

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Methods for In-Situ Fish Length Measurement with a Single Laser and Camera to Support Citizen Science

Permalink

<https://escholarship.org/uc/item/41x511ht>

Author

Hu, Kyle

Publication Date

2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Methods for In-Situ Fish Length Measurement with a Single Laser and Camera to Support
Citizen Science

A thesis submitted in partial satisfaction of the
requirements for the degree Master of Science

in

Electrical Engineering (Intelligent Systems, Robotics and Control)

by

Kyle Sirong Hu

Committee in charge:

Professor Curt Schurgers, Chair
Professor Ryan Kastner, Co-Chair
Professor Nikolay Atanasov

2024

Copyright

Kyle Sirong Hu, 2024

All rights reserved.

The Thesis of Kyle Sirong Hu is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2024

DEDICATION

To Kat, my friends, and my family.

EPIGRAPH

One fish, Two fish, Red fish, Blue fish,
Black fish, Blue fish, Old fish, New fish.
This one has a little car.
This one has a little star.
Say! What a lot of fish there are.

Dr. Seuss

TABLE OF CONTENTS

Thesis Approval Page	iii
Dedication	iv
Epigraph	v
Table of Contents	vi
List of Figures	viii
List of Tables	x
Acknowledgements	xi
Abstract of the Thesis	xiii
Chapter 1 Introduction	1
1.1 Current Methods for Fish Length Studies	2
1.2 Citizen Science	3
1.3 Related Technologies	4
1.3.1 Acoustic Methods	4
1.3.2 Stereo Camera Systems	5
1.3.3 Laser Calipers	6
1.3.4 Laser Rangefinding	7
1.4 Outline	8
Chapter 2 System Overview	9
2.1 Accuracy Bounds	9
2.2 Hardware Setup	9
2.2.1 Camera	10
2.2.2 Wide Angle Lens	10
2.2.3 Laser and Mount	11
2.3 Operation	11
2.3.1 Calibration Procedures	12
2.3.2 Diver Operation	12
2.3.3 Data Offloading	13
2.4 Software	13
2.4.1 Raw Image Processing	13
2.4.2 Finding the Laser Dot	16
2.4.3 Finding Fish Head and Tail Points	17
2.5 Acknowledgements	17
Chapter 3 Algorithm Details	18

3.1	Pinhole Camera Model	18
3.2	Quantities and Conventions	19
3.3	Finding the Distance to the Fish.....	21
3.4	Finding Laser Parameters	24
3.4.1	Method 1: Gauss Newton Optimization.....	24
3.4.2	Method 2: Averaging	27
3.5	Calibration Procedure	28
3.6	Acknowledgements	29
Chapter 4	Experiments	31
4.1	Pool Experiments.....	31
4.1.1	Reference Objects	32
4.1.2	Laser Calibration Algorithm	33
4.1.3	Distance Measurements	34
4.1.4	Length Measurements	35
4.1.5	Contributions of Tilt to Error	39
4.1.6	Laser Attenuation	41
4.2	Field Experiments	42
4.2.1	Salt Water Measurements	43
4.2.2	Laser Calibration With A Slate	44
4.2.3	Fish Behavior Study	46
4.2.4	Points of Mechanical Failure	47
4.3	Summary of Results	48
4.4	Acknowledgements	48
Chapter 5	Conclusion	50
5.1	Summary	50
5.2	Future Work	50
Appendix A	Supplementary Derivations	52
A.1	Derivation of Scale Factor	52
A.2	Derivation Of Jacobian	54
Appendix B	Pool Test List	56
Appendix C	FishSense Lite User Manual	57
C.1	Setup	57
C.2	Use	57
C.3	Teardown	58
C.4	Guidelines	58
C.5	Camera Settings.....	58
Appendix D	Light Attenuation In Water	60
Bibliography	63

LIST OF FIGURES

Figure 2.1.	A fully assembled FishSense Lite hardware unit.....	10
Figure 2.2.	Images of a checkerboard with (left) and without (right) the corrective optic.	11
Figure 2.3.	Flowchart of software processing pipeline.....	14
Figure 2.4.	Illustration of the Bayer pattern on a typical digital camera sensor. From Wikimedia Commons [16].	15
Figure 2.5.	Comparison between camera-processed JPEG image (left) and our processed PNG (right).	16
Figure 3.1.	Illustration of the pinhole camera model. Adapted from [18].	19
Figure 3.2.	Diagram of the image plane, demonstrating pixel pitch w and apparent object length n .	20
Figure 3.3.	Pinhole diagram showing a lionfish with it's image projection.	20
Figure 3.4.	Diagram demonstrating known quantities and unknown parameters. Laser shown in solid green, imaged laser dot in dashed green.	22
Figure 3.5.	An example of a slate calibration image.	29
Figure 4.1.	The aluminum box. The top edge of the black tape across the surface is what we typically measure.	32
Figure 4.2.	Fred, our first fake fish. He, and all the others, are 31cm long.	33
Figure 4.3.	Comparison of two calibration methods, graphing estimated distance against estimated length.	34
Figure 4.4.	Distance estimates compared to a distance reference, measured with two different laser parameter sets. 20% error margin included.	35
Figure 4.5.	Measurements of the fish and the box, plotted over estimated distances away from the camera. 20% error margin included.	36
Figure 4.6.	A demonstration of data where the laser parameters calculated during calibration do not match the parameters during measurement.	37
Figure 4.7.	Measurements taken with the same laser parameters, with differing focal lengths.	38

Figure 4.8.	Illustration of lengths measured on checkerboard. Each square pictured here is 41.3mm wide.	39
Figure 4.9.	Percent error of length measurements over all distances and angles, grouped by true object length.	40
Figure 4.10.	Percent error of length measurements over all distances and angles, plotted over the distance from the camera each image was taken.	40
Figure 4.11.	Variances of percentage length error, with each data point being a single image.	41
Figure 4.12.	Lasers circled in red from left to right: Shark laser (green), Innovation Scuba, Orca, Gold	42
Figure 4.13.	A freshwater vs saltwater comparison of measured fake fish lengths.	44
Figure 4.14.	The length measured in the slate calibration experiment.	45
Figure 4.15.	Comparisons of slate measurements across different devices.	45
Figure 4.16.	A male California Sheephead from the La Jolla Kelp Beds. Note the “lightsaber” effect caused by particulates in the water.	46
Figure 4.17.	Image of one example of mount failure.	47
Figure D.1.	Absorption coefficients of pure water over the electromagnetic spectrum. From Wikimedia Commons [15].	61
Figure D.2.	Absorption coefficients of pure water in just the visible part of the spectrum. Adapted from Wikimedia Commons [14].	61
Figure D.3.	Colored rods photographed on the surface and at depth. From PBS Learning Media [33].	62

LIST OF TABLES

Table 1.1.	Comparison of different underwater length measurement techniques. The approach that FishSense Lite uses is in bold.	7
Table B.1.	List of pool tests conducted, with their purpose	56

ACKNOWLEDGEMENTS

This project would not have been possible without the many hands available to help us throughout these past couple of years.

Thank you to the members of my committee: Curt Schurgers, for his invaluable guidance on writing this document, Ryan Kastner, for guiding us through the bureaucracies of research funding, and Nikolay Atanasov, for helping me smooth out some of the math-heavy portions of this document.

Thank you to Christopher Crutchfield, whose mentorship and companionship has been invaluable to both me personally and to the project overall. Thank you to Nathan Hui, whose stern engineering and management advice was vital to us making it this far. Thank you to all former and current undergraduates who have worked with the FishSense project: Vivaswat Suresh, Raghav Maddukuri, Avik Ghosh, Allen Kan, Ana Perez, Hamish Grant, Harish Vasanth, Shaurya Raswan, Kyle Tran, Sam Prestrelski, Adrian Zugehar, Claire Wang, Jennifer Xu, Louis King, and Grace Liu. Thank you also to the REU students who helped us gather valuable data despite unreasonably short notice: Josie Dominguez, Jordan Reichhart, Ela Lucas, and Nick Reyes.

Thank you to Professor Brice Semmens, Jack Elstner, Alli Candelmo, Jen Loch and the members of REEF for testing this system and providing us with the initial problem we aim to solve. An additional thank you to Patrick Paxson and Peter Tueller, whose work laid the foundation for our current research.

Most importantly, thank you to Fred the fake fish, who we dearly miss. May he enjoy his freedom in the Pacific Ocean.

I would also like to thank the people I've met during my time at UCSD, who have made this journey worth it.

Thank you to Kristina "Kat" Diep, for always being encouraging and supportive during my time in graduate school. Your humor and light-heartedness has kept me going.

Thank you to Matei Gardus, who taught me how to use emacs org-mode several years

ago, and has helped me tremendously with the formatting aspect of writing this document.

Thank you to the members of UCSD Quizbowl, in particular to Alistair Gray, who is a dear friend, and to Praveen Nair, who asked me to include the Dr. Seuss quote in the epigraph.

Thank you the “ibob” crew, who have fostered my love for technology and with who I have created a thriving community of people with similar interests through UCSD’s ACM chapter.

Thank you to Daniel Truong, as well as both the San Diego Smash and Fighting Game communities, for teaching me how to take a loss.

Finally, thank you to my parents, whose continuous support has allowed the opportunity to do all of this.

Chapters 2, 3 and 4 contain material that is being prepared for publication. Christopher L. Crutchfield, Kyle S. Hu, Vivaswat Suresh, Ana I. Perez, Avik Ghosh, Hamish J. Grant, Kyle Tran, Samantha Prestrelski, Ronan Wallace, Nathan Hui, Jack Elstner, Dylan Heppell, Jennifer Loch, Alli Candelmo, Brice X. Semmens, Curt Schurgers, and Ryan Kastner. The thesis author was one of the primary investigators and authors of this material.

I wish the FishSense team the best of luck on the road ahead. Thank you again to everyone for accompanying me on this journey.

ABSTRACT OF THE THESIS

Methods for In-Situ Fish Length Measurement with a Single Laser and Camera to Support
Citizen Science

by

Kyle Sirong Hu

Master of Science in Electrical Engineering (Intelligent Systems, Robotics and Control)

University of California San Diego, 2024

Professor Curt Schurgers, Chair
Professor Ryan Kastner, Co-Chair

Ecologists are interested in studying fish length distributions as a metric for the health of fish populations, both for fishery management and to gauge the effect of policies and worsening ocean conditions. Current methods of gathering these data are time intensive, and require expensive equipment and/or training. Approaching this issue with citizen science – utilizing efforts from the general public – allows us to gather much more data about fish populations than what can currently be acquired, at a reasonable cost. We are working towards developing a device and software platform that allows recreational divers to take pictures of fish, and have the images processed such that length and species information are determined automatically.

This work focuses on the mechanism through which fish length is calculated. The device incorporates a standard laser pointer with a consumer waterproof camera, and uses the location of the laser dot in the object to determine the distance between the object and the camera. We also discuss methods to ensure that the parameters of the laser with respect to the camera are known, as well as how we can ensure that these procedures can be performed in the field. We test these by conducting trial calibrations and length measurements in a controlled setting, and evaluating where sources of error come from. From our testing we are able to verify that our technology is capable of obtaining measurements within 20% of of true length, similar to the error margins of visual estimation.

Chapter 1

Introduction

Marine ecosystems provide humankind many resources which we depend on. These ecosystems an important source of food worldwide, and approximately 3 billion people depend on seafood as a significant source of protein [22]. Data that can be used to evaluate the health of our oceans is therefore essential for us to understand the potential impacts of overfishing and climate change.

Fish in particular are a significant part of this – fish populations are known to be in decline as a result of overfishing and climate change, which has dire consequences for both food stocks and marine ecosystems as a whole [27, 36]. In order to be able to quantify the effects of climate change on fish, we require some metric with which we can study how healthy fish are on a population level.

One such way of doing this is to obtain a fish length distribution [26, 45]. By examining such distributions over time, scientists can identify whether a population is being overfished, or is recovering. There are two primary important applications for these data. The first is fish population conservation – if policies are put into place to assist with the recovery of a fish population, it is within the interest of scientists to document the success of these protective measures. Our collaborators have performed a study of a particular species of Nassau Grouper, and have done this fairly successfully by looking at annual spawning aggregations [26]. The second is fishery management. Fisheries want to ensure that they extract as much catch as

possible without irreparably damaging their stock, meaning that detailed, real-time information about population health is vital, especially where other data are unavailable [3].

1.1 Current Methods for Fish Length Studies

Current global fish length data often come from catch and release programs. In California, the California Collaborative Fisheries Research Program (CCFRP) is one such program, which as of 2024, consists of over 1800 volunteer anglers. These anglers use hook-and-line procedures to catch the fish, from which length is measured directly with a tape measure before the fish is released. This length information is combined with metrics like catch per unit effort to estimate the biomass of a population. However, fishing takes time – in their most recent study, each examined location took 3 days to gather data for [52]. In addition, stress caused by barotrauma and air exposure can cause fish to be less able to find food or evade predators, which reduces their life expectancy [11]. In summary, these expeditions are logistically difficult to set up, and are also harmful to the fish being studied [40, 11].

Catch and release fishing also poses an issue for marine protected areas (MPAs), within which fishing may be restricted to certain species, or completely illegal outright. To study populations within these areas, adjacent fishing activities are sometimes used to infer the abundance of a particular species [52]. Thus, in order to better understand populations within MPAs, a non-invasive method is desirable.

