployed on the Seadragon 2.0 was built upon NDVIDIA JetPack 6.2, an NVIDIA specialized distribution of Ubuntu 22.04 LTS, optimized for the Jetson Orin Nano's hardware acceleration capabilities. Development was primarily carried out in Python for rapid prototyping purposes, with some performance-critical processes being written in c++ to ensure real-time execution.

To achieve efficient and robust communications between the different components and services of the system, a microservice architecture was employed. Each core component or service (e.g., sonar driver, ML-processing pipeline, operator GUI), had

its own dedicated microservice running independently from the others, improving both the modularity and stability of the system. See Fig. 3.12 for an overview of the system architecture.

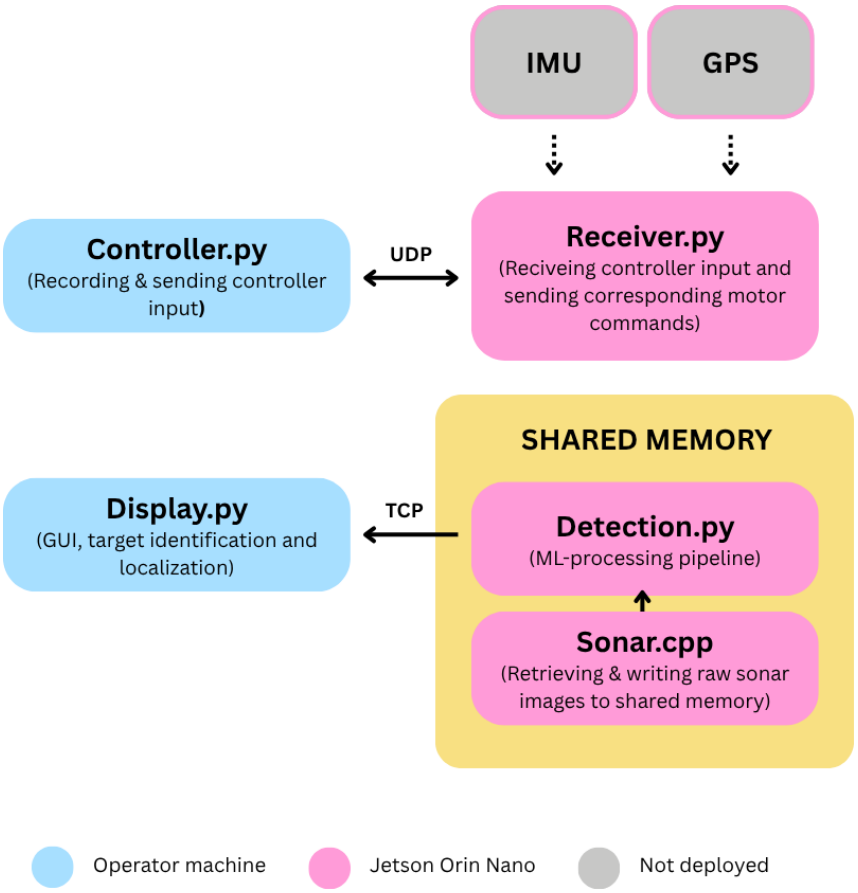

Figure 3.12: Overview of the employed microservice architecture. The IMU and GPS services were planned but never realized.

Inter-service communication (ISC) between the controller input microservice and the receiving motor command microservice was handled via the user datagram protocol (UDP), providing an asynchronous and low-latency messaging framework suitable for real-time commands. To minimize latency and CPU overhead during target detection, inter-process communication (IPC) via fixed-size shared memory was implemented between the sonar microservice and the machine learning processing pipeline. The ML-processed images were later sent over the transmission control protocol (TCP) to ensure reliable and accurate image reconstruction on the operator display.

3.1.6 Remote control

Although the revised Seadragon prototype was ultimately designed to be autonomous, remote control capability was added for testing and fine-control purposes. The triggers and bumpers of a PlayStation 5 DualSense Bluetooth controller was used to control the individual thrusters, while the joysticks of the controller were mapped to uniform thrust output for easier handling. Additional buttons on the controller were used for

thrust safety toggles. See Fig. 3.13 for an illustration of the complete input mapping.

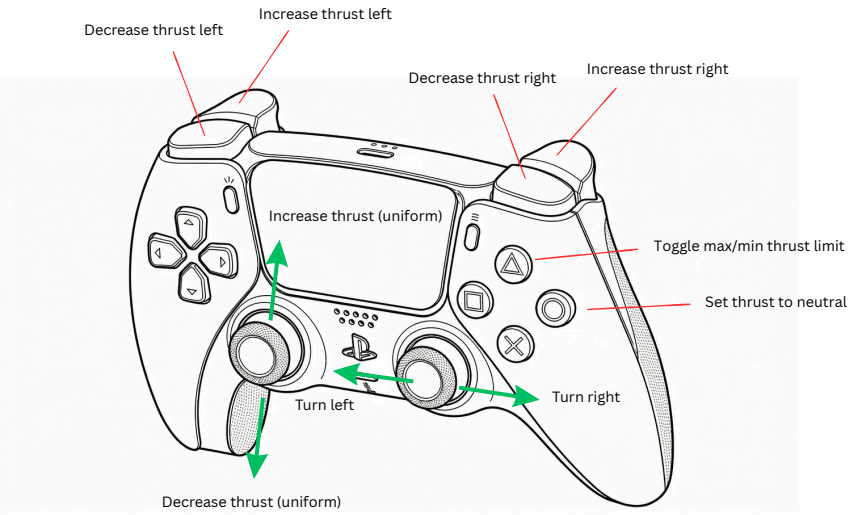

Figure 3.13: Control map of the PlayStation 5 DualSense bluetooth controller during remote control operation.

3.2 Machine learning pipeline

The development of a robust machine learning pipeline for underwater target detection and localization, requires addressing the fundamental differences between acoustic and optical data modalities. While most modern computer vision algorithms are trained on and built for 3-channel coloured images, sonar images are limited to only one intensity channel. This typically requires either modifications to the early layers of to the computer vision architecture itself, or the use of *channel replication*, where duplicates of the same image are stacked in all three channels. In addition, carefully selected transfer learning techniques often have to be employed to ensure that the difference in feature space (e.g., greyscale, speckle noise) is handled properly.

Alternatively, a dedicated purpose built pipeline could be trained from the ground up, designed with the differences in modality (light vs. sound) in mind. However, this generally requires large datasets with high-quality annotations (of which there are few available in the public domain), or extensive use of style transfer techniques like CycleGAN, where simulated or synthetic sonar images are generated to expand the dataset.

For this project, a fine-tuned version of the YOLO26n model from Ultralytics was implemented, incorporating channel replication and a combination of different transfer learning and data augmentation techniques. The model was chosen due to its excellent size-to-performance ratio and native end-to-end predictions, enabling real-time inference and wide hardware support [46].

3.2.1 Dataset acquisition and augmentation

To overcome the data scarcity issue and ensure reliable and accurate inference, a custom acoustic drowning victim dataset (ADVD) was recorded in a controlled 3.5×6.5 m clean pool environment (Fig. 3.14).

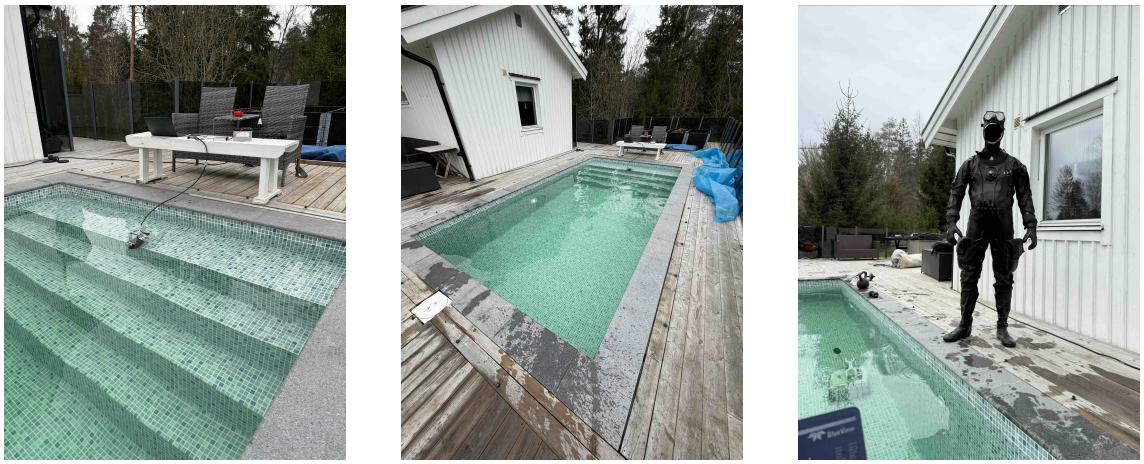

(a) Sonar setup at the edge of the pool.