A non-invasive alternative is the roving diver survey, where a team of divers conducts a visual census of fish species and their lengths within the MPA. Sometimes divers may use a T-bar as a length reference. However, measurements obtained purely from visual estimation are naturally imprecise – humans are only capable of estimating to within 20% of a fish's true length [24]. This is also a skill that must be honed in order to be improved, and requires retraining as frequent as once every 6 months [5].

In summary, state of the art approaches require resources for data collection, whether it

be personnel, time, money or specialized equipment. To aid in efforts to study fish populations, we aim to develop a more accessible data collection method.

1.2 Citizen Science

Our approach seeks to take advantage of efforts from the general public in order to gather as much data as possible. This is known as citizen science. Due to the overhead of the procedures or equipment required to study fish in particular, there are no methods that we know of currently that are accessible to the general public, which we hope to address with our current work.

This approach to data collection has been tried and tested in other fields. Foldit¹ is one such example that has been used to identify new protein structures, using a video-game like interface to get players to solve puzzles related to protein folding. Citizen scientists using Foldit have been shown to be able to recreate existing protein structures, and even provide new ways of creating proteins that professional scientists may not consider [30].

A example related to species identification is iNaturalist, a smartphone application that enables its users to capture images of terrestrial flora and fauna, and upload them to an online platform. A community of naturalists are then able to view these images and identify the species of the subject. This community sourced dataset is showing promise in studies of individual taxa [42], discoveries of new species [49], and species distributions of exotic species [48]. Due to this being a smartphone app, these data primarily consist of terrestrial life, or life that is otherwise only visible above water.

We hope to both provide additional data for fish specifically, as well as provide length data that can be analyzed. In order for such an approach to be useful for citizen scientists, it must be able to be accessed at relatively low expense, and require minimal training to use properly. Subsequent sections will detail how we achieve this.

¹<https://fold.it/>

1.3 Related Technologies

In order for our approach to require minimal training to use, we turn to a technological approach. Training recreational divers to visually estimate fish lengths is impractical, and so we aim to develop a tool that they can carry underwater. Herein we discuss current technologies that exist.

The effort described within this work relates to our solution to the aforementioned problem – a device that we call FishSense Lite. FishSense Lite relies on a single laser as depth reference in order to measure the length of a fish from a captured image, and due to the simplicity and accessibility of the hardware construction, it can be utilized by recreational divers to get much larger data coverage. The primary contributions of this thesis are the algorithms behind this device, i.e. how this single laser is used to measure fish length. However, first we will review the existing state of the art. A summary of these technological approaches can be seen in Table 1.1, including FishSense Lite.

1.3.1 Acoustic Methods

The primary related technology in this category is sonar, which involves actively sending sound waves and using the signal as it echoes back to construct an image of the scene. As we are interested in studying particular fish populations, it is also important that we are able to identify the species of fish being imaged, a relatively recent development when it comes to acoustic methods. Technologies like Dual-frequency Identification Sonar (DIDSON) [4] and Adaptive Resolution Imaging Sonar (ARIS) [29] both take advantage of many individual sonic beams and high frequencies in order to construct a high-fidelity sonic image of the scene. Both DIDSON and ARIS technologies have been used to measure fish lengths [10, 17] and identify fish species [31, 29]. Sonar is far more effective than camera-based techniques when visibility conditions are poor. However, they are by far the most expensive option listed here. A system in 2006 was said to have cost around 20,000USD [35]. We were unable to find specific prices for

modern imaging sonar units, though prices have ranged from tens of thousands to even hundreds of thousands of US dollars.

Passive techniques – monitoring sounds that come from the fish themselves – have also more recently been used as a way to study fish populations [21]. These have shown to be an effective way to study species distribution [47], or for recording sightings of a targeted species [7], though we are unaware of any existing literature that attempts to estimate size using this technology.

1.3.2 Stereo Camera Systems

The approach that is easiest for citizen scientists to carry out is to first capture an image of the fish, using one or more cameras, then use the distance to the fish to scale its apparent size appropriately. The main problem associated with these approaches is finding the distance to the fish.

One way of doing so involves stereo video technology [32]. To measure the actual size of objects in an image, an additional camera can be added to create a stereo camera setup - provided the relationship between the two cameras is known, apparent size can be converted to actual size by leveraging the disparity in pixel location between the same features in both images. These are typically diver-operated (known as stereo diver-operated video or stereo-DOV) [23] or placed in baited remote underwater video systems [32]. While stereo-DOV is a more cost-effective solution than deploying a remote system, the current state of the art still requires purchasing proprietary hardware and software [23], which can be prohibitively expensive for a citizen scientist at a minimum of 4,600USD [43] for a scientist grade stereo video system. In addition, stereo video generates a large amount of data that requires significant effort to store and process [46].

Commercial stereo video solutions include the AQ1 AM100 [44] and the AKVA Vicass HD [13], intended for use in aquaculture. Such systems are also costly and require a tether to a surface-side computer with proprietary software that must be used to manage the system. This limits the regions of the world where the data can be collected as it requires scientists to interact

with the system. The tether also limits the depths at which the data can be collected.

Our team has also attempted to use proprietary stereo camera units to build a FishSense “Pro” device [39, 46]. These attempts involved using an Intel Realsense D455 – a depth camera that used both stereo cameras and structured light in order to infer depth information. This camera was paired with a NVIDIA Jetson TX2 for onboard video processing, and was fully contained in a waterproof enclosure with onboard power and storage. Since the Realsense was up against a flat acrylic port, distortions from Snell’s law caused both the incoming images and the structured infrared light to refract, which yielded erroneous depth map information. On top of this, some of the infrared light reflected back into the lens, adding additional color artefacts in the images produced, rendering species identification impossible. Traditional stereo camera calibration procedures are also difficult to do underwater, if not impossible [50]. We concluded that until we are able to model Snell’s law distortions, an approach using proprietary stereo camera approach would be infeasible.

1.3.3 Laser Calipers

Another solution for length measurement uses laser calipers – two parallel lasers placed a known distance away from each other. When calibrated correctly, the distance between the two laser dots can be used as a reference length to measure the entire fish [41, 26]. For these measurements to be accurate, both lasers must be perfectly parallel with each other and the camera axis. Depending on manufacturing tolerances, such a requirement may mean that lasers must be carefully selected. These systems are calibrated by measuring the distance between the two laser dots at a large distance before a dive [26]. Lengths are then calculated using the known distance between two points and the projection of the fish onto the camera. While length estimation becomes simple using this method, the cost of two lasers and time for the minute readjustments waste valuable resources for researchers. The dual laser mechanism also requires that the object is larger than the distance of the two lasers, meaning fish smaller than the offset cannot be measured.

This system has been used by our collaborators at the Reef Environmental Education Foundation (REEF) to document the recovery of a population of Nassau Grouper over 7 years, in order to quantify the effectiveness of recovery efforts [26, 45]. While such a system was also demonstrated to be more effective than visual estimation, it also requires manufacturing a custom-machined aluminum mount, which is not easily accessible for most divers. Calibration of the system also involves verifying that the beams are the same distance apart up to 15m from the source, which can be challenging to do in a field setting.

1.3.4 Laser Rangefinding

Our system uses only a single laser to measure distance, which removes the need to calibrate two lasers simultaneously and keep them in parallel. This technique is similar to a light projection-based triangulation rangefinder system [38], as it uses spatial information about the laser dot to determine the depth of the subject. This method can be extremely accurate with the right combination of laser and image sensor - up to 10 micrometers [19, 12]. Such sensors have been experimented with as a low-cost solution for robot localization [37], quality assurance in manufacturing [12], and 3D scanning [2].

Single laser range finding also has precedence for use in animal size studies [28, 34, 8]. The primary benefit of this approach is that it is more inexpensive than other solutions and requires less training to operate [34]. Jaquet and Breuer et al. utilize a range finder as a separate module from a regular digital camera [28, 8]. Data from both modules must be combined and processed manually to obtain lengths [28, 34], which our platform improves upon.

Table 1.1. Comparison of different underwater length measurement techniques. The approach that FishSense Lite uses is in bold.

Technique	Est. Cost (USD)	Max. Relative Error	Ease of Use	Range
Acoustic Methods	20k [35]	1.1% - 35.2% [35]	Hard	1-16m [35]
Stereo Video	4600 [43]	2.5% [24]	Intermediate	2-10m [32]
Laser Caliper	600 [6]	12% [45]	Intermediate	2-5m [45]
Laser Rangefinding	1200	15%	Easy	2-5m

1.4 Outline

The rest of this work will be focused on the algorithms behind the length extraction for FishSense Lite, specifically how the laser parameters are determined, how these are used to calculate the coordinates of the laser dot in physical space, and how that in turn is used to get the fish length. Chapter 2 will give a brief overview of FishSense Lite as a whole. Chapter 3 will provide an in-depth description of the laser algorithms. Chapter 4 will show how we have tested our system and where sources of error come from. Finally, Chapter 5 will summarize the content of this work and future work that needs to be done on the system.

Chapter 2

System Overview

This chapter contains an overview of FishSense Lite. Our primary goals with this system are to reduce the cost to gather fish length data, and to make the practice more accessible to recreational divers. Our system relies on the diver taking images with the laser beam pointed on the fish. With just the image, as well as a priori knowledge of the laser beam parameters, we can estimate the distance between the camera and the fish, and in turn use this to get the length of the fish.

2.1 Accuracy Bounds

The primary method that our system aims to improve upon is human-made visual estimates. While humans can collectively estimate the mean length of a population quite accurately, the precision with which this is done can vary wildly between divers [25]. Within the study conducted by Harvey et. al., we see that divers can be incorrect about the length of a fish by up to 20%. As our device is meant to be a supplement for a similar experiment setup, we aim to perform as well or even better than this.

2.2 Hardware Setup

This is the physical device that divers would use to take measurements. It is intended to be composed of materials that many divers already own, which is why we use off the shelf parts.

The only new component that we include is a 3D printed mount. Figure 2.1 shows one device we have built and tested.



Figure 2.1. A fully assembled FishSense Lite hardware unit.

2.2.1 Camera

For our own testing, we use an Olympus TG6, a common camera for divers to own for underwater photography. This camera itself is waterproof, though we use another protective housing around it.

This is just one example of a camera that can be used for this system – theoretically any camera would work.

2.2.2 Wide Angle Lens

With just the waterproof housing, Snell’s law distorts images in such a way that a relationship between lengths in the image and lengths in the real world are difficult to obtain [1].

To mitigate the effects of this, we use a corrective optic, which allows us to use a simple camera model described in Section 3.1. Developed by Backscatter, this corrective optic is designed to attach to the Olympus’s underwater housing, and is typically used to widen the field of view for underwater photographers. Figure 2.2 shows a checkered pattern both with and without this corrective optic installed – note that the distortion without this corrective optic causes lines to be curved, which makes it more difficult to measure lengths. Future work is required for us to understand how to model these Snell’s law distortions better so we can ultimately remove it.




Figure 2.2. Images of a checkerboard with (left) and without (right) the corrective optic.

2.2.3 Laser and Mount

The laser pointer is another off the shelf component. Many divers own laser pointers in order to gesture to others, which further reduces the cost of the overall hardware. This laser is mounted to the camera’s housing with a 3D-printed polylactic acid (PLA) mount. Due to manufacturing defects and small perturbations, that may occur in transit, the laser moves enough to cause significant variability in the measurements. Thus a calibration algorithm is required to determine initial parameters relating to the exact geometry of the system, which is described in more detail in Section 3.4. The color of the laser also has an effect, discussed in Section 4.1.6.

2.3 Operation

This section details how the device would be used in the field.

2.3.1 Calibration Procedures

Camera Calibration

In order to be able to correlate apparent lengths to actual lengths, we use a simple camera model, explained in more detail in Section 3.1. This model is only an approximation of the camera's actual behavior, and the lens dynamics are complex and cannot be easily correlated with the model parameters. Therefore, the parameters of the model must be calibrated. This is done using Zhang's procedure [51], where many images of a checkerboard are taken, and parameters are found such that known dimensions and regularity of the checkerboard are satisfied. This procedure need only be performed once during the device's lifetime.

Currently, this procedure must be done underwater. OpenCV's `cameraCalibrate` function provides a straightforward way to obtain camera intrinsics and lens distortion parameters.

Laser Calibration

As discussed in Section 2.2.3, since our laser mount is not stable enough to guarantee that laser parameters stay consistent between dives, we require a laser calibration at the beginning of every dive. More detail about this procedure can be found in Section 3.5.

2.3.2 Diver Operation

In the field, after calibrating the laser, operating the camera is fairly straightforward. The diver will take a picture of nearby fish as normal, with the restriction that the laser dot must be trained on the fish. Note that we expect the diver to ensure that two conditions are met:

1. The head and tail points of the fish are visible in the image.
2. The laser is both present and visible on the fish.

Excerpts from the field manual are provided in Appendix C.

2.3.3 Data Offloading

After the data has been collected, the operator must offload the data to be processed externally. This software is still in development, though the process is described in the next section.

2.4 Software

Once the data has been collected, it must be processed to extract the lengths of the captured fishes. A summary of the process is as follows: locations of both the fish and the laser dot in the image are currently identified manually. We are currently developing techniques to do this automatically using machine learning. The laser's location in 3D space is determined using the location of the laser dot in the image, along with prior knowledge of the laser's location with respect to the camera. The fish segmentation mask is used to create a polygon of the fish's outline, after which PCA is performed to identify the axis of symmetry of the fish. The intersection of this line is used to find the head and tail of the fish, and the depth obtained from the laser dot is used to determine the locations of the head/tail points in 3D space, described in more detail in Section 3.3. The distance between these points is what is used to calculate the length of the fish. We currently have plans to create a web service that handles all of this processing, and intent for citizen scientists to upload their data.