(b) Clean pool environment.

(c) Human in drysuit acting as drowning victim.

Figure 3.14: Experimental setup for sonar data collection.

To simulate a realistic SAR scenario, a human subject equipped with a drysuit acted as the drowning victim. Additionally, to enhance model specificity and ensure that the system could distinguish human-sized targets from common underwater debris, a variety of additional objects (pool net, dumbbell, kettlebell, lawn chair, pool vacuum, plastic bag) were randomly introduced to the scan area (Fig. 3.15).

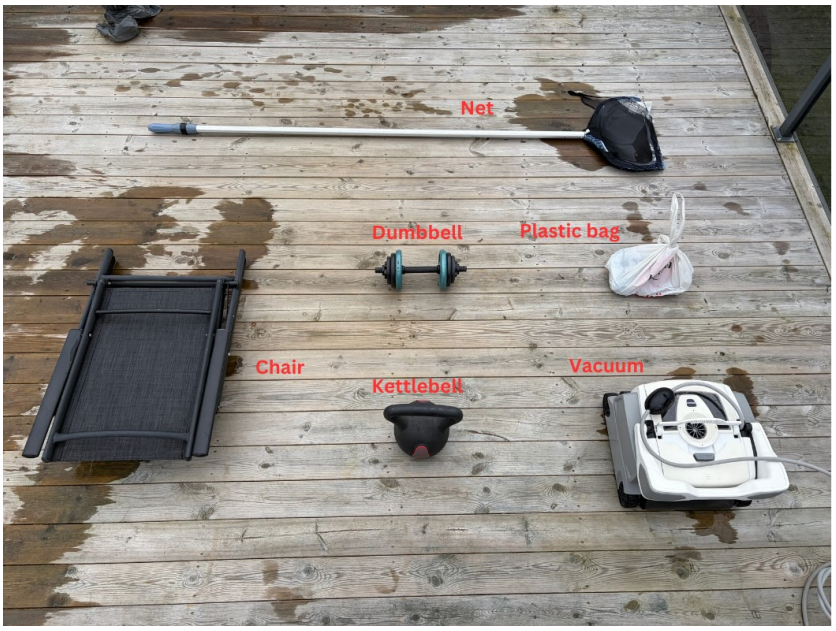

Figure 3.15: Objects added to scan area to improve model specificity.

This variety was intended to minimize false positives by forcing the network to recognize the unique acoustic signatures of human anatomy versus rigid or metallic geometries, such as metal pipes or sunken debris commonly found on the ocean floor.

Moreover, as objects in close proximity often smear and mix together in the reflected sonar pings, visually distinct objects can appear as one in the resulting image. To overcome this and to aid in the subsequent annotation process, regular optical video of the scene was recorded simultaneously. This significantly improved the labelling accuracy and quality as individual frames of the two modalities could be synced and compared to confirm placement and object boundaries.

Furthermore, to bridge the gap between the clean pool environment and the typical noisy conditions experienced in real-world open-water SAR operations, salt and pepper noise of varying intensities was applied to 50% of the collected images to simulate particulate backscatter. A collection of sample images from the resulting augmented dataset, with enhanced contrast for visualisation purposes, is shown in Fig. 3.16 below.

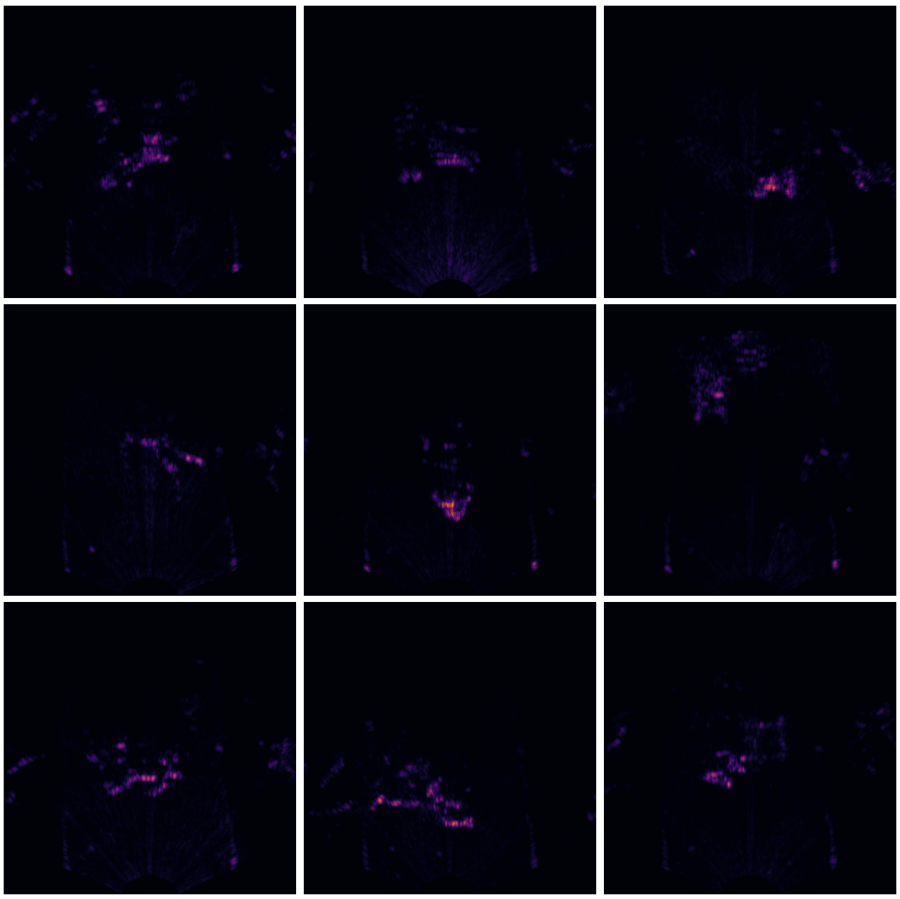

Figure 3.16: Collection of sample images. Contrast enhanced and colour adjusted for easier viewing.

In total, 256 unique sonar pings were extracted and annotated, amounting to 599 labelled object instances, with a class distribution heavily favouring the *body* category to prioritize search accuracy. The final distribution and individual frequency of the objects present in the dataset is shown in Fig. 3.17.

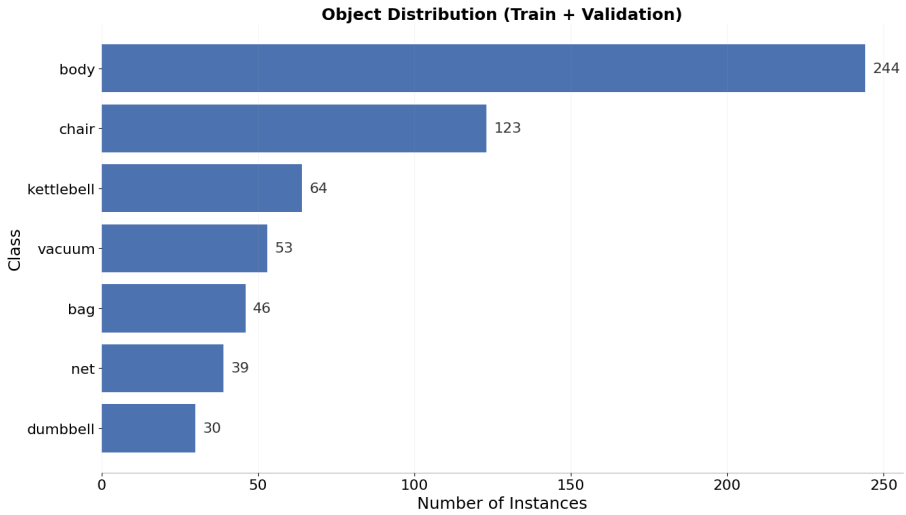

Figure 3.17: Object distribution in the ADVD dataset.

3.2.2 Model architecture and training

Since the YOLO family of models are trained on the COCO (Common Objects in Context) dataset with over 300,000 images spanning 80 categories in their natural environment, they offer excellent general object detection out of the box. However, as the base dataset does not contain any sonar imagery, the models typically perform poorly on acoustic images without further adaptation. To remedy this, a two-stage training pipeline was constructed (Fig. 3.18).