A flowchart of the processing pipeline is included in Figure 2.3. Subsequent subsections will go into more detail regarding the individual stages of this pipeline.

2.4.1 Raw Image Processing

Anecdotally, we have observed that the camera performs some distortion correction internally, and as such, raw images appear to be more distorted than the JPEGs that the camera generates. In order to ensure that we have control over all aspects of the image processing pipeline, we convert the raw camera images to a PNG instead of working with JPEGs processed

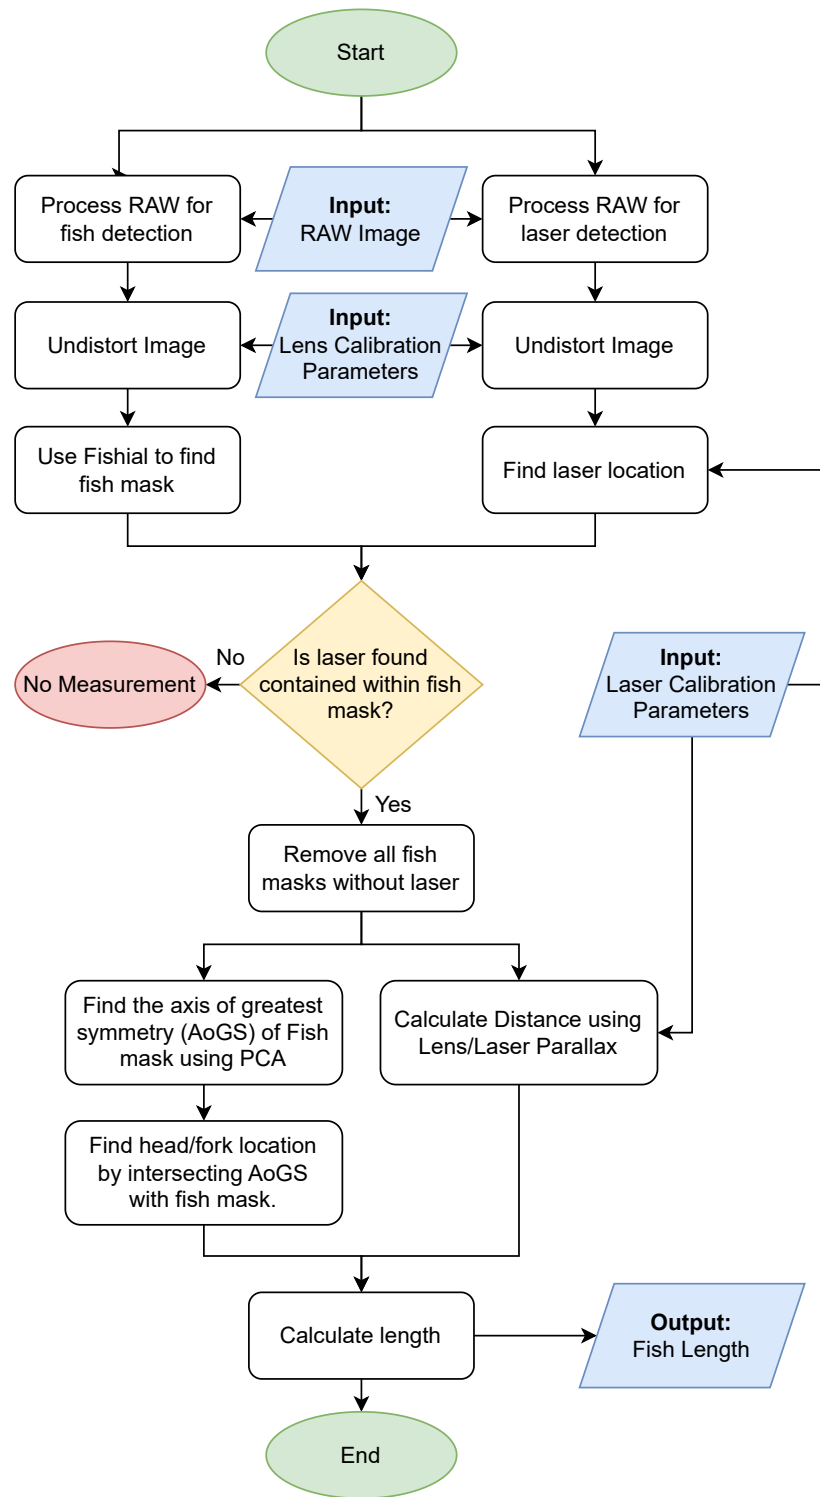


Figure 2.3. Flowchart of software processing pipeline.

by the camera.

To convert from raw images to a PNG, we must first apply a debayering algorithm. Debayering is the process of using a Bayer pattern, shown in Figure 2.4, to approximate a pixel’s “true” color. Pixels on an image sensor cannot inherently detect color, so colored filters are applied in front of digital camera sensors. The raw image consists of cells that only detect red, green, or blue light, and algorithms must be employed to approximate the actual color of each pixel.




Figure 2.4. Illustration of the Bayer pattern on a typical digital camera sensor. From Wikimedia Commons [16].

After the image is debayered, we must apply color correction to ensure that the fish being photographed are still identifiable. Typically, this requires a known “white point”, which may change in environments with different lighting.

To circumvent this issue, we use the grey world assumption [9]. This assumes that the color of the image can be adjusted by scaling each color by a uniform factor, for which we use the average of each channel, as described in the transformation below:

$$(\alpha R, \beta G, \gamma B) \rightarrow \left(\frac{\alpha R}{\frac{\alpha}{n} \sum_i R}, \frac{\beta G}{\frac{\beta}{n} \sum_i G}, \frac{\gamma B}{\frac{\gamma}{n} \sum_i B} \right)$$

A comparison with this form of color correction on a sample image is displayed in Figure 2.5.




Figure 2.5. Comparison between camera-processed JPEG image (left) and our processed PNG (right).

Note that within this work, unless stated otherwise, the JPEG images are displayed instead of the PNGs, as they are usually better looking. However, the PNGs are what we currently use for our algorithms.

2.4.2 Finding the Laser Dot

In order to identify where the laser dot is in the image we usually hand-label them. Automatic approaches are still under development. The naive approach of simply picking the brightest spot in the image does not always work, as sometimes the laser beam may attenuate so much underwater that the laser dot is no longer the brightest spot in the image, or perhaps sunlight may reflect off a bubble into the camera. Our current attempt relies on a convolutional neural network (CNN) that is trained on 20×20 tiles within the image. When searching for the laser dot within an image, the entire image is broken up into tiles, and the tile with the highest probability of being the laser dot is selected.

We can further constrain the tiles we need to feed into the network by also relying on the laser parameters to select a smaller region of the image. Since we know the laser is restricted to a single line in 3D space, it will also be restricted to a single line in the image.

2.4.3 Finding Fish Head and Tail Points

These are also usually hand-labeled, and all data presented within this work is hand-labeled. For automatic detection, we currently rely on an open source network called Fishial [20] to segment the fish. This is a MaskRCNN that both identifies the species of the fish and draws a polygon surrounding the fish's outline. Once we identify the outline, we perform Principal Component Analysis (PCA) to identify the line along which the location of outline points varies the most, which is usually along the body of the fish. The intersections of this line with the fish outline are chosen as the head and tail points.

Thus far we have described the system setup. The next chapter will describe the main contribution of this work – the algorithms underlying distance estimation that utilize the laser.

2.5 Acknowledgements

The author would like to acknowledge the work of Christopher Crutchfield, Vivaswat Suresh, Ana Perez, Hamish Grant and Kyle Tran for assisting in the development of the software processing pipeline.

This chapter contains material that is being prepared for publication. Christopher L. Crutchfield, Kyle S. Hu, Vivaswat Suresh, Ana I. Perez, Avik Ghosh, Hamish J. Grant, Kyle Tran, Samantha Prestrelski, Ronan Wallace, Nathan Hui, Jack Elstner, Dylan Heppell, Jennifer Loch, Alli Candelmo, Brice X. Semmens, Curt Schurgers, and Ryan Kastner. The thesis author was one of the primary investigators and authors of this material.

Chapter 3

Algorithm Details

The specific contribution detailed in this work are the algorithms relating to the laser pointer – both how the parameters of the laser with respect to the camera are determined, and, given those parameters are known, how the location of the intersection of the laser and fish can be determined.

3.1 Pinhole Camera Model

All of the calculations in this chapter assume that our camera follows the pinhole camera assumption. In air, this assumption is extremely common and is used in stereo cameras for object triangulation.

The pinhole camera model assumes that all light that comes into the camera passes through a single point known as the *focal point*, or *optical center*, of the camera. All incoming rays fall onto the image sensor, which is a plane perpendicular to the camera's axis, at a fixed distance from the focal point known as the *focal length*. These parameters must be calibrated for, and we follow the procedure described in Section 2.3.1.

The image projected onto the sensor is a flipped version of the real life scene, as the light rays are inverted after passing through the optical center. Instead of looking at the image on the image sensor, it is typically the convention to mirror the image sensor plane in front of the optical sensor, such that the distance from the focal point to this new plane is also the focal

length. This plane is designated as the *image plane*. This convention allows us to more easily translate from pixel coordinates to coordinates in the image plane, as we need only know the spacing between individual pixels on the sensor to map to physical units, without needing to flip the image. A pictorial representation of this is show in Figure 3.1.




Figure 3.1. Illustration of the pinhole camera model. Adapted from [18].

Underwater, it is more difficult to make the assumption that imaging systems follow the pinhole camera model. A image taken with a camera in an underwater housing is distorted, as the flat port of the camera housing causes incoming light to refract. Thus we use the corrective optic mentioned in Section 2.2.2, which allows the pinhole camera assumption to hold.

3.2 Quantities and Conventions

Our axis conventions in this chapter are as follows: the x axis points to the right in the image, the y axis points downward in the image, and the z axis points forward, away from the camera. The z axis is commonly known as the *optical axis* or *camera axis*. This system has the added bonus that the x and y axes coincide with the directions of pixel coordinates. The origin will be defined at the camera’s optical center, where all light rays intersect. This coordinate system is commonly known as the “optical frame”, and any mention of 3D coordinates in this

work will be within this coordinate system.

We will make use of the following quantities: n represents the length of the fish in pixels, w represents the “pixel pitch” (distance between two pixels on the image sensor), and f represents the focal length of the camera. A diagram of the image sensor is shown in Figure 3.2. Note that the product nw can be interpreted as the length of the fish on the image plane. Figure 3.3 shows a diagram of both the live fish and its projection onto the image plane.

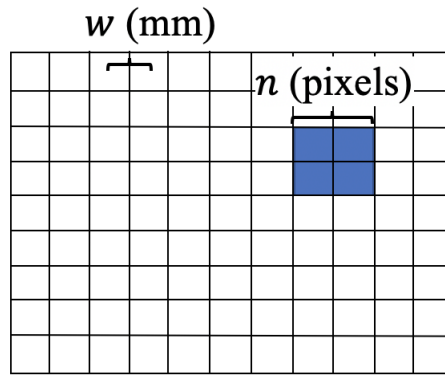


Figure 3.2. Diagram of the image plane, demonstrating pixel pitch w and apparent object length n .




Figure 3.3. Pinhole diagram showing a lionfish with it’s image projection.

If the apparent length of the fish is known, using similar triangles we are able to obtain

the true length using the following relationship:

$$L = \frac{Dnw}{f} \quad (3.1)$$

Our only unknown on the right hand side is D , the distance of the fish from the camera. The key part of our system is the mechanism with which we obtain this D .

Note that the above equation makes the following assumptions:

1. The fish is parallel with the image plane.
2. The laser dot is at the same depth as head and tail points on the fish.

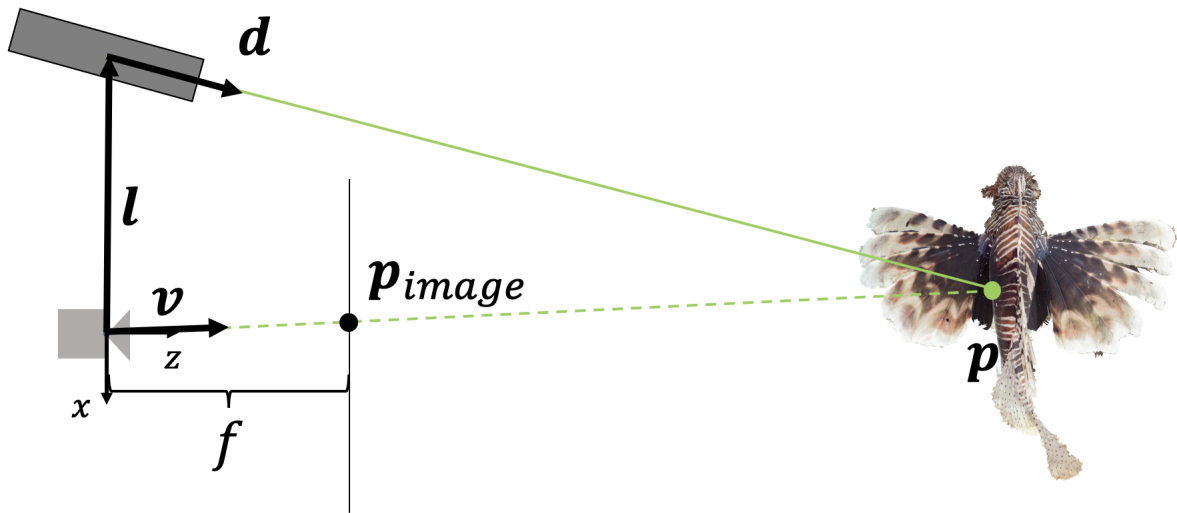
In practice, neither of these are always, if ever, true, and the implications of this are discussed more in Chapter 4. However, these tend to be reasonable assumptions for most fish.

We can define parameters \mathbf{l} and \mathbf{d} that describe the laser's position and orientation with respect to the camera, where \mathbf{l} is the 3D vector from the origin to the laser beam in the xy plane, and \mathbf{d} is a 3D vector of unit length that points in the direction of the laser beam. These are shown in Figure 3.4. Note that our restriction that \mathbf{l} lies purely in the xy plane means that the z component of \mathbf{l} is always zero.

3.3 Finding the Distance to the Fish

In order to find the distance between the camera and fish, we must assume that \mathbf{l} and \mathbf{d} are known. How these are obtained is explained in the next section.

We assume that this laser beam intersects the fish at an unknown point \mathbf{p} . As imaged by the camera, this laser dot will have known pixel coordinates \mathbf{p} . Note that the origin of these pixel coordinates is not the top left corner of the image, but the center or *principal point* of the image, which is also found during camera calibration. The point \mathbf{p} corresponds to a point $\mathbf{p}_{\text{image}}$ in the optical frame, defined as



$$\|d\| = \|v\| = 1$$

Figure 3.4. Diagram demonstrating known quantities and unknown parameters. Laser shown in solid green, imaged laser dot in dashed green.

$$\mathbf{p}_{\text{image}} = \begin{bmatrix} p_x \\ p_y \\ f \end{bmatrix}$$

We arbitrarily decide to scale this vector to be of unit length, defining

$$\mathbf{v} = \frac{\mathbf{p}_{\text{image}}}{\|\mathbf{p}_{\text{image}}\|}$$

We can identify parameters λ_1 and λ_2 such that when the laser and \mathbf{v} are scaled out, they intersect, or more formally,

$$\mathbf{l} + \lambda_1 \mathbf{d} = \lambda_2 \mathbf{v}$$

We can refactor the above relationship to be the form

$$\begin{bmatrix} \mathbf{d} & -\mathbf{v} \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} = -\mathbf{1} \quad (3.2)$$

In practice, however, this relationship does not necessarily hold, because the laser dot as observed in the image may be in a particular pixel that does not match up perfectly with where $\lambda_2 \mathbf{v}$ intersects with the image plane. In three dimensions, this often causes the rays to come very close to each other, but not touch, meaning Equation 3.2 has no solution.

We can, however, obtain λ_1 and λ_2 values that minimize the difference between the left and right hand sides – this is the least squares solution, i.e. the solution to the minimization problem

$$\operatorname{argmin}_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2,$$

where $A \in \mathbb{R}^{m \times n}$, $m > n$, and $\mathbf{x}, \mathbf{b} \in \mathbb{R}^n$. In our case, we define the following:

$$\begin{aligned} \mathbf{A} &= \begin{bmatrix} \mathbf{d} & -\mathbf{v} \end{bmatrix} \\ \mathbf{x} &= \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \\ \mathbf{b} &= -\mathbf{1} \end{aligned}$$

In general, we can solve a least squares problem with the following formula:

$$\mathbf{x} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$$

Using this, we can obtain a closed form solution for both λ_1 and λ_2 . We only require one

of them, so we use λ_2 , for which we can obtain this closed form solution:

$$\lambda_2 = \frac{-\mathbf{d}^T \mathbf{l} \mathbf{d}^T \mathbf{v} + \mathbf{v}^T \mathbf{l}}{1 - (\mathbf{d}^T \mathbf{v})^2}$$

A derivation of the above is included in Appendix A.1.

Obtaining λ_1 or λ_2 allows us to obtain \mathbf{p} , and the z-component of \mathbf{p} , along with Equation 3.1, gives us the corresponding fish length.

3.4 Finding Laser Parameters

So far, we have assumed that the parameters of the laser \mathbf{d} and \mathbf{l} are known. These cannot be measured directly for two reasons:

1. Precise measurements relative to the optical center of the camera are hard to obtain, since in reality the location of the optical center of the camera is not known.
2. Since the laser mount is not perfectly stable, and can change in between dives and over time, the parameters need to be recalculated.

Thus, we must use a calibration procedure. This involves capturing k total images, where from any particular i -th image we obtain a 3D laser point \mathbf{p}_i . Details on how this point is obtained are discussed in Section 3.5.

Here we describe two possible algorithms to leverage this information to determine \mathbf{d} and \mathbf{l} . A comparison of the two methods is presented in Section 4.1.2.

3.4.1 Method 1: Gauss Newton Optimization

We leverage the fact that the parameterized laser beam must intersect with the laser dot point to give us the following series of equations:

$$\mathbf{p}_i = \|\mathbf{p}_i - \mathbf{l}\| \mathbf{d} + \mathbf{l} \tag{3.3}$$

The above equation states that the point \mathbf{p}_i should be able to be retrieved from scaling the laser ray out by a factor determined by the difference between \mathbf{p}_i and \mathbf{l} .

We stack these points into a single vector, defining the following:

$$P = \begin{bmatrix} \mathbf{p}_1 \\ \mathbf{p}_2 \\ \vdots \\ \mathbf{p}_k \end{bmatrix}$$

We then once again define a parameter vector \mathbf{x} that contains all of our parameters:

$$\mathbf{x} = \begin{bmatrix} \mathbf{d} \\ l_x \\ l_y \end{bmatrix}$$

We also define a function in terms of our parameters for the right hand side:

$$\mathbf{g}_i(\mathbf{x}) = \|\mathbf{p}_i - \mathbf{l}\| \mathbf{d} + \mathbf{l}$$

$$G(\mathbf{x}) = \begin{bmatrix} \mathbf{g}_1(\mathbf{x}) \\ \mathbf{g}_2(\mathbf{x}) \\ \vdots \\ \mathbf{g}_k(\mathbf{x}) \end{bmatrix}$$

We then formulate this in terms of the following optimization problem:

$$\begin{aligned} & \operatorname{argmin}_{\mathbf{x}} \|\mathbf{r}(\mathbf{x})\|, \\ & \mathbf{r}(\mathbf{x}) = P - G(\mathbf{x}) \end{aligned} \tag{3.4}$$

We have a relationship that relates our known quantities \mathbf{p}_i and unknown quantities \mathbf{d}

and \mathbf{l} , though in this case the relationship is non-linear. We must therefore choose a non-linear optimization method to find the best candidates for \mathbf{d} and \mathbf{l} that satisfy this.

The method we currently choose is the Gauss-Newton method, which involves the following steps:

1. Find the Jacobian $J_{\mathbf{r}}$ of the minimizing function w.r.t to \mathbf{x} .
2. Take iterative steps of the following form:

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - (J_{\mathbf{r}}^T J_{\mathbf{r}})^{-1} J_{\mathbf{r}} \mathbf{r}(\mathbf{x}^{(k)})$$

According to Equation 3.4, only the second term depends on our parameters, so we can rework our iterative step into

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + (J_G^T J_G)^{-1} J_G \mathbf{r}(\mathbf{x}^{(k)})$$

Here we set $\mathbf{x}^{(0)}$ to a rough estimate of where the laser is with respect to the camera, taken with a ruler. For the particular test system mentioned in this work, the laser's starting point is assumed to be -4cm in the x direction and -11cm in the y direction, with the laser parallel to the camera axis.

There are many other methods of this kind that we could have used, such as the Levenberg-Marquardt algorithm, though from our experiments this calibration method has been sufficient.

The Jacobian of $G(\mathbf{x})$ is given by the following, and the full derivation is detailed in Appendix A.2:

$$J_G = \begin{bmatrix} J_{Gd}^1 & J_{G1}^1 \\ J_{Gd}^2 & J_{G1}^2 \\ \vdots & \vdots \\ J_{Gd}^k & J_{G1}^k \end{bmatrix} \in \mathbb{R}^{3k \times 5}$$

$$J_{Gd}^i = \|\mathbf{p}_i - \mathbf{l}\| I \in \mathbb{R}^{3 \times 3} \quad (3.5)$$

$$J_{G1}^i = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} \left(I_{3 \times 3} - \mathbf{d} \frac{(\mathbf{p}_i - \mathbf{l})^T}{\|\mathbf{p}_i - \mathbf{l}\|} \right) \in \mathbb{R}^{3 \times 2} \quad (3.6)$$

3.4.2 Method 2: Averaging

Once again assuming we have n images with image i corresponding with laser dot point \mathbf{p}_i , our laser trajectory \mathbf{d} can be obtained from a normalized sum over the differences between all points:

$$\mathbf{d} = \frac{\sum_{i=1}^k \sum_{j>i} \mathbf{p}_j - \mathbf{p}_i}{\|\sum_{i=1}^n \sum_{j>i} \mathbf{p}_j - \mathbf{p}_i\|}$$

Note that here we take all pairwise differences where $j > i$, as without loss of generalization, here we assume that the points are ordered in ascending distance from the optical center. Hence we only sum up differences that point away from the laser origin.

Once \mathbf{d} is obtained, we take the centroid of all of our points, which we define as μ , and set the intersection of the ray defined by μ and \mathbf{d} with the xy plane to be \mathbf{l} . Formally, this is defined as

$$\mathbf{l} = \mu - \frac{\mu_z}{d_z} \mathbf{d}$$

This method can be performed much faster than the former method, though while developing our device we have stuck to using the Gauss Newton method since the constraints and minimizing function are more explicit.

3.5 Calibration Procedure

As mentioned in Section 3.4, to calibrate the laser, we assume that we know the optical frame coordinates of each laser dot \mathbf{p}_i . While the laser parameters \mathbf{d} and \mathbf{l} are still unknown, the vector \mathbf{v} is always known, since this only depends on the location of the laser dot in the image \mathbf{p} . By intersecting \mathbf{v} with a known plane, we can precisely determine the location of each \mathbf{p}_i .

In order to create such a known plane, we require a calibration target that is both flat and possesses distinct features. If an image is taken of this calibration target, since the relationship between features on the target are known, a plane can be defined that best fits the calibration object.

The same checkerboard pattern mentioned in Section 2.3.1 is one example plane that is ideal for this. The main reasoning for this method was that there are two parts that OpenCV provides functionality for: detecting corners of specifically the checkerboard pattern, and calculating a 3D transformation given object points and corresponding points in an image. However, in the field, requiring that divers carry a large and heavy checkerboard for every dive is a large ask.

We have developed a procedure that does not utilize a checkerboard, and can instead be done with a dive slate – something that most divers will carry. An example of one of these is shown in Figure 3.5. We add pieces of duct tape in an arbitrary pattern on one side to make the slate more featureful. We also make the assumption that the slate remains close to parallel with the image plane.

We assume that we have a scanned copy of the dive slate, and hence the physical measurements of the duct tape pattern. Assuming that a correspondence can be drawn between the corners of the scan representation and the corners within the image, a transformation between the two can be found. This is known as a Perspective-N-Point (PnP) problem, and OpenCV has a `solvePnP` function that makes this relatively straightforward. The same problem exists using the checkerboard, but OpenCV's `findChessboardCorners` function abstracts this away from the



Figure 3.5. An example of a slate calibration image.

developer by taking advantage of the specific structure of the checkerboard (number of squares per side, side length of a square).

Results demonstrating this method compared to the checkerboard are shown in Section 4.2.2.

In this chapter, we have discussed the assumptions we make for this system to work, and outlined the algorithm for obtaining laser parameters, and using these to obtain the length of a fish. In the next chapter, we will discuss experiments done to test the device.

3.6 Acknowledgements

The author would like to thank Christopher Crutchfield, Nathan Hui and Nikolay Atanasov for assistance with the math in this chapter.

This chapter contains material that is being prepared for publication. Christopher L. Crutchfield, Kyle S. Hu, Vivaswat Suresh, Ana I. Perez, Avik Ghosh, Hamish J. Grant, Kyle

Tran, Samantha Prestrelski, Ronan Wallace, Nathan Hui, Jack Elstner, Dylan Heppell, Jennifer Loch, Alli Candelmo, Brice X. Semmens, Curt Schurgers, and Ryan Kastner. The thesis author was one of the primary investigators and authors of this material.

Chapter 4

Experiments

In this section we describe the testing that we have done so far to quantify how well the system works. A total of 24 pool tests and 2 field deployments were conducted in the process of developing FishSense Lite.

4.1 Pool Experiments

The vast majority of our experiments have been performed in swimming pools, due to our ability to control as many variables as possible, and for convenience. A full list of pool tests is included in Appendix B. All data in this section was taken on November 17th, 2023 and May 9th, 2024.

In order to test the system, we follow variations on the following procedure:

1. If calibration parameters do not exist or we do not trust the current parameters, calibrate the camera.
2. Calibrate the laser by taking pictures of one or both the planar objects mentioned in Section 3.5.
3. Take images of objects with a known length.
4. Analyze how the estimated lengths compare with the true length of the object.

The majority of these experiments were performed in a pool setting, as this was the most accessible approach for us.

4.1.1 Reference Objects

The experiments described in this section involve one of three measurement objects, described below.

Checkerboard

The checkerboard itself is useful for the same reason that makes it a good object to calibrate the laser with – it is relatively straightforward to identify the corners of a checkerboard in an image. In our case, we use the checkerboard as a convenient way to measure the orientation of the object when we measure lengths from it, while also using it to get many different measurements from a single image.