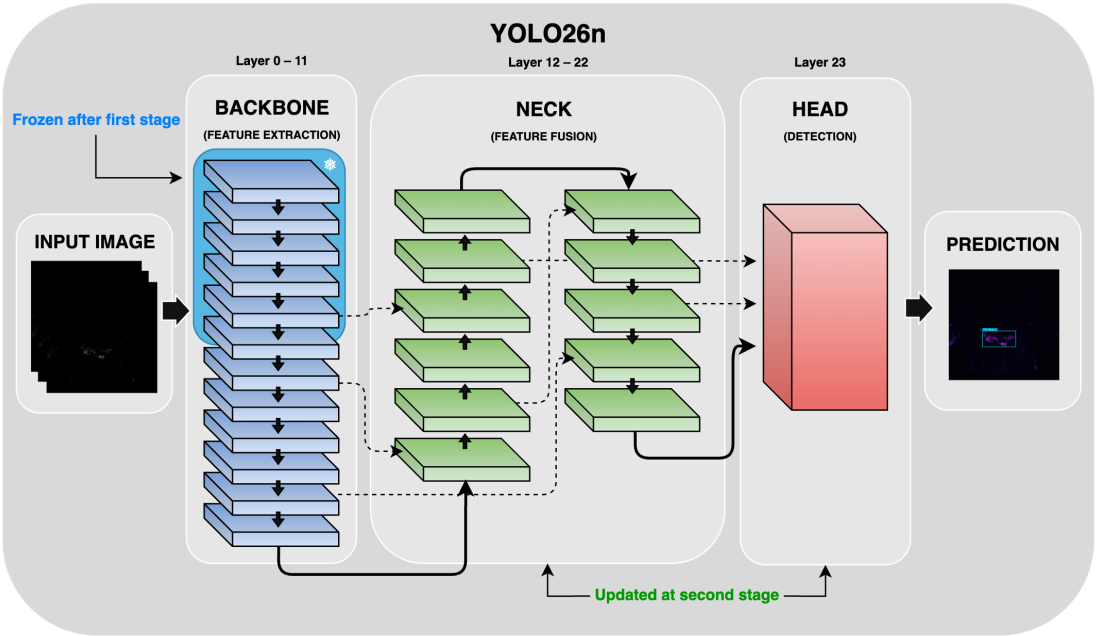

Figure 3.18: Dual-stage fine-tuning pipeline for YOLO26n. In stage 1, the model is trained for 20 epochs on the UATD dataset to learn general sonar features. In stage 2, the neck and head are fine-tuned for 100 additional epochs on the custom ADVD dataset to specialize in target-specific textures and shadows.

To overcome the lack of generality across the modalities, the first stage was designed to teach the model the fundamental characteristics of sonar imagery. This was achieved by fine-tuning the model on the underwater acoustic target detection (UATD) dataset by Xie et al. [47]. The dataset consists of a diverse collection of over 9,000 raw multibeam forward-looking sonar images across 10 different classes of objects, collected in a similar fashion to the ADVD dataset. The initial training was performed for 20 epochs with a learning rate of 0.01.

To subsequently prepare the model for the specific sonar images expected during inference, a final fine-tuning stage was implemented on the ADVD dataset for an additional 100 epochs. During this stage, the early layers of the model, which typically are responsible for simple geometry and basic feature detection, were frozen. This preserved the general sonar features learned in the first stage while allowing the head of the model to specialize in the exact textures and shadows produced by the BlueView M900 Mk2 sonar. Additionally, the learning rate was decreased to 0.001 to ensure stable convergence and prevent overwriting the knowledge acquired during the first stage.

4

Results

The development and testing of the revised Seadragon prototype resulted in a partially functional USV, offering a substantial upgrade in capability and robustness over the previous design. Although time constraints prevented experimental validation of the fully assembled prototype before the conclusion of this report, all major subsystems were successfully validated individually. These included the propulsion system, sonar sensing, remote comunication and inter-process communication.

The following sections detail the specific metrics of the prototype's physical performance, the detection accuracy of the machine learning model, and its overall computational and operational efficiency in simulated SAR scenarios.

4.1 Final assembled prototype

While most of the major subsystems of the Seadragon 2.0 were independently validated, full integration of all planned features was deferred to focus on core prototype testing. Although the GPS and IMU modules were connected to the Jetson Orin Nano, they were never functionally tested to evaluate integration performance and accuracy. Furthermore, while waterproofing tests were successfully performed using simulated weights, no marine tests were conducted with the electronics onboard. The final assembled prototype is shown in Fig. 4.1 below.

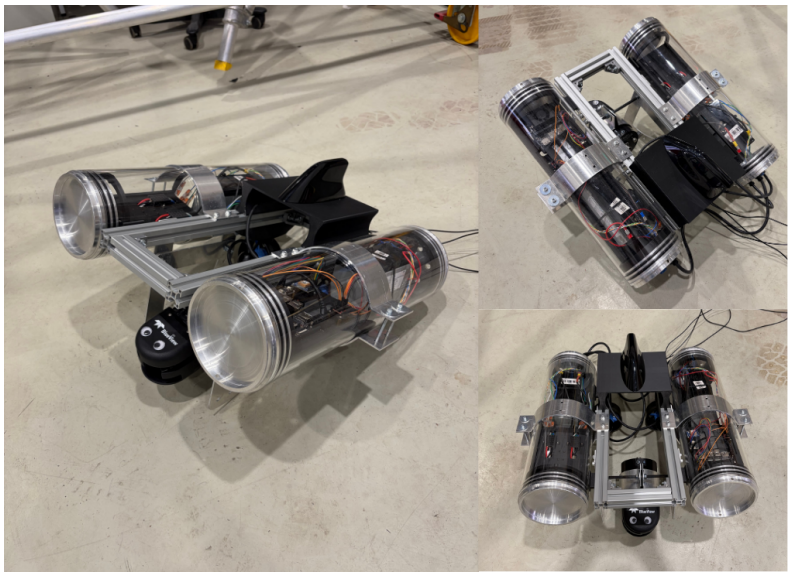

Figure 4.1: Final assembled prototype.

4.1.1 Mass and buoyancy analysis

In Table 4.1, an overview of the masses involved in the prototype is presented. Smaller components such as the custom machined clamps, mounts, electronics and 3D-printed parts are grouped together into a single miscellaneous category.

Table 4.1: Seadragon 2.0 mass analysis.

Component	**Qty**	**Unit mass (air)** [kg]	**Total mass (air)** [kg]	**Mass in water** [kg]
3030 Al frame	1	1.054	1.054	0.664[a]
Sonar	1	2.500	2.500	1.200
Thruster	2	0.427	0.854	0.478
Antenna	1	0.500	0.500	N/A
Acrylic tube[b]	2	1.301	2.602	0.415[c]
Al end cap	4	0.650	2.600	1.637
12V Battery	2	2.600	5.200	N/A
Other misc.	1	2.500	2.500	N/A
Totals			**17.81 kg**	—

[a] Calculated as $m_{air} - m_{wdisp}$, where $m_{wdisp} = \frac{m_{air}}{\rho_{Al}} \times \rho_{water}$.
Densities used: $\rho_{Al} = 2700\,\mathrm{kg/m^3}$, $\rho_{water} = 1000\,\mathrm{kg/m^3}$.
[b] Material volume calculated using $OD = 150\,\mathrm{mm}$, $ID = 140\,\mathrm{mm}$, $L = 480\,\mathrm{mm}$.
[c] Submerged mass of the plastic wall only ($m_{air} - V_{wall} \times \rho_{water}$).

To evaluate the expected buoyancy of the prototype, we use Archimedes principle, which states that the upward buoyant force is equal to the weight of the fluids displaced by the object.

The volume (V) of one tube can be calculated as:

$$V_{one} = \pi \times r^2 \times L = 0.00848\,m^3,$$

where r and L are the radius and length of the tube, respectively.

The resulting buoyant force (B_{one}), if sealed and 100% submerged, can then be calculated as:

$$B_{one} = V_{one} \times \rho_{water} \times g = 83.211\,N,$$

where the density of fresh water $\rho_{water} \approx 1000\,\mathrm{kg/m^3}$ and $g = 9.81\,\mathrm{m/s^2}$. Combined, two sealed tubes thus provide a total of 166.42 N buoyant force, or 16.96 kg of equivalent lift.

However, simulated buoyancy tests performed on the fully assembled prototype, show a net positive payload capacity in excess of 2.0 kg. This indicates that the volume displaced by the remaining structural components, such as the aluminium frame and end caps, successfully compensates for the prototype's dry mass, ensuring it remains sufficiently buoyant for normal operation in calm waters.

4.2 Search coverage and operational time

As the Seadragon 2.0 was not ready for real-world trials at completion of this thesis, the following performance metrics are based on theoretical calculations of the expected search coverage and operational endurance. These metrics were derived from experimental tests, estimated hydrodynamic performance, and the power consumption profiles of the onboard systems.

4.2.1 Operational speed

To estimate the prototype operational speed v in the water, we use the drag equation, defined as:

$$F_d = \frac{1}{2}\rho v^2 C_d A, \tag{4.1}$$

where F_d is the hydrodynamic drag, ρ the density of water and A the projected frontal area of the craft.

Rearranging and solving for the drag coefficient C_d, we get:

$$C_d = \frac{2F_d}{\rho v^2 A}. \tag{4.2}$$

Using this, the drag coefficient of the Seadragon 2.0 was estimated experimentally by performing a tow test and measuring the speed v and the force F_d used to pull the prototype. The force was measured using a dynamometer attached to the tow line, and the speed was determined by optical recording analysis over a fixed distance.

With a pulling force of 2.5 kgf and tow velocity of 1.1 m/s, the resulting drag coefficient C_d was estimated to be 0.62.

By applying the maximum thrust of the combined T200 thrusters (7.42 kgf) and solving for the speed v_{max}, we can then estimate the expected maximum velocity as:

$$v_{max} = \sqrt{\frac{2F_d}{\rho C_d A}} = 1.90 \ m/s, \tag{4.3}$$

where $F_d = 7.42 \times 9.81 = 72.79 \ N$, $A = 0.065 \ m^2$, and $\rho = 1000 \ kg/m^3$.

4.2.2 Operational endurance

The operational endurance of the USV is dictated by the capacity of the two 12V lead-acid batteries and the combined power draw of the internal and external components. At full throttle, each T200 thruster consumes around 205 W (17.0 A), resulting in a maximum propulsion-side power draw of 410 W/34 A. Meanwhile, on the sensing and computation side, the Teledyne BlueView sonar and the Jetson Orin Nano draw a combined 35 W, totalling less than 5 A. Consequently, the prototype consumes an estimated 445 W of continuous power when operating at its maximum capability.

However, as the two subsystems are powered independently, the difference in their power consumption leads to uneven battery discharge. With each battery providing a capacity of 8.5 Ah, and the *Peukert effect* causing a decreased capacity at high

discharge rates, the battery on the propulsion-side will be fully depleted after less than 15 minutes of operation. In contrast, the lower draw of the computational and sensing side allows it to remain operational for an additional 80 minutes, eventually resulting in an asymmetrical power loss.

4.2.3 Theoretical search coverage

The search effectiveness of the rapid deployable USV concept is primarily governed by the platform's velocity and effective range of the onboard sonar. Utilizing the estimated top speed of $v = 1.90$ (as derived in section 4.2.1), and an assumed effective range of 30 m for accurate and reliable target identification, a theoretical search coverage rate (C_{rate}) for the Seadragon 2.0 can be calculated as:

$$C_{rate} = v \cdot W_{swath}, \tag{4.4}$$

where W_{swath} is the swath width.

Given that the horizontal field of view (FOV) of the sonar is 130° (see Fig. 4.2), the effective swath width at 30 m can be calculated as:

$$W_{swath} = 2 \cdot 30 \cdot tan(65°) = 128.6\ m, \tag{4.5}$$

The search coverage rate (C_{rate}) can then be estimated to be:

$$C_{rate} = 1.90 \cdot 128.6 = 244.3\ m^2/s. \tag{4.6}$$

While this is significantly faster than traditional rescue divers (see section 4.2.4), the coverage rate is limited by the relative high drag coefficient (0.62) associated with the prototype's un-streamlined hull and flat end caps. Moreover, as the battery on the propulsion side drains in under 15 min under full throttle, the effective utilization time of the prototype is heavily limited.

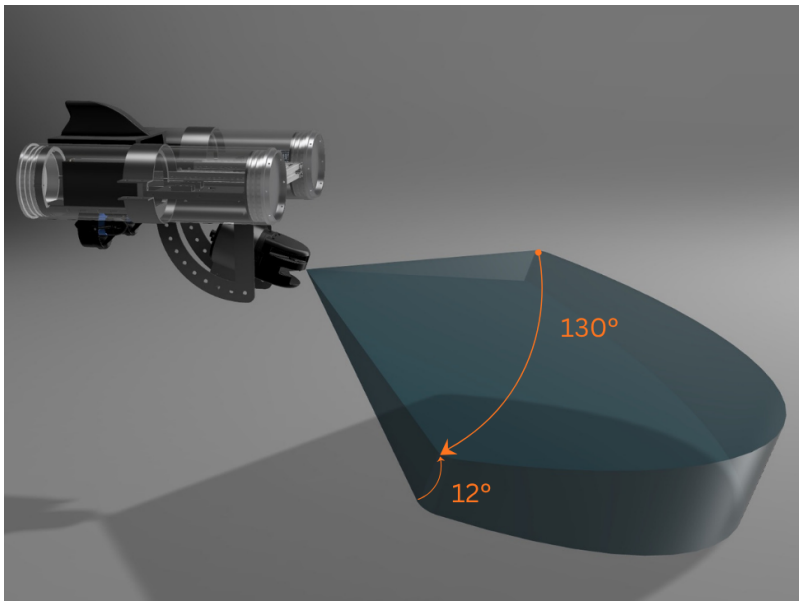

Figure 4.2: Illustration of the Seadragon 2.0 search swath.

4.2.4 Comparison to conventional methods

Assuming a standard swim pace of 0.5 m/s and moderate line of sight of 2 meters, a comparable search coverage rate C_{rate_d} for a rescue diver can be calculated as:

$$C_{rate_d} = 0.5 \cdot 4.0 = 2.0 \ m^2/s, \tag{4.7}$$

given an approximate effective sweep width twice that of the visible range.

For a standard commercial or SAR side-scan operation, vessels typically operate at 3 to 5 knots ($\approx 2.0 \ m/s$) to maintain high-resolution imagery [48]. Given a nominal swath width of 70 m [49], this results in an approximate search coverage rate of:

$$C_{rate_v} = 2.0 \cdot 70.0 = 140.0 \ m^2/s. \tag{4.8}$$

See table 4.2 for a summary of each method's operational metrics.

Table 4.2: Comparison of search coverage and operational metrics.

Method	**Velocity (v)**	**Swath width (W)**	**Coverage rate (C_{rate})**
Rescue diver	0.5 m/s	4.0 m	2.0 m^2/s
Vessel (Side-Scan)	2.0 m/s	70.0 m	140.0 m^2/s
Seadragon 2.0	1.90 m/s	128.6 m	244.3 m^2/s

However, as the search coverage rate does not take into account the deployment time of each, nor the effective resolution of a human diver compared to typical sonar devices, these numbers alone do not indicate which method is more efficient in real-world SAR scenarios.

4.3 Detection and localization accuracy

After completion of the second stage fine-tuning process, the trained model achieved a recall score of 0.812, a precision of 0.868 and a mean average precision (mAP50) of 0.859 on the *body* object class (see Table 4.3). These result indicate that the model successfully learned to recognize human specific features, and was able to correctly identify most target bodies in the previously unseen validation set images. The lower performance (0.430) on the much stricter mAP50-95 score, which averages the mAP over ten different IoU (intersection over union) thresholds, from 0.5 up to 0.95, indicates that while the model is capable of finding the victims, the exact boundaries are slightly off.

To get a quantitative measure of the harmonic mean of the precision and recall on the body object class, we can calculate the F_1-score as follows:

$$F_1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} = 0.839, \tag{4.9}$$

indicating that the accuracy of the model remains high without an excessive number of false positives or negatives.

Table 4.3: Validation results per class. P = Precision, R = Recall.

Class	Images	Instances	P	R	mAP50	mAP50-95
All	52	114	0.769	0.686	0.753	0.448
Bag	9	10	0.574	0.600	0.655	0.465
Body	**48**	**48**	**0.868**	**0.812**	**0.859**	**0.430**
Chair	24	24	0.834	0.792	0.845	0.462
Dumbbell	5	5	0.939	0.600	0.649	0.403
Kettlebell	9	9	0.690	0.778	0.794	0.411
Net	9	9	0.795	0.667	0.802	0.459
Vacuum	9	9	0.685	0.556	0.667	0.506

The highest registered F_1-score across all objects was $F_1 = 0.71$, achieved at a confidence level of 0.267 (Fig. 4.3). This result reveal that while the model can achieve a high accuracy across different object categories, it is not very confident in its final predictions, likely due to the inherent noise present the sonar images.