Box

Historically we have also used an aluminum box, with a 15cm section of tape being used as the reference length that we measure. The box is shown in Figure 4.1.




Figure 4.1. The aluminum box. The top edge of the black tape across the surface is what we typically measure.

Fake Fish

In order to fully simulate the purpose for which the device was intended, we use a dummy rainbow trout as a reference object. Each dummy is made from the same model with the same dimensions, though different techniques have been used to make them slightly negatively buoyant. Three generations of dummy fish, named Fred, George, and Ginny, have been used for these tests¹. Fred is shown in Figure 4.2. Henceforth in this section, we will refer to all three of them as just “the fake fish”.




Figure 4.2. Fred, our first fake fish. He, and all the others, are 31cm long.

4.1.2 Laser Calibration Algorithm

We test the accuracy of the calibration by comparing reference length measurements with different laser calibration parameters. The most common way in which we will show these measurements is be a plot of distance versus estimated length, as pictures of each subject were taken over a range of distances. This allows us to observe where errors occur, as will be shown in Section 4.1.4.

Figure 4.3 shows the results obtained from both calibrations – as can be seen, they are functionally identical. Future work still needs to be done to formally verify that the two methods are equivalent.

¹If you would like to purchase a fish of your own, go to <https://www.loftus.com/items/LF-0167>.




Figure 4.3. Comparison of two calibration methods, graphing estimated distance against estimated length.

4.1.3 Distance Measurements

To evaluate purely the distance estimates we can obtain from the system, we take images of objects of known length (in this case, all three of the reference objects listed in Section 4.1.1), at known distances, and comparing this known distance with the distance we get from the algorithm described in Section 3.3. These data were collected by taking pictures at fixed increments – in our case, we taped 20cm marks to a long fishing rod, and took images at each mark. There were originally 19 marks along the rod, though many of the pieces of tape came off underwater. Results are shown in Figure 4.4, which demonstrates that the actual distance and the measured distance are well correlated.

Note that this is heavily dependent on the laser calibration being accurate – the data in Figure 4.4 follow two different trends, each of which comes from a different set of laser parameters. Both of these parameters were obtained from images in the same test – one

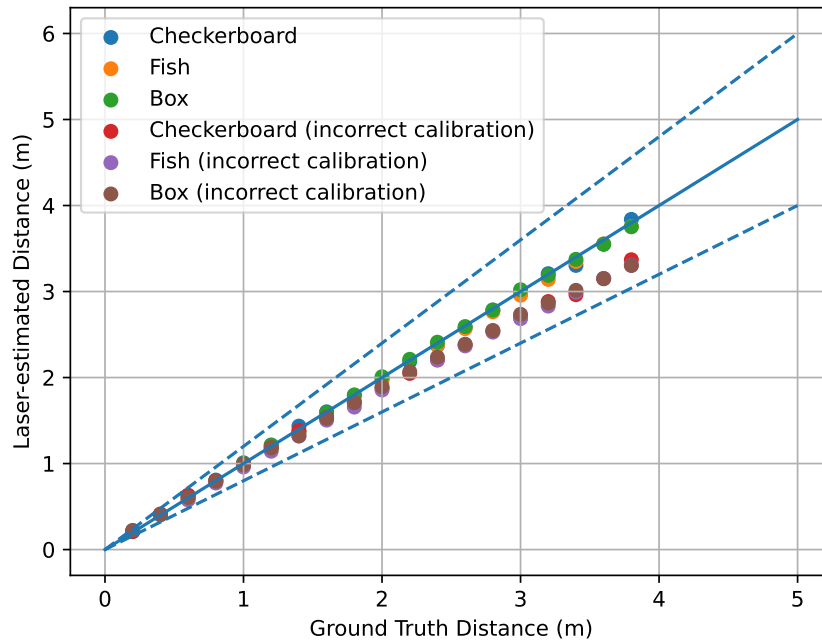


Figure 4.4. Distance estimates compared to a distance reference, measured with two different laser parameter sets. 20% error margin included.

calibration run was performed before and the other performed after the length measurements. In this particular case, the incorrect laser parameters were obtained before all measurements were taken, and the correct parameters were obtained afterward. This implies that the mount or the laser itself was moved slightly before measurements began.

In practice, there is no way that the diver would be able to tell if the mount has been moved or not, which makes verifying the measurements difficult. Section 4.1.4 will contain further discussion on how calibration parameters affect our measurements.

4.1.4 Length Measurements

Figure 4.5 demonstrates measurements of our fake fish and the box during a pool test. We can see that with varying distances away, the estimate remains relatively consistent, though there are several factors that cause error. These are listed in the subsequent subsections.

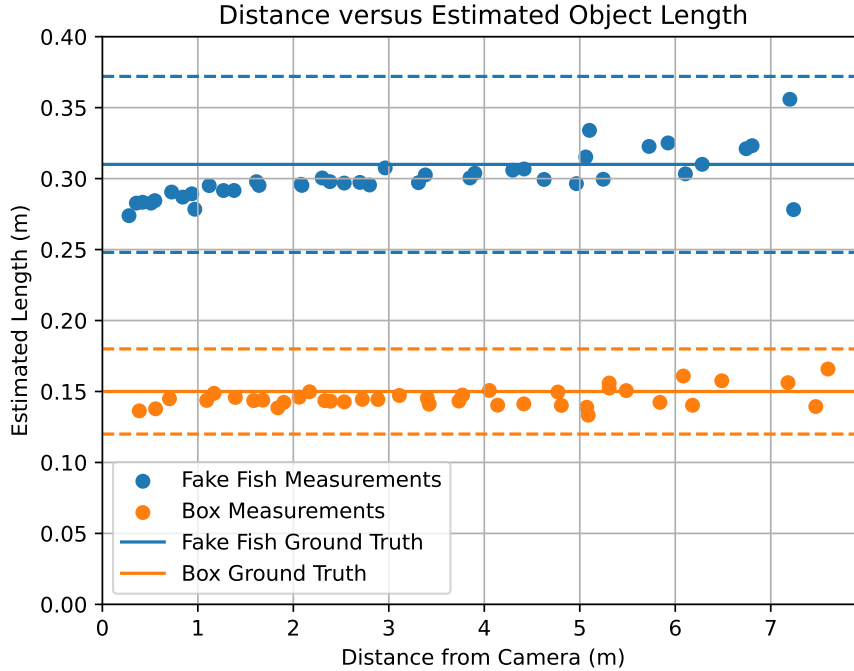


Figure 4.5. Measurements of the fish and the box, plotted over estimated distances away from the camera. 20% error margin included.

Object Thickness

If the object is not flat, as is the case with most fish, then the laser dot may be at a different distance from the camera than the head and tail points. This causes closeup measurements of the object to be shorter than in images taken further away, due to the fact that the difference in distance between the laser dot and the head/tail points makes up a smaller percentage of the total distance as the camera becomes further away. This effect can be seen in Figure 4.5, as the fish’s measurements trend upward until around 1m away. It is likely that in field use, there will not be many measurements taken that close to the fish anyway, as that may scare the fish away.

Deviations in Laser Calibration

Section 4.1.3 showed that an error in the laser calibration parameters can cause deviation in measured distance. This creates a deviation in measured length that increases the further away the object is from the camera. Figure 4.6 shows the effect that these incorrect parameters have

on object length measurements.

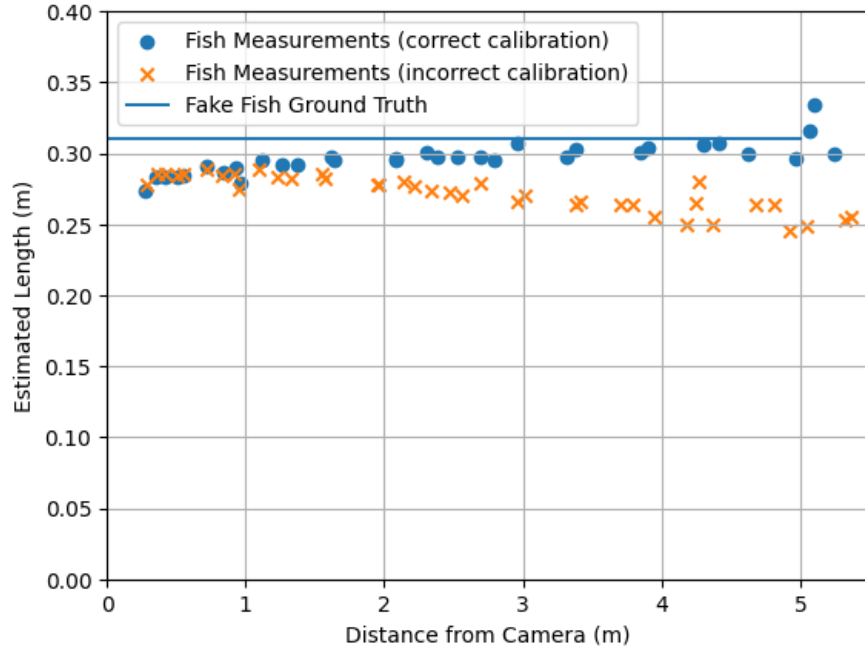


Figure 4.6. A demonstration of data where the laser parameters calculated during calibration do not match the parameters during measurement.

Notably, a change in the laser orientation \mathbf{d} results in much more error than a change in the laser location \mathbf{l} , especially as the laser dot moves further away from the camera.

Labelling Errors

All the data in this chapter were hand labeled, and this is especially noticeable at longer distances from the camera. Figure 4.5 has particularly obvious instances of large amounts of noise being introduced at longer distances. When the subject within the image is farther away, smaller deviations in the point being labeled matter much more in the final length measurement, meaning that with a hand-labeled approach, the effect of noise increases.

Deviations in Camera Parameters

An incorrect calibration of the camera's internal parameters can also cause a deviation in length. Specifically, if the focal length f is off by a certain percentage, this affects both the

process that determines distance as described in Section 3.3, as well as the conversion of distance into length using Equation 3.1. Consequently, as shown in Figure 4.7, the relationship between focal length and measurement error is complex. Concrete examples of this occurring are given in Section 4.2.1.

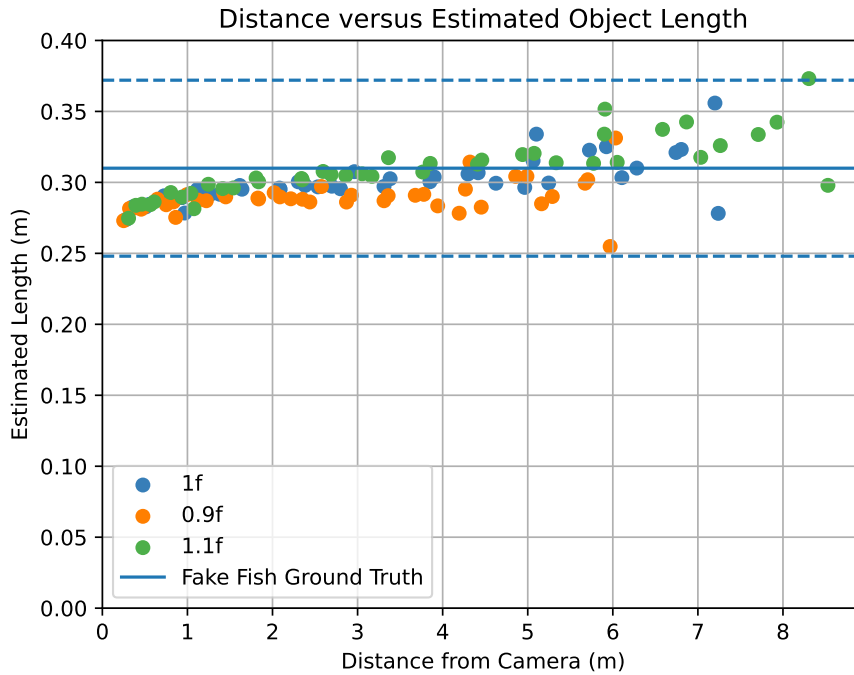


Figure 4.7. Measurements taken with the same laser parameters, with differing focal lengths.

This means that for now, camera calibration must be done carefully and verified to show ensure the obtained parameters are as accurate as possible.

Object Tilt

An important case we must account for is when the fish is not perfectly perpendicular to the camera axis. If the fish is rotated around the y axis (recall that this corresponds with the vertical axis in the image plane), the fish may look shorter within the image than it actually is, and as such the measurements of this fish would be more inaccurate. An experiment that quantifies this is described in the next section.

4.1.5 Contributions of Tilt to Error

To evaluate the system’s accuracy with respect to different tilt angles, we used the following experiment setup. We captured a total of 183 images of a checkerboard pattern with the laser dot on it, with varying tilt angles and depths. A subset of these images were used to obtain the laser position and orientation. Using these data, from each image we extracted 13 lengths, as shown in Figure 4.8. Figures 4.9 and 4.10 shows the results of this experiment.




Figure 4.8. Illustration of lengths measured on checkerboard. Each square pictured here is 41.3mm wide.

Both figures show that the relative error in measurement decreases dramatically when the angle becomes lower than 15 degrees, which supports previous studies [26]. Since our system is diver operated, we expect the vast majority of images taken will fit this criteria, as humans are quite good at observing this threshold.