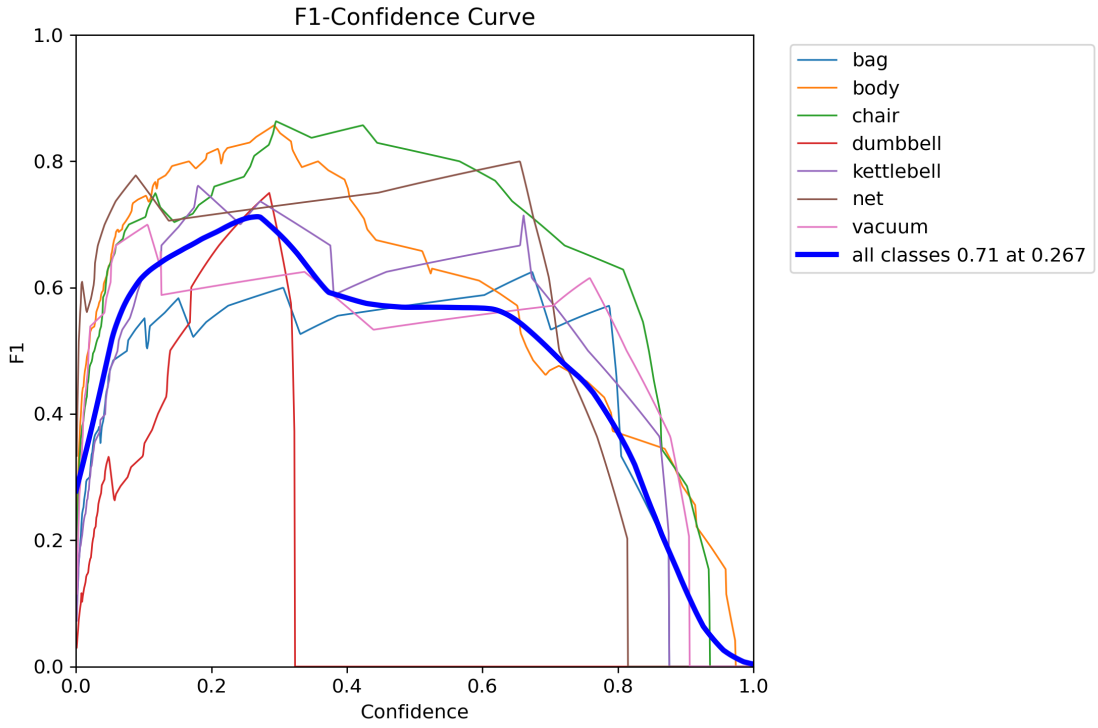

Figure 4.3: F_1-score vs confidence across all object classes.

In Fig. 4.4, a confusion matrix breaking down the model's performance by comparing the actual target values with those predicted by the machine learning algorithm is shown.

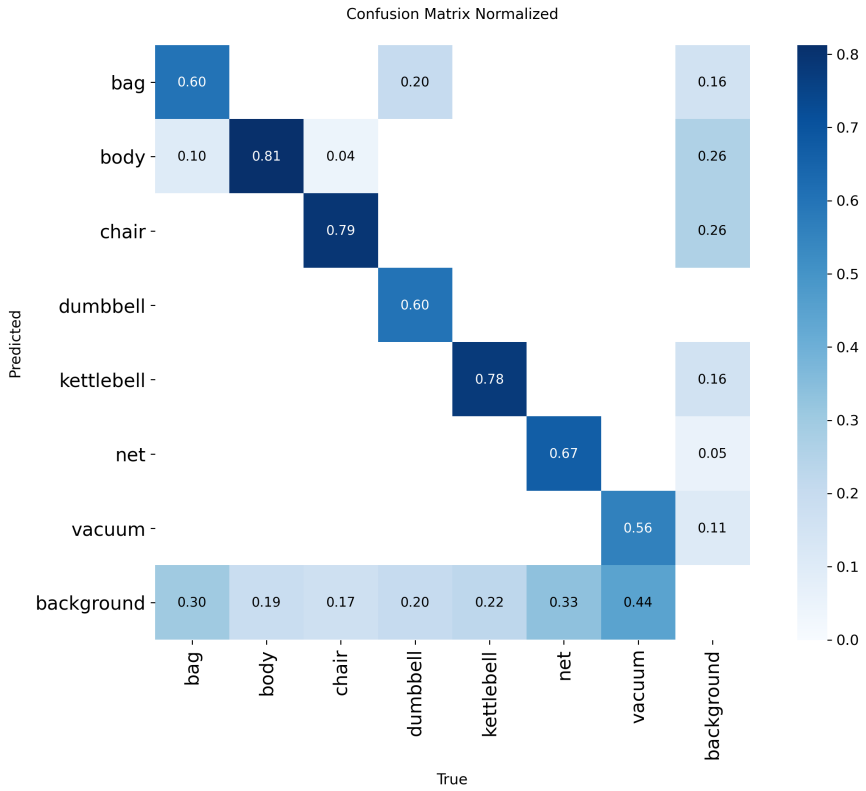

Figure 4.4: Consfusion matrix. y-axis: Predicted objects, x-axis: Ground truth.

As can be seen from the figure, the model is most effective at identifying the *body* object class, correctly identifying 81% of the actual human targets.

An example of a successful body identification with 99% confidence can be seen in Fig. 4.5.

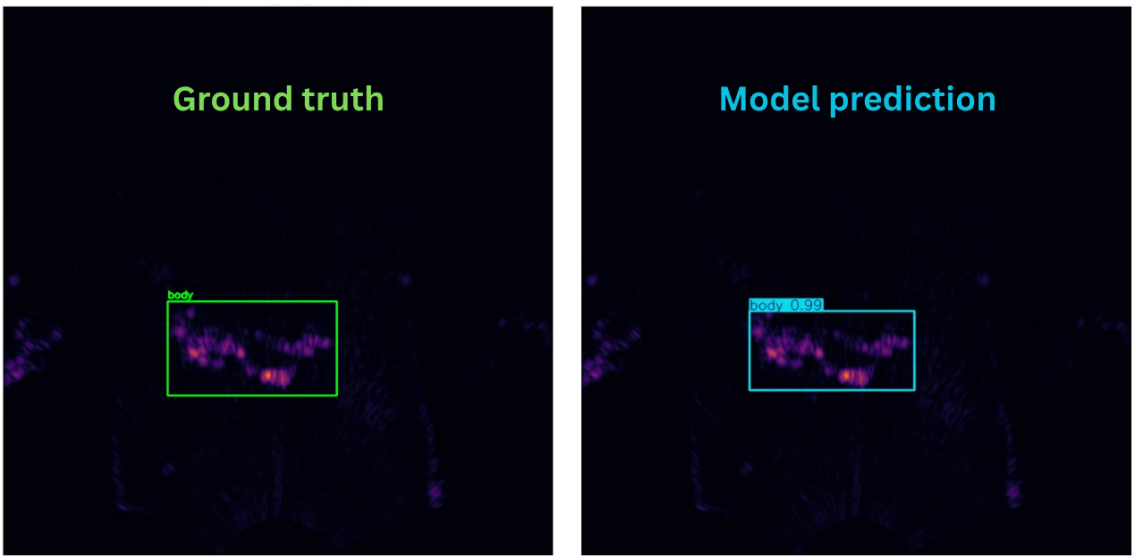

Figure 4.5: Ground truth vs. prediction.

While it never wrongly identified an actual body as another object, it still missed 19% of true targets, incorrectly assuming they were part of the background. Conversely, 26% of all body predictions were false predictions of background noise.

4.4 Computational efficiency

To evaluate the performance and suitability of the fine-tuned YOLO26n model for real-time SAR operations, a speed benchmark was conducted on the Jetson Orin Nano. The model was exported to a TensorRT engine with 16fp precision to maximize the utilization of the 1,024 CUDA cores and 32 Tensor cores available on the platform. The final benchmark was averaged over the 100 raw sonar images to accurately determine the expected real-time inference latency.

With the TensorRT engine, the system achieved a total average latency of 25.6 ms per image, compared to 643.2 ms while running on the CPU only. The equivalent throughput of approximately 39.1 FPS, well exceeds the 25 hz update rate of the BlueView sonar, ensuring no discernible lag for the remote operator during search manoeuvrers and leaving ample room for any additional overhead added during real-world deployment.

5

Discussion

In this chapter, the overall effectiveness of the developed USV platform will be discussed, highlighting both the potential operational benefits it could bring to current SAR operations as well as challenges facing real-life deployment. In addition, the performance of the deployed ML pipeline and design considerations of the architecture will be explored, as well as the inherent limitations of the acoustic modality. Lastly, a roadmap for future development is presented, outlining how the current prototype can evolve from a minimum viable product (MVP) into a fully autonomous and robust SAR asset.

5.1 Effectiveness of the lightweight USV concept

The revised Seadragon prototype demonstrates several characteristics that support the viability of lightweight USVs for time-critical underwater SAR operations. With a total dry weight of 17.81 kg and volumetric footprint below 0.16 m^3, the platform is sufficiently compact to be carried and assembled by a single operator. Compared to larger sonar-equipped vessels, this may provide practical advantages in situations where rapid deployment and operation in confined or shallow environments are required.

The modular construction also simplifies transport and maintenance during development. However, while the compact design improved portability, it also introduced constraints related to buoyancy, internal space, and power capacity. In particular, the limited volume introduced hydrostatic balancing concerns, especially if used in rougher waters.

The theoretical search coverage rate of 244.3 m^2/s suggests that the platform could realistically survey larger areas more efficiently that manual divers under certain conditions. At the same time, this estimate should be interpreted cautiously, as it assumes continuous operations under controlled conditions and does not account for environmental disturbances, operator workload, or reduced visibility. Additionally, although side-scan sonar vessels generally require more logistical support and may be less suitable for shallow environments, they remain significantly more mature and operationally validated than the prototype developed in this project.