Figure 4.11 shows the variances of the percentage measurement errors in all images,

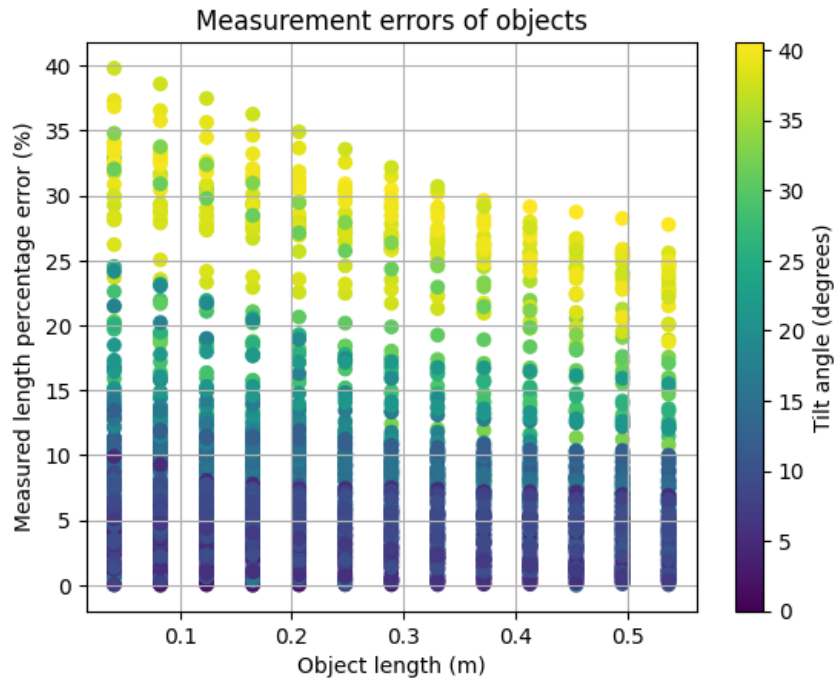


Figure 4.9. Percent error of length measurements over all distances and angles, grouped by true object length.




Figure 4.10. Percent error of length measurements over all distances and angles, plotted over the distance from the camera each image was taken.

where each data point represents a single image. We see that the variance in measurements is much higher when the tilt angle is higher and the camera is close to the image. Our current hypothesis for this is that when the calibration board is closer to the camera, the distance covered by the tilting checkerboard is a more significant percentage of the overall distance, which causes the error of the distance measured to be larger.

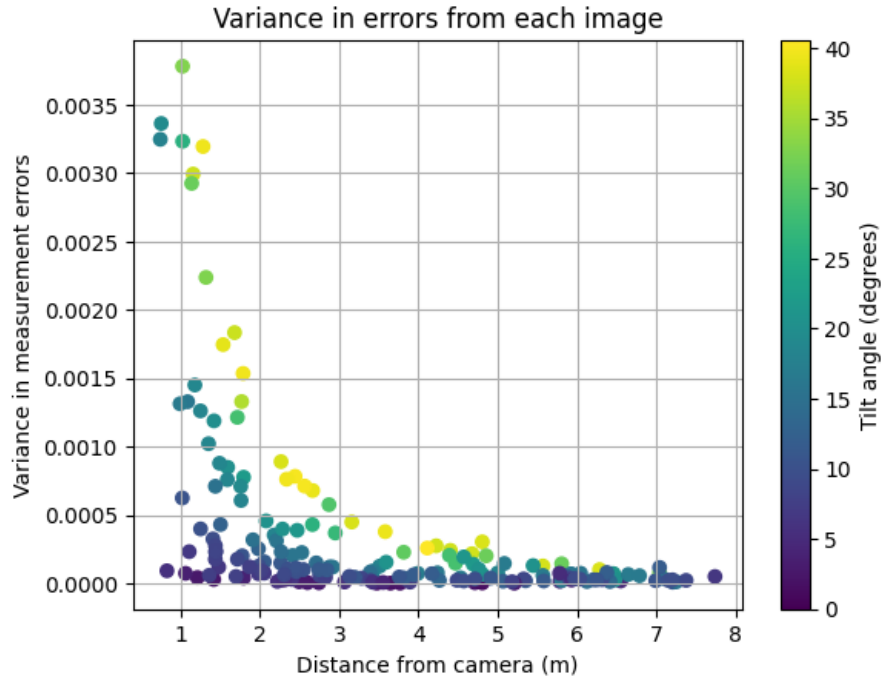


Figure 4.11. Variances of percentage length error, with each data point being a single image.

4.1.6 Laser Attenuation

The bulk of our experiments have been with two different laser pointers: the Innovative Scuba Concepts laser pointer², and the Shark Laser³. Both lasers are rated for less than 5mW of power. The former model emits light at roughly 700nm, and the latter at roughly 532nm.

Figure 4.12 shows an example of a comparative range test with the Shark laser and three different red lasers, of which the innovative Scuba laser is second from the left. Our goal with

²<https://dealer.innovativescuba.com/tc-101-aluminum-underwater-laser-pointer.html>

³https://waterprooflaser.com/contents/en-us/d22_TO-ORDER.html

this test was to qualitatively understand how attenuation of different laser colors would affect how it would be used. In pool water, we were able to observe the green laser from almost the entire width of the pool (roughly 15m), while the red laser was effectively unusable at around 5m. This did not take into account the fact that the the diver would sometimes not able to keep the camera steady enough to perceive the red laser dot at distances of roughly 4m.




Figure 4.12. Lasers circled in red from left to right: Shark laser (green), Innovation Scuba, Orca, Gold

A more detailed discussion of light attenuation in water is included in Appendix D.

4.2 Field Experiments

There are two main locations in which Fishsense Lite modules have been field tested. The first was conducted off the La Jolla coast. Water conditions in La Jolla are turbid, so this was an opportunity to test how our system would fare in low visibility conditions. The main goal of this test was to evaluate the red and green lasers, looking both at attenuation and how fish

would react.

Secondly, in August 2023, we deployed six of our FishSense Lite units (FSL-01 to FSL-06) to be tested by REEF staff in the Florida Keys. The water conditions in this region were much clearer than in the La Jolla Kelp forests, so despite the shorter attenuation of red lasers, the data from this test are predominantly generated using red lasers. These deployments allowed us to look at how the system would fare during more long term use, and in the hands of real recreational divers.

4.2.1 Salt Water Measurements

During the La Jolla deployment, we compared fake fish measurements taken in fresh water (pool water in our case) and salt water. Much like in our pool tests, the laser was first calibrated, then a fake fish was measured and lengths were calculated in post-processing. Three sets of measurements were taken: one before the ocean test, one during the ocean test, and one after the ocean test. Both the camera calibration parameters and laser calibration parameters were the same across all three measurement sets.

As seen in Figure 4.13, both the measurements before and after the ocean tests were similar, while the salt water measurements exceeded both the freshwater measurement sets. Since the cameras themselves are calibrated in fresh water, the change in refractive index in salt water is likely what causes the discrepancy in measurements. However, this is still within our accepted error margin of 20%, as we defined earlier.

We currently choose to calibrate the camera in pool water despite this error, as asking divers to calibrate in the field would be an extremely tedious process. We recognize that to minimize as much error as we can, we will need to be able to obtain the camera calibration parameters in salt water, and are currently looking into methods of doing so that do not involve using a checkerboard.

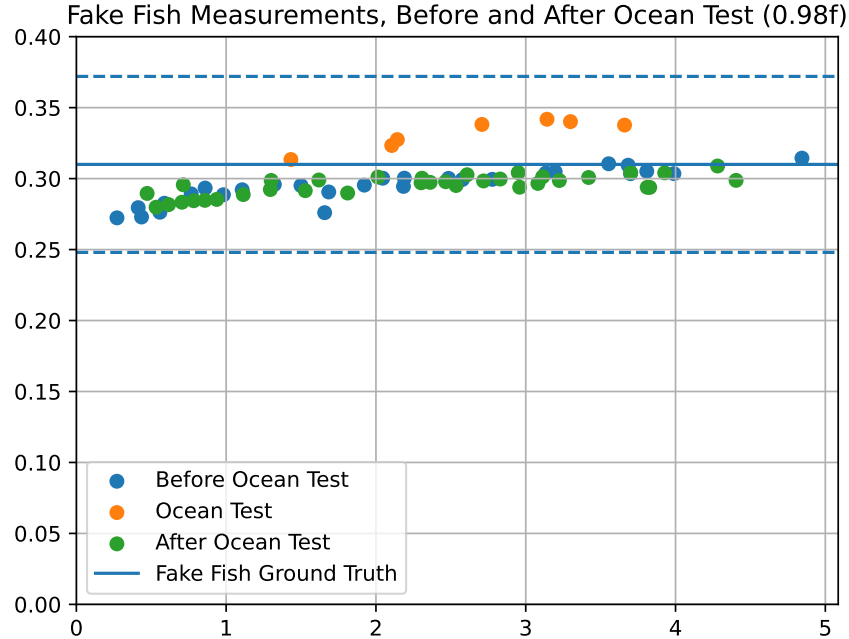


Figure 4.13. A freshwater vs saltwater comparison of measured fake fish lengths.

4.2.2 Laser Calibration With A Slate

We have been using the checkerboard method as the standard of calibration, as this can be easily automated. However, we want to assess how effective slate calibration is in theory, as it would be a huge asset to people using the device in the field. Luckily for us, the divers at REEF were able to test the slate calibration with a slate of their own construction. The methodology was as follows: each of their seven FishSense Lite units were calibrated using the same slate, and the length of the slate, show in Figure 4.14, was measured using the newly obtained laser parameters. Several measurements were taken per image to create a box and whisker plot, shown in Figure 4.15. We can see that the measurements are fairly consistent, which demonstrates that the calibration method is indeed viable.




Figure 4.14. The length measured in the slate calibration experiment.

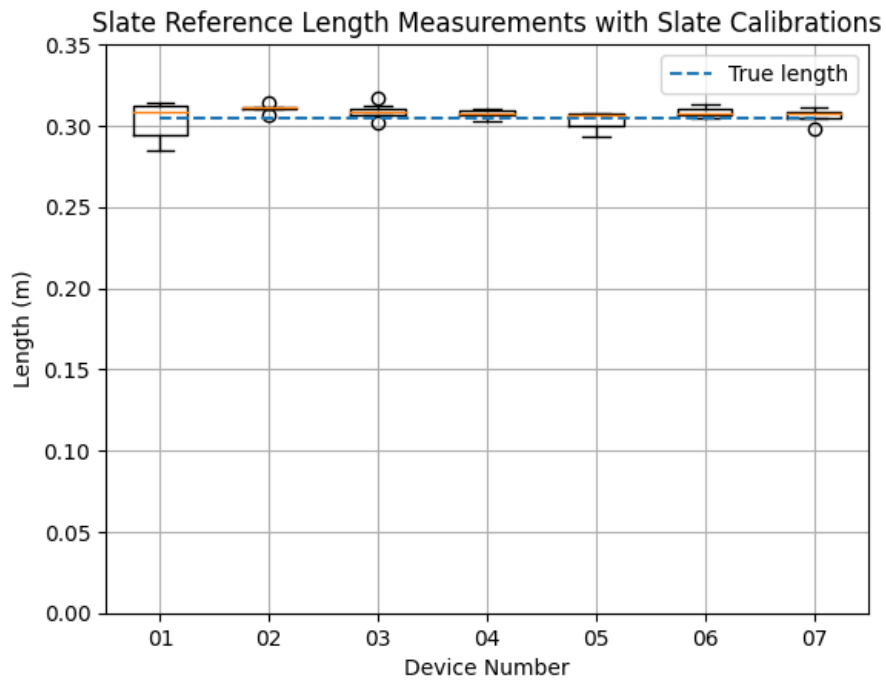


Figure 4.15. Comparisons of slate measurements across different devices.

4.2.3 Fish Behavior Study

From our testing in the La Jolla Kelp Forest, divers observed that fish tended to avoid the green laser. Our main hypothesis for why this occurs is that fish are able to see green light far better than red. In especially turbid water, added scattering meant that the beam was visible even when the beam was not pointed directly at the fish's eyes. An example of this is shown in Figure 4.16. This meant that in La Jolla, fish were especially skittish around the green laser. We realized this because after the divers made a safety stop further off the bottom, when pointing the green laser at the fish they seemed less skittish.



Figure 4.16. A male California Sheephead from the La Jolla Kelp Beds. Note the “lightsaber” effect caused by particulates in the water.

As of now, for all pool tests, we use a green laser, since we are not concerned about scaring fish away while testing. Green lasers in the field have been shown to scare fish away in turbid water, though testing is currently being conducted by REEF to assess whether a green laser is viable in the Florida Keys, where the water is relatively clear.

4.2.4 Points of Mechanical Failure

During field testing, we identified two main points where the device would tend to mechanically fail, listed below.

Laser Mount

The laser mounts have been shown to break after relatively brief periods of use (several dives). This could be due to a combination of absorption of water, degradation from exposure to ultraviolet light, or the large amount of stress being put on certain parts by the screws. Figure 4.17 shows the most common stress point which causes the mount to fail. This cracking is caused by the two screws used to keep the laser secured to the mount.




Figure 4.17. Image of one example of mount failure.

Current work is being done to create these mounts out of a different material which does not absorb water and resists corrosion.

Laser Pointer

The laser pointers themselves have also been known to fail. There are two main reasons for this:

1. The O-rings fail due to lack of proper maintenance or debris getting in the way, flooding the battery enclosure.

2. A mechanical failure from the mechanism which activates the laser. Wear from turning them on so many times causes them to chip.

Discussions with Backscatter are in the works to create a custom solution that does not need to be unscrewed, and also includes the possibility of having multiple lasers. Multiple lasers could also allow us to estimate the angle of the fish relative to the camera, which further increases the accuracy of our estimations.

4.3 Summary of Results

Thus far, we have shown that in ideal scenarios, the measurements of our device can be quite accurate. However, the “ideal” scenario seems to be rare – in practice, it seems to be quite easy for the mount to go out of calibration, whether it be due to weakness in the mount material, twisting or moving the laser out of alignment during operation, etc. There is currently no way to know for sure when this happens, which makes the results from this device more unreliable than they should be. Our next priority is to create a mount of sturdier construction such that the mount is less likely to move out of calibration during use, and that less calibration runs are required.

While this system still has many issues with durability and reliability, we have still shown that it is capable of obtaining fish lengths to within our desired error margins.

4.4 Acknowledgements

The author would like to thank Nathan Hui and Jack Elstner for testing our system off the La Jolla coast. The author would also like to thank Alli Candelmo, Jen Loch, and Dylan Heppell from REEF for testing our system extensively in the Florida Keys.

This chapter contains material that is being prepared for publication. Christopher L. Crutchfield, Kyle S. Hu, Vivaswat Suresh, Ana I. Perez, Avik Ghosh, Hamish J. Grant, Kyle Tran, Samantha Prestrelski, Ronan Wallace, Nathan Hui, Jack Elstner, Dylan Heppell, Jennifer Loch, Alli Candelmo, Brice X. Semmens, Curt Schurgers, and Ryan Kastner. The thesis author

was one of the primary investigators and authors of this material.

Chapter 5

Conclusion

5.1 Summary

In this work we have discussed the construction and inner workings of FishSense Lite. We described the main assumption – the pinhole camera assumption, which allows us to draw simple relationships between apparent and actual size. We established that distance to the fish is necessary to calculate it's actual length. We discussed both procedures to obtain the laser parameters, and to obtain the distance to the laser dot from the camera once those parameters are known. We presented results that verify the accuracy of this system is within 20% of a fish's true length.

5.2 Future Work

There is much work to be done to improve the system. We are working towards removing some of the current requirements of the subject being imaged. In order to remove the requirement that the fish be parallel to the image plane, we are looking into either using multiple lasers, or relying on a machine learning method to estimate the fish's orientation.

We also aim to build a sturdier mount so that less frequent calibration is required, and so that our device will be able to better survive the elements. Current work is being done to manufacture a mount from aluminum.

One area which we are especially interested in advancing is removing the corrective

optic, which will be possible once we are able to properly model the distortions caused by the port of the underwater housing. This will reduce the cost of our current device significantly, and will also be a significant step in allowing us to use stereo cameras with Snell's law distortions.

Lastly, we also need to refine some of the automatic detection algorithms for both the fish and the calibration object, and continue to build the infrastructure to support fully automatic data processing. To ensure that citizen scientists can easily use our platform, all processing must eventually be done automatically. One of our long term goals is to also be able to perform identification, so that scientists can use these data to draw conclusions about specific populations of a certain species or even specific individuals.

Appendix A

Supplementary Derivations

A.1 Derivation of Scale Factor

Recall from Section 3.3 that we have

$$A\mathbf{x} = -\mathbf{1},$$

where we define the following matrices:

$$A = \begin{bmatrix} d_x & -v_x \\ d_y & -v_y \\ d_z & -v_z \end{bmatrix}$$
$$\mathbf{x} = \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix}$$

By multiplying A^T on the left and right hand sides, we get

$$A^T A\mathbf{x} = -A^T \mathbf{1}.$$

We also note the following:

$$A^T A = \begin{bmatrix} \|\mathbf{d}\|_2^2 & -\mathbf{d}^T \mathbf{v} \\ -\mathbf{d}^T \mathbf{v} & \|\mathbf{v}\|_2^2 \end{bmatrix}$$

$$-A^T \mathbf{1} = \begin{bmatrix} -\mathbf{d}^T \mathbf{1} \\ -\mathbf{v}^T \mathbf{1} \end{bmatrix}$$

Since this is a system of two linear equations, we can leverage standard row reduction methods to solve for λ_2 . First we create an augmented matrix:

$$\left[\begin{array}{cc|c} \|\mathbf{d}\|_2^2 & -\mathbf{d}^T \mathbf{v} & -\mathbf{d}^T \mathbf{1} \\ -\mathbf{d}^T \mathbf{v} & \|\mathbf{v}\|_2^2 & -\mathbf{v}^T \mathbf{1} \end{array} \right]$$

Then we apply row reduction:

$$\left[\begin{array}{cc|c} \|\mathbf{d}\|_2^2 & -\mathbf{d}^T \mathbf{v} & -\mathbf{d}^T \mathbf{1} \\ 0 & 1 & \left(-\mathbf{d}^T \mathbf{1} + \frac{\mathbf{v}^T \mathbf{1} \|\mathbf{d}\|_2^2}{\mathbf{d}^T \mathbf{v}} \right) \frac{1}{\left(\frac{\|\mathbf{v}\|_2^2 \|\mathbf{d}\|_2^2}{\mathbf{d}^T \mathbf{v}} - \mathbf{d}^T \mathbf{v} \right)} \end{array} \right]$$

The above provides the solution to λ_2 :

$$\lambda_2 = \frac{-\mathbf{d}^T \mathbf{1} \mathbf{d}^T \mathbf{v} + \mathbf{v}^T \mathbf{1} \|\mathbf{d}\|_2^2}{\|\mathbf{v}\|_2^2 \|\mathbf{d}\|_2^2 - \mathbf{d}^T \mathbf{v} \mathbf{d}^T \mathbf{v}}$$

Since both \mathbf{v} and \mathbf{d} are both unit vectors, we then get

$$\lambda_2 = \frac{-\mathbf{d}^T \mathbf{1} \mathbf{d}^T \mathbf{v} + \mathbf{v}^T \mathbf{1}}{1 - (\mathbf{d}^T \mathbf{v})^2}.$$

A.2 Derivation Of Jacobian

Recall the function $G(\mathbf{x})$:

$$G(\mathbf{x}) = \begin{bmatrix} \|\mathbf{p}_1 - \mathbf{l}\| \mathbf{d} + \mathbf{l} \\ \|\mathbf{p}_2 - \mathbf{l}\| \mathbf{d} + \mathbf{l} \\ \vdots \\ \|\mathbf{p}_k - \mathbf{l}\| \mathbf{d} + \mathbf{l} \end{bmatrix},$$

where \mathbf{p}_i is the laser dot in optical coordinates from calibration image i .

We can split J_G into two block columns representing the Jacobians with respect to different vectors:

$$J_G = \begin{bmatrix} J_{G\mathbf{d}} & J_{G\mathbf{l}} \end{bmatrix}$$

Calculating $J_{G\mathbf{d}}$ is fairly trivial, as the function is linear in \mathbf{d} :

$$J_{G\mathbf{d}} = \begin{bmatrix} \|\mathbf{p}_1 - \mathbf{l}\| I_{3 \times 3} \\ \|\mathbf{p}_2 - \mathbf{l}\| I_{3 \times 3} \\ \vdots \\ \|\mathbf{p}_k - \mathbf{l}\| I_{3 \times 3} \end{bmatrix}$$

We can take the Jacobian with respect to \mathbf{l} by taking the Jacobians of each row individually.

For a particular image i we have:

$$\begin{aligned} \mathbf{g}_i(\mathbf{d}, \mathbf{l}) &= \|\mathbf{p}_i - \mathbf{l}\| \mathbf{d} + \mathbf{l} \\ \frac{\partial}{\partial \mathbf{l}} \mathbf{g}_i(\mathbf{d}, \mathbf{l}) &= -\frac{1}{\|\mathbf{p}_i - \mathbf{l}\|} \begin{bmatrix} (p_x - l_x)d_x + 1 & (p_y - l_y)d_x & (p_z - l_z)d_x \\ (p_x - l_x)d_y & (p_y - l_y)d_y + 1 & (p_z - l_z)d_y \\ p_x - l_x & p_y - l_y & p_z - l_z \end{bmatrix} \\ &= I - \frac{1}{\|\mathbf{p}_i - \mathbf{l}\|} \mathbf{d}(\mathbf{p}_i - \mathbf{l})^T \end{aligned}$$

We only want the first two columns of this, so we get

$$J_{G\mathbf{l}}^i = \begin{bmatrix} I_{2 \times 2} \\ 0 \end{bmatrix} \left(I_{3 \times 3} - \frac{1}{\|\mathbf{p}_i - \mathbf{l}\|} \mathbf{d}(\mathbf{p}_i - \mathbf{l})^T \right) \in \mathbb{R}^{3 \times 2}$$

Appendix B

Pool Test List

Table B.1 shows the full list of pool tests conducted while testing FishSense Lite.

Table B.1. List of pool tests conducted, with their purpose

Date	Purpose
2023-02-24	Lens calibration of first TG6
2023-03-03	Evaluate accuracy of OpenCV’s calibration (Olympus vs. Go Pro, housing vs. no housing)
2023-03-10	Test attenuation of red vs. green laser
2023-03-20	First attempt of laser calibration
2023-04-05	Repeat of both camera calibration and laser calibration
2023-04-12	Quantify errors of length measurement from both object tilt and thickness
2023-04-19	Collect length-referenced data for ground-truth comparison
2023-05-03	Evaluate deviation in laser calibration parameters after travel
2023-05-08	Obtain laser parameters with new mount
2023-07-13	Calibrate FSL-01 camera and laser
2023-07-29	Determine whether FSL-01 has maintained calibration after ocean test
2023-08-07	Camera calibrations for FSL-02, 03, 04, 05, 06 and 07
2023-08-14	Laser calibrations for FSL-02, 03, 04, 05, 06 and 07
2023-08-18	Recalibration of laser mounts due to faulty mechanical design. First testing of slate calibration
2023-10-28	Flat port data
2023-11-10	Trial for obtaining more fine-grain length data for different tilt angles
2023-11-17	Second trial for obtaining more fine-grain length data for different tilt angles
2024-02-12	Testing if rougher edges and corners affect the quality of slate calibration
2024-02-26	Laser calibration and length measurements with flat port
2024-03-04	Data for testing automatic slate calibration
2024-04-25	Determine accuracy of range measurements
2024-05-02	Better determine accuracy of range measurements with more data points
2024-05-09	Redo the range accuracy test, attempt new camera calibration

Appendix C

FishSense Lite User Manual

C.1 Setup

1. Take the camera underwater and burp the lens. When reattaching the lens, ensure that it is tightly attached. The lens needs to remain in the same position it originally calibrated.
2. Turn on the laser by turning the back end clockwise.
3. Take a validation photo using the method specified in Section

C.2 Use

1. When the laser is in use, aim it towards the center of the body mass of the fish. Ensure that the laser is not pointed directly toward a person.
2. In order to image fish, it is not necessary to look through the camera's viewport. If the laser is visible on the fish, it should also be visible in the camera.
3. Refer to the appropriate camera manual for photo and video collection procedures. Take photos with fish as close to flat in the image as possible; otherwise, the length measurement may be inaccurate.

C.3 Teardown

1. When data collection is complete, turn off the laser.
2. Offload data to a hard disk if necessary.

C.4 Guidelines

1. Do not use the zoom function on the camera. This changes the focal length, rendering our calibration useless.
2. Avoid touching or moving the laser directly. Do not use the laser as a handle. The laser must be rigidly attached and unmoved for calibration values to hold true.
3. The laser must be pointed directly at the fish, as close to its center of body mass as possible.
4. The camera's optical axis should establish approximately a right angle with the vertical axis of the fish.
5. Ensure the fish is located within the effective range of 0.5 to 5 meters from the camera. With 3 meters being the ideal distance.
6. The fish must be illuminated.

C.5 Camera Settings

The following settings are the ones for which the camera was tested, and to keep the data consistent, we recommend that you use them as well. Unless specified, all these settings can be found by pressing the “OK” button.

- Program auto (set the wheel to P)
- Flash off

- ISO-A 1600
- Exposure bias 0
- Burst mode
- Underwater shallow white balance
- ESP metering
- Face priority off
- Accessory off
- JPEG mode underwater
- Autofocus
- Aspect ratio 4000 x 3000
- Output format LF + raw
- Still image stabilization on

Appendix D

Light Attenuation In Water

The optical properties of water create some unique challenges in underwater optical imaging. Attenuation of light underwater is influenced by two main factors: absorption, the property of water itself to absorb light, and scattering, where water can re-emit incident light in other directions. Formally, the *attenuation coefficient* c of a medium is defined as a sum of the *absorption coefficient* a and *scattering coefficient* b , all of which are functions of the wavelength of light λ .

$$c(\lambda) = a(\lambda) + b(\lambda) \quad (\text{D.1})$$

All of the above quantities are in units of m^{-1} , and represent the fraction of incident power dissipated per meter.

Figure D.1 shows the absorption coefficients of pure water corresponding to different wavelengths of light – the least absorbed portion of the spectra falls within the visible region.

In particular, water permits green and blue light to pass up to an order of magnitude more easily than red light. Divers experience this the deeper they go into the ocean – as sunlight gets attenuated by ocean water, objects that are red begin to appear black, while blue and green objects remain the same. Underwater photographers must either bring filters or lights with them underwater to ensure that the scene they photograph is reflective of its true colors. Figure D.3 demonstrates the difference in color as sunlight is absorbed. Also note that ultraviolet and near IR frequencies are absorbed significantly more than within the visual spectrum – this is why




Figure D.1. Absorption coefficients of pure water over the electromagnetic spectrum. From Wikimedia Commons [15].




Figure D.2. Absorption coefficients of pure water in just the visible part of the spectrum. Adapted from Wikimedia Commons [14].

underwater imaging is almost never done outside the visual spectrum.

Referring back to the laser comparison in Section 4.1.6, the red laser sits with the 650-700nm wavelength range, and the green laser sits within the 510-540nm range. As can be seen from Figure D.2, the green laser is absorbed around 10 times less than the red laser.



Figure D.3. Colored rods photographed on the surface and at depth. From PBS Learning Media [33].

The coefficients in Equation D.1 also change with dissolved or particulate matter in the water, and will generally increase both a and b .

For a detailed guide on the optical properties of water and how they are studied, the Ocean Optics Web Book¹ is a great resource.

¹<https://www.oceanopticsbook.info/>

Bibliography

- [1] Amit Agrawal, Srikumar Ramalingam, Yuichi Taguchi, and Visesh Chari. A theory of multi-layer flat refractive geometry. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, jun 2012.
- [2] Mitsuru Baba, Kozo Ohtani, Makoto Imai, and Tadatoka Konishi. New laser rangefinder for three-dimensional shape measurement of specular objects. *Optical Engineering*, 40(1):53 – 60, 2001.
- [3] Bocar Sabaly Baldé, Massal Fall, Justin Kantoussan, Fambaye Ngom Sow, Malick Diouf, and Patrice Brehmer. Fish-length based indicators for improved management of the sardinella fisheries in senegal. *Regional Studies in Marine Science*, 31:100801, 2019.
- [4] E. Belcher, W. Hanot, and J. Burch. Dual-frequency identification sonar (didson). In *Proceedings of the 2002 International Symposium on Underwater Technology (Cat. No.02EX556)*, pages 187–192, 2002.
- [5] JD Bell, GJS Craik, DA Pollard, and BC Russell. Estimating length frequency distributions of large reef fish underwater. *Coral Reefs*, 4:41–44, 1985.
- [6] PATRICK BERGERON. Parallel lasers for remote measurements of morphological traits. *Journal of Wildlife Management*, 71(1):289–292, feb 2007.
- [7] Marta Bolgan, Eric Parmentier, Marta Picciulin, Louis Hadjioannou, and Lucia Di Iorio. Use of passive acoustic monitoring to fill knowledge gaps of fish global conservation status. *Aquatic Conservation: Marine and Freshwater Ecosystems*, 33(12):1580–1589, 2023.
- [8] Thomas Breuer, Martha M. Robbins, and Christophe Boesch. Using photogrammetry and color scoring to assess sexual dimorphism in wild western gorillas (*gorilla gorilla*). *American Journal of Physical Anthropology*, 134(3):369–382, 2007.
- [9] G. Buchsbaum. A spatial processor model for object colour perception. *Journal of the Franklin Institute*, 310(1):1–26, 1980.
- [10] Debby L. Burwen, Steven J. Fleischman, and James D. Miller. Accuracy and precision of salmon length estimates taken from didson sonar images. *Transactions of the American Fisheries Society*, 139(5):1306–1314, 2010.

- [11] Matthew D. Campbell, Reynaldo Patino, James Tolan, Richard Strauss, and Sandra L. Diamond. Sublethal effects of catch-and-release fishing: measuring capture stress, fish impairment, and predation risk using a condition index. *ICES Journal of Marine Science*, 67(3):513–521, 11 2009.
- [12] Federico Cavedo, Michele Norgia, Alessandro Pesatori, and Gabriel E. Solari. Steel pipe measurement system based on laser rangefinder. *IEEE Transactions on Instrumentation and Measurement*, 65(6):1472–1477, 2016.
- [13] James Churnside, Michael Jech, Eirik Tenningen, Howard McElderry, Richard O’Driscoll, Tim E Ryan, Mark R Shortis, Stephen J Smith, and WW aldo Wakefield. Fishery applications of optical technologies. 2012.
- [14] Wikimedia Commons. File:absorption coefficient of water.svg — wikimedia commons, the free media repository, 2023. [Online; accessed 26-May-2024].
- [15] Wikimedia Commons. File:absorption spectrum of liquid water.png — wikimedia commons, the free media repository, 2023. [Online; accessed 26-May-2024].
- [16] Wikimedia Commons. File:bayer pattern on sensor.svg — wikimedia commons, the free media repository, 2023. [Online; accessed 8-June-2024].
- [17] Denham Cook, Karen Middlemiss, Peter Jaksons, William Davison, and Alistair Jerrett. Validation of fish length estimations from a high frequency multi-beam sonar (aris) and its utilisation as a field-based measurement technique. *Fisheries Research*, 218:59–68, 2019.
- [18] Elijs Dima. *Multi-Camera Light Field Capture : Synchronization, Calibration, Depth Uncertainty, and System Design*. PhD thesis, 06 2018.
- [19] Mostafa Abdel-Bary Ebrahim. 3d laser scanners’ techniques overview. *Int J Sci Res*, 4(10):323–331, 2015.
- [20] Fishial. Home — fishial, 2023.
- [21] John A Fornshell and Alessandra Tesei. The development of sonar as a tool in marine biological research in the twentieth century. *International Journal of Oceanography*, 2013, 2013.
- [22] World Wildlife Fund. Sustainable seafood. [Online; accessed 27-May-2024].
- [23] Jordan S. Goetze, Todd. Bond, Dianne L. McLean, Benjamin J. Saunders, Tim J. Langlois, Steve Lindfield, Laura. A. F. Fullwood, Damon Driessen, George Shedrawi, and Euan S. Harvey. A field and video analysis guide for diver operated stereo-video. *Methods in Ecology and Evolution*, 10(7):1083–1090, 2019.
- [24] Euan Harvey, D. Fletcher, and Mark Shortis. A comparison of the precision and accuracy of estimates of reef-fish length made by divers and a stereo-video system. *Fisheries Bulletin*, 36:38–49, 01 2001.

- [25] Euan Harvey, David Fletcher, and Mark Shortis. Estimation of reef fish length by divers and by stereo-video: a first comparison of the accuracy and precision in the field on living fish under operational conditions. *Fisheries Research*, 57(3):255–265, 2002.
- [26] Scott A. Heppell, Brice X. Semmens, Stephanie K. Archer, Christy V. Pattengill-Semmens, Philippe G. Bush, Croy M. McCoy, Selina S. Heppell, and Bradley C. Johnson. Documenting recovery of a spawning aggregation through size frequency analysis from underwater laser calipers measurements. *Biological Conservation*, 155:119–127, oct 2012.
- [27] Jeffrey A. Hutchings and John D. Reynolds. Marine Fish Population Collapses: Consequences for Recovery and Extinction Risk. *BioScience*, 54(4):297–309, 04 2004.
- [28] Nathalie Jaquet. A simple photogrammetric technique to measure sperm whales at sea. *Marine Mammal Science*, 22(4):862–879, 2006.
- [29] Robyn E. Jones, Ross A. Griffin, and Richard K.F. Unsworth. Adaptive resolution imaging sonar (aris) as a tool for marine fish identification. *Fisheries Research*, 243:106092, 2021.
- [30] Brian Koepnick, Jeff Flatten, Tamir Husain, Alex Ford, Daniel-Adriano Silva, Matthew J Bick, Aaron Bauer, Gaohua Liu, Yojiro Ishida, Alexander Boykov, Estep D Roger, Susan Kleinfelder, Toke Nørgård-Solano, Linda Wei, Foldit Players, Gaetano T Montelione, Frank DiMaio, Zoran Popović, Firas Khatib, Seth Cooper, and David Baker. De novo protein design by citizen scientists. *Nature*, 570(7761):390–394, 2019.
- [31] M. C. LANGKAU, H. BALK, M. B. SCHMIDT, and J. BORCHERDING. Can acoustic shadows identify fish species? a novel application of imaging sonar data. *Fisheries Management and Ecology*, 19(4):313–322, 2012.
- [32] Delphine Mallet and Dominique Pelletier. Underwater video techniques for observing coastal marine biodiversity: A review of sixty years of publications (1952-2012). *Fisheries Research*, 154:44–62, 2014.
- [33] PBS Learning Media. Colors underwater. [Online; accessed 26-May-2024].
- [34] Graham G Monkman, Kieran Hyder, Michel J Kaiser, and Franck P Vidal. Accurate estimation of fish length in single camera photogrammetry with a fiducial marker. *ICES Journal of Marine Science*, 77(6):2245–2254, 03 2019.
- [35] Robert P. Mueller, Richard S. Brown, Haakon Hop, and Larry Moulton. Video and acoustic camera techniques for studying fish under ice: a review and comparison. *Reviews in Fish Biology and Fisheries*, 16(2):213–226, sep 2006.
- [36] Ransom A. Myers and Boris Worm. Rapid worldwide depletion of predatory fish communities. *Nature*, 423(6937):280–283, may 2003.
- [37] Hoa G Nguyen and Michael R Blackburn. A simple method for range finding via laser triangulation. *Technical Document*, 2734, 1995.

- [38] S. Parthasarathy, J. Birk, and J. Dessimoz. Laser Rangefinder For Robot Control And Inspection. In Azriel Rosenfeld, editor, *Robot Vision*, volume 0336, pages 2 – 11. International Society for Optics and Photonics, SPIE, 1982.
- [39] Patrick Richard Paxson. Systems design of underwater 3-d camera modules. mathesis, University of California San Diego, 2022.
- [40] JM Ramsay, V Watral, CB Schreck, and ML Kent. Husbandry stress exacerbates mycobacterial infections in adult zebrafish, danio rerio (hamilton). *Journal of Fish Diseases*, 32(11):931–941, 2009.
- [41] C. A. Rohner, A. J. Richardson, A. D. Marshall, S. J. Weeks, and S. J. Pierce. How large is the world’s largest fish? measuring whale sharks rhincodon typus with laser photogrammetry. *Journal of Fish Biology*, 78(1):378–385, 2011.
- [42] Rafael Masson Rosa, Daniel Caracanhas Cavallari, and Rodrigo Brincalepe Salvador. inaturalist as a tool in the study of tropical molluscs. *PLOS ONE*, 17(5):1–15, 05 2022.
- [43] SeaGIS. Pricing of popular products, 2023.
- [44] Faisal Shafait, Euan S. Harvey, Mark R. Shortis, Ajmal Mian, Mehdi Ravanbakhsh, James W. Seager, Philip F. Culverhouse, Danelle E. Cline, and Duane R. Edgington. Towards automating underwater measurement of fish length: a comparison of semi-automatic and manual stereo-video measurements. *ICES Journal of Marine Science*, 74(6):1690–1701, 02 2017.
- [45] Brian C Stock, Scott A Heppell, Lynn Waterhouse, India C Dove, Christy V Pattengill-Semmens, Croy M McCoy, Phillippe G Bush, Gina Ebanks-Petrie, and Brice X Semmens. Pulse recruitment and recovery of cayman islands nassau grouper (*epinephelus striatus*) spawning aggregations revealed by in situ length-frequency data. *ICES Journal of Marine Science*, 78(1):277–292, jan 2021.
- [46] Peter Tueller, Raghav Maddukuri, Patrick Paxson, Vivaswat Suresh, Arjun Ashok, Madison Bland, Ronan Wallace, Julia Guerrero, Brice Semmens, and Ryan Kastner. FishSense: Underwater RGBD imaging for fish measurement. In *OCEANS 2021: San Diego – Porto*. IEEE, sep 2021.
- [47] Rebecca V. Van Hoeck, Avery B. Paxton, DelWayne R. Bohnenstiehl, J. Christopher Taylor, F. Joel Fodrie, and Charles H. Peterson. Passive acoustic monitoring complements traditional methods for assessing marine habitat enhancement outcomes. *Ecosphere*, 12(11):e03840, 2021.
- [48] Jann E. Vendetti, Cedric Lee, and Pat LaFollette. Five new records of introduced terrestrial gastropods in southern california discovered by citizen science. *American Malacological Bulletin*, 36(2):232 – 247, 2018.
- [49] Shaun L Winterton. A new bee-mimicking stiletto fly (therevidae) from china discovered on inaturalist. *Zootaxa*, 4816(3):361–369, 2020.

- [50] Emily Wong, Isabella Humphrey, Scott Switzer, Christopher Crutchfield, Nathan Hui, Curt Schurgers, and Ryan Kastner. Underwater depth calibration using a commercial depth camera. In *Proceedings of the 16th International Conference on Underwater Networks & Systems*, WUWNet '22, page 15, New York, NY, USA, 12 2022. Association for Computing Machinery.
- [51] Z. Zhang. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(11):1330–1334, Nov 2000.
- [52] Shelby L. Ziegler, Rachel O. Brooks, Lyall F. Bellquist, Jennifer E. Caselle, Steven G. Morgan, Timothy J. Mulligan, Benjamin I. Ruttenberg, Brice X. Semmens, Richard M. Starr, Joe Tyburczy, Dean E. Wendt, Andre Buchheister, Jose R. Marin Jarrin, Christina Pasparakis, Salvador J. Jorgensen, Jennifer A. Chiu, Jordan Colby, Connor L. Coscino, Leon Davis, Francine de Castro, Jack T. Elstner, Christopher Honeyman, Erica T. Jarvis Mason, Erin M. Johnston, Sadie L. Small, Jay Staton, Grant T. Waltz, Bonnie Basnett, Erin V. Satterthwaite, Helen Killeen, Connor D. Dibble, and Scott L. Hamilton. Collaborative fisheries research reveals reserve size and age determine efficacy across a network of marine protected areas. *Conservation Letters*, n/a(n/a):e13000.