The results therefore do not suggest that lightweight USVs such as the Seadragon 2.0 can replace existing SAR methods, but rather that they may serve as a complementary tool. In particular, a rapidly deployable platform may help close the capability gap between handheld sonar systems and larger vessel-based operations, especially in areas where rescue divers or specialized sonar crews are not immediately

available.

A major limitation of the current prototype is its operational endurance, which is estimated to be less than 15 minutes when operating at its maximum capability. This substantially limits the practical usefulness of the system in real SAR scenarios, where transportation between search zones, repeated sweeps, and prolonged operation often are required. Frequent battery replacement is also likely to interrupt operations and reduce the overall efficiency in the field, potentially negating all other benefits of the system.

Furthermore, the prototype was never fully validated in open-water rescue conditions. While the propulsion system, sonar pipeline, and machine learning components were individually tested, the absence of full-scale field trials means that several important factors remain unverified, including long-range communication stability, environmental robustness, localization accuracy, and operator usability during realistic missions.

Despite these limitations, the project demonstrates that a relatively low-cost and portable USV platform can integrate real-time sonar imaging and embedded machine learning within a compact form factor. Although considerable engineering and operational challenges remain, the concept appears sufficiently promising to justify further investigation and refinement.

5.2 Machine learning performance

While the trained YOLO26n model demonstrated strong performance on the validation set, achieving a mAP of 0.859 and an F1-score of 0.839, a qualitative analysis of the false positives and negatives suggests the model's practical utility exceed these statistical performance metrics.

In all five reported false positives, the model correctly localized portions of the target body but failed to satisfy the 0.5 IoU criterion for successful detection. Similarly, 4 of the 11 false negatives were marginally below this threshold. Additionally, as complete full body accuracy is not needed in a real life deployment scenario, all of these cases would correctly trigger an accurate body identification, likely resulting in a successful victim localization.

Furthermore, four additional missed detections involved scenarios where only the victim's legs were visible, significantly increasing the complexity of the identification task. Manual review of the source footage also revealed two instances where the model correctly identified targets that had been mislabelled during the annotation process (Fig 5.1), further indicating that the reported metrics represent a conservative estimate of the model's true real-world capabilities.

This also highlights the inherent difficulty in curating high-fidelity sonar datasets. Accurate labelling of acoustic imagery is often complicated by the tendency for reflections from disparate objects to merge or smear, obscuring clear boundaries. While integrating secondary modalities (e.g., optical cameras) during data collection, can provide ground-truth confirmation, this approach significantly increases the temporal and labour intensive demands of the annotation process. Furthermore, it introduces complex synchronization requirements across sensors, necessitating rigorous planning to ensure accurate alignment between the different modalities.

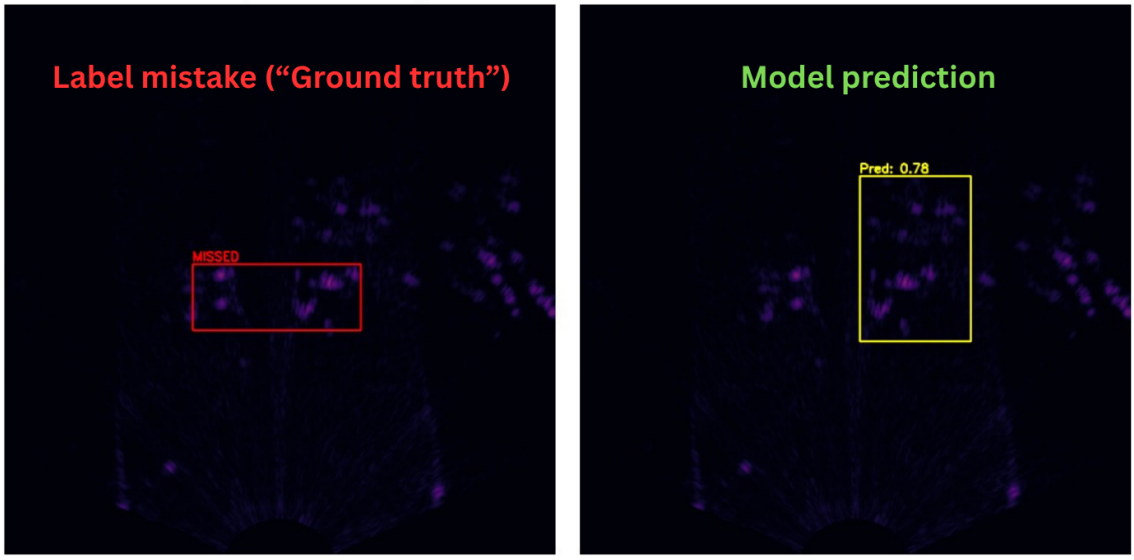

Figure 5.1: Example of labelling mistake (left) and correct prediction (right).

5.2.1 Limits and improvements of the deployed model

Although the deployed YOLO26n model achieved promising validation results within the scope of the collected dataset, its real-world robustness remains untested. Since the model was trained and evaluated exclusively on images from a clean pool environment with limited augmentations, its performance is likely to decrease in real-world deployment scenarios. Previously unseen seafloor objects may be mistaken for a human, and increased particulate backscatter could degrade image quality, altering the characteristics of the reflected sonar pings the model have learned to recognize. As a result, the detection accuracy observed during validation is unlikely to fully translate to operational deployment without further adaptation and testing.

The limited environmental diversity of the ADVD dataset also increases the risk of overfitting to specific acoustic characteristics of the recording environment. While inclusion of additional objects and noise augmentations likely improved generalization to some extent, the model may still rely on subtle scene specific features unintentionally correlated with the target class. The restricted size of the recording pool may for example have limited the models ability to learn how targets deform acoustically at greater distances. Unintentional echoes reflected from the narrow walls may also have introduced artifacts and noise to the objects that are unlikely to be present in real-world deployment environments. Recording additional datasets across larger distances and more varied viewing angles would therefore likely improve the localization consistency and overall robustness of the model.

Moreover, collecting data from real sea or lake environments would likely also contribute to improving the real-world accuracy of the model. However, such recordings also introduce challenges in high-quality data capture. The same environmental noise that confuses the model may also increase the risk of human error during the annotation process. Ambiguous sonar reflections and lower visual interpretability could drastically complicate the process of ensuring accurate object boundaries and may itself introduce noise into the training pipeline.

A more pragmatic approach may therefore involve increasing the diversity of the existing dataset through additional augmentations techniques. This could be achieved by applying further polarization filtering, adding stronger speckle noise and intensity distortions, or by utilizing more advanced style transfer techniques like CycleGAN. This could potentially expose the model to a wider range of acoustic appearances without requiring a large-scale field data collection.

Furthermore, the YOLO26n architecture was primarily chosen due to its excellent size-to-performance ratio. However, given that the measured 39.1 average FPS (via the TensorRT engine) significantly exceed the update rate of the sonar itself, the onboard Jetson Orin Nano likely has the computational headroom for larger and more capable models. Although transitioning to a larger model (e.g., YOLO26s or YOLO26m) likely would not have a discernible impact on the real-time detection process, it remains unclear whether increasing the model complexity alone meaningfully would improve real-world performance, particularly if the available training data remains limited in scale and diversity.

5.3 Future work

While the revised Seadragon prototype demonstrates the feasibility of a lightweight sonar-equipped USV platform, the system remain an early-stage proof-of-concept with several partially implemented or experimentally unverified subsystems. Future work should therefore focus not only on expanding functionality, but also on validating the reliability and operational practicality of the existing design.

5.3.1 Enhanced autonomy and localization

In its current state, the Seadragon 2.0 platform relies heavily on manual remote control and lacks the end-to-end autonomous capabilities set as a final goal of the project. While this was sufficient for subsystem testing and rapid prototyping, manual operation may become challenging in real-world SAR scenarios where operators must simultaneously monitor sonar imagery, maintain situational awareness, and coordinate with other responders.

Future iterations of the project could therefore focus on implementing more advanced localization and position algorithms, such as EKF-SLAM (extended kalman filter simultaneous localization and mapping), waypoint following, and automated raster search patterns. Successful integration of these would drastically increase the autonomy of the platform and minimize the cognitive load on first responders. Furthermore, it would allow the USV to be deployed with minimal procedural changes to existing routines, ensuring that the focus remains on the rescue.

5.3.2 Connectivity and communication stability

The current communication architecture was developed primarily for short-range testing in controlled environments. Consequently, its robustness during longer-range or obstructed operation remains largely untested.

Potential future improvements therefore include integration of LTE or sattellite connectivity, as well as more robust FHSS (frequency hopping spread spectrum) communication systems. Such upgrades could greatly increase the operational radius and ensure a high-bandwidth and stable data link for real-time sonar streaming, even in cluttered or electromagnetically challenging environments.

5.3.3 Hardware and user interface refinement

While the revised Seadragon design substantially improved the structural rigidity and modularity of the platform compared to the previous prototype, several potential areas of improvement were identified during development. The exposed aluminium frame and externally mounted components may be vulnerable to repeated impacts and corrosion. The flat and unstreamlined end caps add unnecessary drag, decreasing the overall efficiency of the craft. Future iterations could therefore focus on improving hydrodynamic efficiency and ruggedness of the USV through a more integrated hull design. Such modifications may improve endurance and handling characteristics, particularly in more challenging water conditions.

The current power system also presents several limitations. The use of separate lead-acid batteries in each hull simplified early development but resulted in a relatively high mass and limited operational endurance. Transitioning to lithium-ion batteries and implementing a unified power distribution system would considerably improve the energy density, reduce the overall weight, and simplify charging and maintenance of the prototype.

Furthermore, the sonar mounting system currently require manual angular positioning before deployment. This prohibits the operator from making live angular adjustments, which potentially hinder search operations where a target of interest would be better viewed from a different angle. To remedy this, the mount of the sonar could be motorized, allowing the operator to set the search angle as needed without pausing or further delaying the search operation.

Lastly, the user interface (UI) was developed solely for testing purposes and therefore lacks many features expected of a robust and intuitive deployable system. In its current form, the interface consists of a basic graphical display with target highlighting, while heartbeat and telemetry data are transmitted separately. Potential improvements include, enhanced contrast and visualization settings for improved human viewing, dedicated manual control functionality, and more unified presentation of system telemetry and system status data.

6
Conclusion

This thesis investigated the potential of a lightweight, sonar-equipped USV to enhance time-critical SAR operations. Existing subsurface search methods face deployment and efficiency challenges, limiting their effective use. The work aimed to to evaluate whether a rapidly deployable platform, combined with machine learning based sonar image processing, could improve detection and localization of submerged human-sized targets.

The developed Seadragon 2.0 prototype demonstrates that a compact and modular USV can feasibly support rapid deployment and wide-area search operations. The redesigned system achieves a low overall weight and a form factor suitable for single-operator handling. Theoretical estimates of search coverage suggest a substantial improvement over traditional diver-based methods, while requiring substantially less resources than side-scan sonar vessel alternatives. However, as the estimates are derived from controlled experiments and analytical models rather than full-scale field trials, the operational effectiveness of the system in real-world environments remains to be validated.

The results further show that a USV-mounted imaging sonar, combined with a fine-tuned deep learning model, can successfully detect and localize submerged human-sized targets in controlled conditions. The trained YOLO-based model achieved a strong detection performance on the custom acoustic dataset and demonstrated real-time inference capability on the embedded hardware, indicating it effectively can extract relevant features from the acoustic imagery.

A major limitation of the work involves the lack of a full system validation. The complete assembled system was never tested in real open-water conditions, limiting key performance metrics such as search coverage and operational endurance to be based partly on theoretical estimates. Moreover, the machine learning model was trained and evaluated exclusively on data collected in a controlled pool environment, which does not fully capture the variability and noise expected in natural bodies of water. In combination with the relative small custom dataset size and limited range of object variations, the robustness and generalizability of the detection results can be restricted.

Future work should prioritize comprehensive field validation to asses system performance under realistic environmental conditions. Expanding the dataset with more diverse sonar recordings and synthetic images would improve the model robustness and should be considered to enable more reliable real-world deployment. To transition the platform from a proof-of-concept to a practical SAR asset, further development of more energy efficient hardware, autonomous navigation capabilities, and improved user interfaces would also be necessary.

In conclusion, this thesis demonstrates that a lightweight, sonar-equipped USV constitutes a viable approach to improving underwater search and rescue operations, with the potential to improve search efficiency in time-critical scenarios, while also highlighting the need for further development and real-world validation.

Bibliography

[1] W. H. Organization, "World Drowning Prevention Day," WHO Campaigns, 2024, accessed: 2024-05-22. [Online]. Available: https://www.who.int/campaigns/world-drowning-prevention-day

[2] U. N. G. Assembly, "Resolution 75/273: Global drowning prevention," UN Doc. A/RES/75/273, apr 2021, adopted on 28 April 2021. [Online]. Available: https://docs.un.org/en/A/RES/75/273

[3] W. H. Organization, "Global report on drowning: preventing a leading killer," World Health Organization, Geneva, Switzerland, Tech. Rep., 2014, accessed: 2024-05-22. [Online]. Available: https://iris.who.int/handle/10665/143893

[4] S. Livräddningssällskapet, "Drunkningsstatistik," https://svenskalivraddningssallskapet.se/drunkningsstatistik/, 2025, accessed: 2025-11-28.

[5] Rättsmedicinalverket, "Rättsmedicinalverket presenterar unik forskningsstudie om drunkningar i samarbete med svenska livräddningssällskapet," https://www.rmv.se/aktuellt/rattsmedicinalverket-presenterar-unik-forskningsstudie-om-drunkningar-i-samarbete-med-svenska-livraddningssallskapet/, 2025, accessed: 2025-11-28.

[6] J. E. Johansson, "Search optimization rescue assistance during drowning incidents," Master's thesis, Umeå University, Umeå Institute of Design, 2021, independent thesis, Advanced level (Master Two Years, 80 credits / 120 HE credits). [Online]. Available: https://www.diva-portal.org/smash/record.jsf?pid=diva2:1577369

[7] L. Quan, C. D. Mack, and M. A. Schiff, "Association of water temperature and submersion duration and drowning outcome," *Resuscitation*, vol. 85, no. 6, pp. 790–794, 6 2014, epub 2014 Mar 4. Erratum in: Resuscitation. 2014 Sep;85(9):1304.

[8] A. Claesson, J. Lindqvist, P. Ortenwall, and J. Herlitz, "Characteristics of lifesaving from drowning as reported by the swedish fire and rescue services 1996–2010," *Resuscitation*, vol. 83, no. 9, pp. 1072–1077, 2012.

[9] RäddSam F, *Riktlinjer för vattenlivräddning inom RäddSam F*, RäddSam F, 2024, gäller för RäddSam F:s medlemskommuner i Jönköpings län. [Online]. Available: https://raddsamf.se/app/uploads/2019/09/Riktlinjer-vattenlivraddning-raddsamf-2024.pdf

[10] Statens räddningsverk, *Räddningsdykning*, D. Hedberg, Ed. Statens räddningsverk (nuvarande MSB), 2004. [Online]. Available: https://rib.msb.se/Filer/pdf/19618.pdf

[11] Sjöfartsverket, “Samverkan för att rädda liv,” Sjöfartsverket, May 2024, senast uppdaterad: 2024-05-08. [Online]. Available: https://www.sjofartsverket.se/sv/tjanster/sakert-batliv/samverka-for-att-radda-liv/

[12] S. Radio, “Sonarutrustning används för att rädda liv vid drunkning,” Sveriges Radio, February 2026, accessed: 2026-02-13. [Online]. Available: https://www.sverigesradio.se/artikel/sonarutrustning-anvands-for-att-radda-liv-vid-drunkning

[13] ——, “Brist på räddningsdykare – länet satsar på fridykare,” Sveriges Radio, February 2026, accessed: 2026-02-16. [Online]. Available: https://www.sverigesradio.se/artikel/brist-pa-raddningsdykare-lanet-satsar-pa-fridykare

[14] A. J. Murphy, M. J. Landamore, and R. W. Birmingham, “The role of autonomous underwater vehicles for marine search and rescue operations,” *Underwater Technology*, vol. 27, no. 4, pp. 195–205, June 2008.

[15] B. R. Inc., “Blue robotics: Marine robotics components and systems,” 2025, accessed: February 24, 2025. [Online]. Available: https://bluerobotics.com/

[16] D. T. Inc., “Remotely operated vehicles (ROVs),” 2025, accessed: February 24, 2025. [Online]. Available: https://www.deeptrekker.com/shop/categories/rovs

[17] V. LLC, “VideoRay underwater ROV products and systems,” 2025, accessed: February 24, 2025. [Online]. Available: https://videoray.com/products/

[18] ——, “Search & rescue success in colorado: Summit county case study,” VideoRay LLC, Pottstown, Pennsylvania, USA, Tech. Rep., 2018, case study on the use of VideoRay Pro 4 ROV by Summit County Water Rescue Team. [Online]. Available: https://videoray.com/wp-content/uploads/2021/09/Summit-County-Case-Study-2018.pdf

[19] R. Doornekamp, “Search and recovery: The houston county rescue,” Deep Trekker News, October 2019. [Online]. Available: https://www.deeptrekker.com/news/houston-county-rescue

[20] Sterling Communications, “Videoray pro 4 rack base rov system,” Sterling Communications Product Page, 2025, beskriver konfigurationen av Pro 4 ROV för rackmontering i fartyg eller ledningscentraler. [Online]. Available: https://sterlingcomms.com/product/videoray-pro-4-rack-base-rov-system-2/

[21] D. T. Inc., “Rov buying guide,” 2 2026, accessed: February 24, 2026. [Online]. Available: https://www.deeptrekker.com/resources/rov-buying-guide

[22] VideoRay LLC, *VideoRay Pro 4 Operator’s Manual*, VideoRay LLC, 2022, copyright © 2022, VideoRay LLC. [Online]. Available: https://download.videoray.com/documentation/pro_4/html/overview_.html

[23] Hydronalix, “E.m.i.l.y. | emily rescue robot,” Official Website, 2025, emergency Integrated Lifesaving Lanyard (EMILY) for beach, river, and flood rescue. [Online]. Available: https://www.emilyrobot.com/

[24] Agistar, “Bx-usv ii | next-generation usv,” Official Product Page, 2025, autonomous surface vehicle for water quality monitoring, sonar mapping, and bathymetry. [Online]. Available: https://www.agistar.fr/en/bx-usv2

[25] Seaber, “Marvel-scan | compact uuv with sss and acoustic positioning,” Official Product Page, 2025, autonomous underwater vehicle equipped with Side Scan Sonar (SSS) and DVL for bathymetric surveys and search and rescue. [Online]. Available: https://seaber.fr/marvel-uuv-scan

[26] R. J. Urick, *Principles of Underwater Sound*, 2nd ed. New York: McGraw-Hill, 1975.

[27] Blueprint Subsea, *Oculus M-Series Single & Dual Frequency Sonars Datasheet*, Blueprint Subsea, Ulverston, Cumbria, UK, 2023, rev: DA-148-P01443-04. [Online]. Available: https://download.videoray.com/documentation/mss/pdf/blueprint_subsea/oculus/Oculus_M-Series_Datasheet_EN.pdf

[28] E. O. Belcher, W. H. Hanot, and J. Burch, "Dual-frequency identification sonar (didson)," in *Proceedings of the 2002 International Symposium on Underwater Technology*, 2002, pp. 187–192.

[29] H.-T. Nguyen, E.-H. Lee, and S. Lee, "Study on the classification performance of underwater sonar image classification based on convolutional neural networks for detecting a submerged human body," *Sensors*, vol. 20, no. 1, p. 94, 2020.

[30] K. O. Amer, M. Elbouz, A. Alfalou, C. Brosseau, and J. Hajjami, "Enhancing underwater optical imaging by using a low-pass polarization filter," *Optics Express*, vol. 27, no. 2, pp. 621–643, 2019.

[31] Y. Hao, Y. Yuan, H. Zhang, and Z. Zhang, "Underwater optical imaging: Methods, applications and perspectives," *Remote Sensing*, vol. 16, no. 20, p. 3831, 2024.

[32] L. T. Hatch and A. J. Wright, "A brief review of anthropogenic sound in the oceans," *International Journal of Comparative Psychology*, vol. 20, no. 2, pp. 121–133, 2007. [Online]. Available: https://escholarship.org/uc/item/5cj6s4r9

[33] G. Yang, L. Dai, and Z. Wei, "Challenges, threats, security issues and new trends of underwater wireless sensor networks," *Sensors*, vol. 18, no. 11, p. 3907, 2018. [Online]. Available: https://www.mdpi.com/1424-8220/18/11/3907

[34] F. Francisco, A. Bender, and J. Sundberg, "Use of multibeam imaging sonar for observation of marine mammals and fish on a marine renewable energy site," *PLOS ONE*, vol. 17, no. 12, p. e0275978, 2022. [Online]. Available: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0275978

[35] K. Xie, J. Yang, and K. Qiu, "A dataset with multibeam forward-looking sonar for underwater object detection," *Scientific Data*, vol. 9, no. 1, p. 739, 2022.

[36] S. Hu and T. Liu, "Underwater rescue target detection based on acoustic images," *Sensors*, vol. 24, no. 6, p. 1780, 2024.

[37] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, F. Pereira, C. Burges, L. Bottou, and K. Weinberger, Eds., vol. 25. Curran Associates, Inc., 2012. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf

[38] I. D. Mienye, T. G. Swart, G. Obaido, M. Jordan, and P. Ilono, "Deep convolutional neural networks in medical image analysis: A review," *Information*, vol. 16, no. 3, p. 195, 2025.

[39] L. K. Sahoo and V. Varadarajan, "Deep learning for autonomous driving systems: technological innovations, strategic implementations, and business implications - a comprehensive review," *Complex Engineering Systems*, vol. 5, no. 2, p. 2, 2025.

[40] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.

[41] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 5998–6008. [Online]. Available: https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf

[42] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 580–587.

[43] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788.

[44] J. Lam, A. Chen, A. Bennett, and M. Triantafyllou, "Propeller characterization testing of a blue robotics t200 thruster," in *OCEANS 2023 - Limerick*. IEEE, 2023.

[45] B. S. S. Inc., "Blueview m900-mk2 high-resolution 2d imaging sonar," Product Overview, 2024, accessed: 2026-03-07. [Online]. Available: https://beyond-srv.com/products/blueview-m900-mk2-high-resolution-2d-imaging-sonar

[46] Ultralytics LLC, "Ultralytics yolo26," https://docs.ultralytics.com/models/yolo26, 2026, accessed: 2026-05-16.

[47] K. Xie, J. Yang, and K. Qiu, "A dataset with multibeam forward-looking sonar for underwater object detection," *Scientific Data*, vol. 9, no. 1, Dec. 2022. [Online]. Available: http://dx.doi.org/10.1038/s41597-022-01854-w

[48] Blueprint Design Engineering Ltd. (2017) Tips for good imagery – blueprint subsea starfish 453oem system manual. Accessed: 2026-05-05. [Online]. Available: https://www.manualslib.com/manual/1573430/Blueprint-Subsea-Starfish-453oem.html?page=17

[49] iRov Underwater Services. (n.d.) Starfish 990f advanced towed sidescan system. Accessed: 2026-05-05. [Online]. Available: https://www.irov.it/en/products/sonars/starfish-sidescan-sonars/starfish-990f-advanced-towed-sidescan-system/

A

Appendix 1

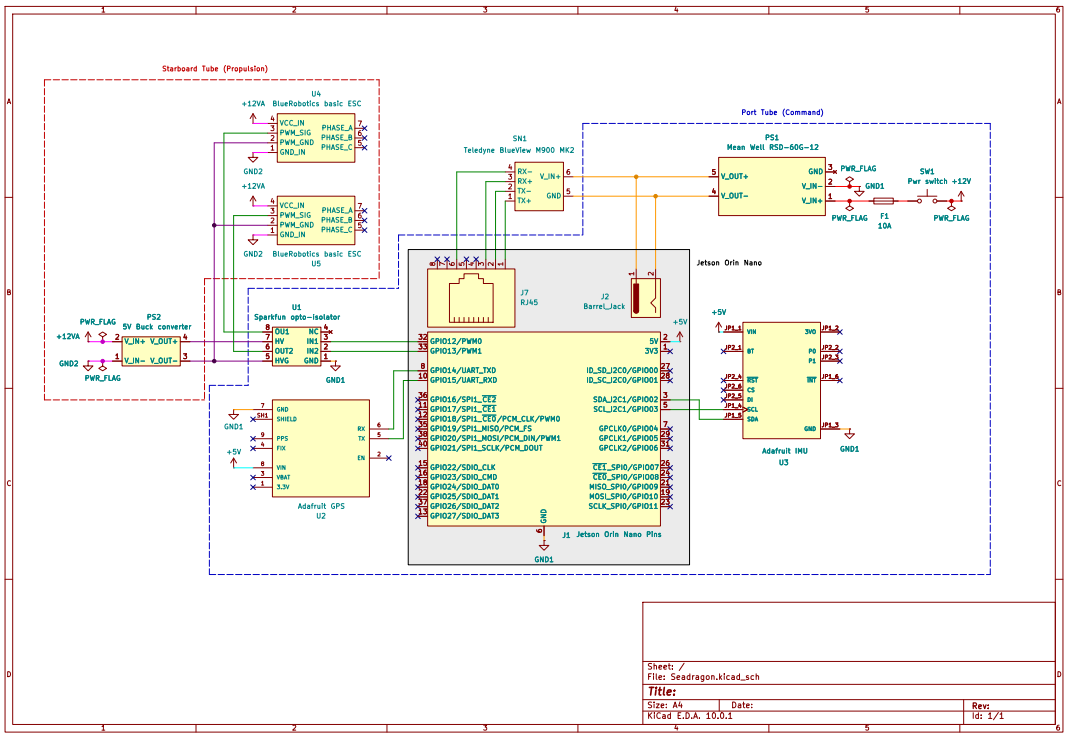

Figure A.1: Complete schematic of the electrical system of the Seadragon 2.0 prototype.

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY