



UNIVERSITY OF TASMANIA  
SCHOOL OF COMPUTING AND INFORMATION SYSTEMS

---

# Benthic Habitat Mapping by Autonomous Underwater Vehicles

---

SUBMITTED IN FULFILMENT OF  
THE REQUIREMENTS FOR THE DEGREE OF

**Doctor of Philosophy**

NOVEMBER 2013

BY

**Andrew Donald DAVIE**

*B.Comp (2003, Monash)*

*GCertResCom (2013, UniSA)*

SUPERVISORS:

Prof. Paulo DE SOUZA, Prof. Neil BOSE, Dr. Robert OLLINGTON



## **Declaration of Originality**

---

This thesis contains no material which has been accepted for a degree or diploma by the University or any other institution, except by way of background information and duly acknowledged in the thesis, and to the best of my knowledge and belief no material previously published or written by another person except where due acknowledgement is made in the text of the thesis, nor does the thesis contain any material that infringes copyright.

## **Authority of Access**

---

This thesis may be made available for loan and limited copying and communication in accordance with the Copyright Act 1968. The Commonwealth Scientific and Industrial Research Organisation (CSIRO) owns intellectual property arising as a result of this work.



---

## Acknowledgements

---

I would like to thank my supervisors Professor Paulo de Souza, Professor Neil Bose, and Doctor Robert Ollington for their patient and instructive guidance during my research. It has been my honour to have them as my mentors.

This work was supported by the CSIRO's Intelligent Sensing and Systems Laboratory (ISSL), and the CSIRO Wealth from Oceans National Research Flagship. I would like to specifically acknowledge and thank Simon Allen for his continual support of the autonomous underwater vehicle (*Searise*) and my work with this machine.

I also send thanks to the many distinguished scientists who have patiently donated their time and expertise to clarify my thinking and approach. These include Professor Christoph Stiller, Professor Alberto Elfes, Professor Branco Celler and Professor Arnold Dekker. I am privileged and very grateful to have received advice from such knowledgeable people.

Dr. Jeremy Breen spent countless hours assisting me during field-deployments of *Searise* and I acknowledge his patience and persistence during those times when nothing seemed to be working, and thank him for the many discussions we had which clarified my approach and plugged the holes in my logic.

Finally, I would like to tell the world of my unbounded love for my wife and best friend Ann, and thank her for her unshakable belief and support.



## **Abstract**

This work presents a functional system for performing unsupervised classification and mapping of benthic habitats using an autonomous underwater vehicle.

Traditionally, fine-scale underwater mapping has been an expensive process, inaccurate, and is performed infrequently. This work provides contributions in three main areas: the control systems allowing the underwater vehicle to perform autonomous measurements using inexpensive sensors; implementation of a memetic algorithm for unmixing hyper-spectral signals to identify and classify habitat types; detection of new end-member types, and the production of accurate maps, while mitigating for constraints in vehicle capabilities. Meta-data is recorded and associated with detected end-member types to act as an aid in expert classification.

The production of multi-layer maps is demonstrated. These maps account for uncertainty in the vehicle's position, and variability in unmixing accuracy and confidence in data quality. This system produces end-member library spectra, maps showing end-member location and abundance, bathymetry, coverage maps, and a confidence map. An analysis of the accuracy of the algorithms using image-processing techniques showing the strength of the mapping algorithms in the presence of noise is presented.

This work has developed techniques to support fine-scale benthic classification using underwater vehicles. This capability can complement existing remote sensing techniques, allowing mapping where benthic habitat is obscured from airborne and satellite sensors because of water clarity and depth. The techniques for accurate and fine-scale measurement of benthic mixtures proposed here offer alternative tools for the monitoring and management of estuarine environments.



# Contents

**Abstract**

**List of Figures** **iv**

**List of Tables** **vi**

|          |                                      |           |
|----------|--------------------------------------|-----------|
| <b>1</b> | <b>Introduction</b>                  | <b>1</b>  |
| 1.1      | Purpose of the Study . . . . .       | 2         |
| 1.2      | Research Questions . . . . .         | 3         |
| 1.3      | Justification for Research . . . . . | 4         |
| 1.4      | The Approach Taken . . . . .         | 5         |
| 1.5      | Thesis Structure . . . . .           | 6         |
| 1.6      | Contributions . . . . .              | 7         |
| 1.7      | Summary . . . . .                    | 9         |
| <b>2</b> | <b>Autonomous Underwater Vehicle</b> | <b>11</b> |
| 2.1      | Introduction . . . . .               | 11        |
| 2.2      | System Overview . . . . .            | 11        |
| 2.2.1    | Power . . . . .                      | 13        |
| 2.2.2    | Software . . . . .                   | 14        |
| 2.2.3    | Connecting Sensors . . . . .         | 14        |
| 2.2.4    | Communication . . . . .              | 16        |
| 2.3      | Camera . . . . .                     | 16        |
| 2.4      | Navigation . . . . .                 | 18        |
| 2.5      | Control System . . . . .             | 20        |
| 2.5.1    | Control Language . . . . .           | 20        |
| 2.5.2    | The MOVETO Command . . . . .         | 20        |
| 2.5.3    | Moving to a New Location . . . . .   | 21        |
| 2.5.4    | Diving . . . . .                     | 24        |
| 2.6      | Limitations . . . . .                | 27        |
| 2.7      | Limitations . . . . .                | 29        |
| 2.7.1    | Suitability of Vehicle . . . . .     | 29        |
| 2.8      | Summary . . . . .                    | 29        |
| <b>3</b> | <b>Measuring Light</b>               | <b>31</b> |
| 3.1      | Introduction . . . . .               | 31        |

## CONTENTS

---

|          |   |           |
|----------|---|-----------|
| 3.2      | Light Propagation in Water . . . . .          | 31        |
| 3.3      | Describing the Seafloor . . . . .             | 33        |
| 3.3.1    | Community of Light . . . . .                  | 34        |
| 3.3.2    | Formal Description . . . . .                  | 35        |
| 3.4      | Optical Spectrometers . . . . .               | 36        |
| 3.4.1    | Dark Spectra . . . . .                        | 38        |
| 3.4.2    | Spectrometer Efficiency . . . . .             | 40        |
| 3.5      | Fibre-Optic Cable . . . . .                   | 41        |
| 3.6      | End-Member Spectra Collection . . . . .       | 42        |
| 3.7      | Dependency on Particular Hardware . . . . .   | 44        |
| 3.7.1    | Efficiency of the Spectrometer . . . . .      | 44        |
| 3.8      | Factors Affecting Light . . . . .             | 45        |
| 3.8.1    | Format . . . . .                              | 45        |
| 3.9      | Factors Affecting Unmixing Accuracy . . . . . | 46        |
| 3.9.1    | End-Member Library . . . . .                  | 48        |
| 3.9.2    | Calibration . . . . .                         | 49        |
| 3.9.3    | Reflective Spectra . . . . .                  | 49        |
| 3.10     | Limitations of End-Member Libraries . . . . . | 49        |
| 3.11     | Summary . . . . .                             | 50        |
| <b>4</b> | <b>Unmixing</b> . . . . .                     | <b>51</b> |
| 4.1      | Introduction . . . . .                        | 51        |
| 4.2      | End-Member Extraction . . . . .               | 52        |
| 4.2.1    | Linear Mixing Model . . . . .                 | 52        |
| 4.3      | Formal Description . . . . .                  | 53        |
| 4.3.1    | Reflectance and Reflectivity . . . . .        | 53        |
| 4.4      | Techniques . . . . .                          | 53        |
| 4.4.1    | Reconstruction of Spectra . . . . .           | 55        |
| 4.5      | Fitness Criteria . . . . .                    | 55        |
| 4.6      | Unmixing Software . . . . .                   | 56        |
| 4.6.1    | Methods . . . . .                             | 56        |
| 4.6.2    | Presentation of Results . . . . .             | 57        |
| 4.7      | The Unmixing Algorithm . . . . .              | 59        |
| 4.8      | Software Testing . . . . .                    | 62        |
| 4.9      | Autonomous End-Member Detection . . . . .     | 62        |
| 4.9.1    | Validation of Methodology . . . . .           | 64        |
| 4.9.2    | Pruning the Search . . . . .                  | 72        |
| 4.9.3    | Expert Review . . . . .                       | 73        |
| 4.9.4    | Visual Representation of Fitness . . . . .    | 73        |
| 4.10     | Analysis of Unmixing Results . . . . .        | 74        |
| 4.10.1   | Noise and Sampling . . . . .                  | 74        |

|  |            |
|--|------------|
| 4.11 Summary . . . . .                                   | 75         |
| <b>5 Mapping</b>   | <b>79</b>  |
| 5.1 Introduction . . . . .                               | 79         |
| 5.2 Habitat Maps . . . . .                               | 80         |
| 5.2.1 Format . . . . .                                   | 80         |
| 5.3 Description of Maps . . . . .                        | 81         |
| 5.3.1 Coverage Map . . . . .                             | 82         |
| 5.3.2 Bathymetry Map . . . . .                           | 82         |
| 5.3.3 Abundance Map . . . . .                            | 83         |
| 5.3.4 Confidence Map . . . . .                           | 83         |
| 5.4 Observation Distribution . . . . .                   | 85         |
| 5.4.1 Normalising the Sub-sampled Distribution . . . . . | 86         |
| 5.4.2 Distribution to Cells . . . . .                    | 88         |
| 5.5 Validating the Algorithm . . . . .                   | 91         |
| 5.5.1 Procedure . . . . .                                | 91         |
| 5.5.2 Sum-to-one Constraint . . . . .                    | 92         |
| 5.5.3 Proportion Preservation . . . . .                  | 92         |
| 5.6 Mapping Lena . . . . .                               | 93         |
| 5.6.1 Transects . . . . .                                | 94         |
| 5.6.2 Squaring the Circle . . . . .                      | 98         |
| 5.6.3 The Effect of sigma . . . . .                      | 101        |
| 5.6.4 Confidence . . . . .                               | 101        |
| 5.7 Application . . . . .                                | 101        |
| 5.8 Summary . . . . .                                    | 102        |
| <b>6 Field Testing</b>                                   | <b>103</b> |
| 6.1 Introduction . . . . .                               | 103        |
| 6.2 Cloudy Bay Lagoon . . . . .                          | 103        |
| 6.3 The Mission at Cloudy Bay Lagoon . . . . .           | 105        |
| 6.3.1 Seafloor Depth Map . . . . .                       | 112        |
| 6.3.2 Combined Abundance Map . . . . .                   | 113        |
| 6.3.3 Confidence Map . . . . .                           | 114        |
| 6.3.4 Coverage Map . . . . .                             | 115        |
| 6.3.5 End-member Abundance Layers . . . . .              | 116        |
| 6.3.6 Holotype Images and Spectra . . . . .              | 121        |
| 6.3.7 Visual Reconstruction of the Seafloor . . . . .    | 126        |
| 6.3.8 Comparative Abundance Map . . . . .                | 127        |
| 6.4 Summary . . . . .                                    | 130        |
| <b>7 Review and Conclusion</b>                           | <b>131</b> |

## CONTENTS

---

|          |   |            |
|----------|---|------------|
| 7.1      | Review of This Work . . . . .               | 131        |
| 7.2      | Revisiting the Research Questions . . . . . | 132        |
| 7.2.1    | Unmixing Spectral Data . . . . .            | 132        |
| 7.2.2    | Autonomous End-member Detection . . . . .   | 133        |
| 7.2.3    | Autonomous Mapping . . . . .                | 133        |
| 7.3      | Contribution . . . . .                      | 134        |
| 7.4      | Further Work . . . . .                      | 134        |
| 7.5      | Conclusions . . . . .                       | 135        |
| <b>A</b> | <b>GPS Algorithm</b>                        | <b>137</b> |
| A.1      | GPS Accuracy . . . . .                      | 137        |
| A.1.1    | Dilution of Precision . . . . .             | 139        |
| A.1.2    | Algorithm . . . . .                         | 139        |
| A.1.3    | DOP Bias . . . . .                          | 140        |
| A.1.4    | Proximity Bias . . . . .                    | 142        |
| A.1.5    | Weighted Average . . . . .                  | 142        |
| A.1.6    | Results . . . . .                           | 143        |
| A.1.7    | Conclusion . . . . .                        | 144        |
|          | <b>Appendices</b>                           | <b>137</b> |
| <b>B</b> | <b>Instrumentation Software</b>             | <b>145</b> |
| <b>C</b> | <b>MOVETO Flags</b>                         | <b>147</b> |
| <b>D</b> | <b>Optical Calibration</b>                  | <b>149</b> |
|          | Fibre Optic Cable Calibration . . . . .     | 150        |
|          | Fibre-Optic Cable Attenuation . . . . .     | 151        |
|          | Optical Spectrometer Calibration . . . . .  | 151        |
|          | Spectrometer Efficiency . . . . .           | 156        |
|          | <b>References</b>                           | <b>156</b> |
|          | <b>Acronyms</b>                             | <b>161</b> |

# List of Figures

|      |   |    |
|------|---|----|
| 2.1  | ISSL's AUV, <i>Searise</i> . . . . .  | 13 |
| 2.2  | Block diagram of software architecture on <i>Searise</i> . . . . .          | 15 |
| 2.3  | <i>Searise</i> operating in the Hawkesbury River (NSW). . . . .             | 17 |
| 2.4  | Three types of steering behaviour. . . . .                                  | 22 |
| 2.5  | Dive profile and vehicle behaviour. . . . .                                 | 24 |
|      |   |    |
| 3.1  | Conceptual view of habitat mapping. . . . .                                 | 32 |
| 3.2  | Miniature fibre-optic spectrometer. . . . .                                 | 36 |
| 3.3  | Internal mount holding two spectrometers. . . . .                           | 37 |
| 3.4  | Optical fibre guide. . . . .  | 37 |
| 3.5  | Instrument noise in spectrometers. . . . .                                  | 39 |
| 3.6  | Typical grating optical efficiency curve. . . . .                           | 40 |
| 3.7  | Fibre-optic cable efficiency curve. . . . .                                 | 41 |
| 3.8  | Combined fibre-optic and spectrometer efficiency curve. . . . .             | 43 |
| 3.9  | Circumsolar reference spectrum. . . . .                                     | 43 |
|      |   |    |
| 4.1  | Unmixing and evaluating result. . . . .                                     | 54 |
| 4.2  | Separate runs of the linear algebra unmixing algorithm. . . . .             | 56 |
| 4.3  | Outputs from the unmixing software. . . . .                                 | 58 |
| 4.4  | Snapshots of unmixing algorithm's solutions. . . . .                        | 60 |
| 4.5  | Ground truth and unmixed estimate. . . . .                                  | 68 |
| 4.6  | Expert analysis of end-member. . . . .                                      | 69 |
| 4.7  | Expression of solution in terms of pure end-member spectra. . . . .         | 70 |
| 4.8  | Last pass of the unmixing with nine end-members. . . . .                    | 71 |
| 4.9  | Effect of band reduction on algorithm accuracy. . . . .                     | 75 |
| 4.10 | Memetic best-fit solutions for example observation. . . . .                 | 77 |
|      |   |    |
| 5.1  | Bathymetry and coverage layers. . . . .                                     | 82 |
| 5.2  | Contour plots of abundance layers. . . . .                                  | 84 |
| 5.3  | Normal distributions with $\sigma$ of 2 and 4. . . . .                      | 86 |
| 5.4  | Normalised Gaussian before/after comparison. . . . .                        | 87 |
| 5.5  | Sample layer before and after normalisation. . . . .                        | 88 |
| 5.6  | Sum of abundance layers is one. . . . .                                     | 89 |
| 5.7  | Validation of Gaussian distributions. . . . .                               | 90 |
| 5.8  | Standard 'ground truth' image; Lena. . . . .                                | 94 |
| 5.9  | Two types of transects, showing horizontal and vertical components. . . . . | 95 |

## LIST OF FIGURES

---

|      |   |     |
|------|---|-----|
| 5.10 | Increase in transect ordinality improves coverage and quality. . . . .  | 96  |
| 5.11 | Biasing the relative contributions of observations. . . . .   | 97  |
| 5.12 | Variation in $\sigma$ affects coverage, but not significantly the sharpness<br>of reconstructed images. . . . . | 99  |
| 5.13 | Low-density sampling on order-5 transect. . . . .   | 100 |
| 6.1  | Photo of Cloudy Bay Lagoon, Bruny Island, Tasmania. . . . .   | 104 |
| 6.2  | Satellite image of Cloudy Bay Lagoon, Bruny Island, Tasmania . .  | 104 |
| 6.3  | Location of Cloudy Bay Lagoon . . . . .   | 105 |
| 6.4  | Samples of unmix results for two observations. . . . .  | 107 |
| 6.5  | Cloudy Bay Lagoon habitat map. . . . .  | 111 |
| 6.6  | Cloudy Bay Lagoon seafloor depth map. . . . .   | 112 |
| 6.7  | Cloudy Bay Lagoon combined abundance map. . . . .   | 113 |
| 6.8  | Cloudy Bay Lagoon confidence map. . . . .   | 114 |
| 6.9  | Cloudy Bay Lagoon coverage map and path overlay. . . . .  | 115 |
| 6.10 | Abundance and distribution layers for end-members NEW1 and NEW2.  | 117 |
| 6.11 | Abundance and distribution layers for end-members NEW3 and NEW4.  | 118 |
| 6.12 | Abundance and distribution layers for end-members NEW5 and NEW6.  | 119 |
| 6.13 | Abundance and distribution layers for end-members NEW7 and NEW8.  | 120 |
| 6.14 | Images and holotype spectra for detected end-members. . . . .   | 122 |
| 6.15 | Images and holotype spectra for detected end-members. . . . .   | 123 |
| 6.16 | Images and holotype spectra for detected end-members. . . . .   | 124 |
| 6.17 | Images and holotype spectra for detected end-members. . . . .   | 125 |
| 6.18 | Visual reconstruction of the seafloor . . . . .   | 126 |
| 6.19 | Reconstructions using different end-members. . . . .  | 128 |
| 6.20 | Comparative abundance map. . . . .  | 129 |
| A.1  | Converting from HDOP or distance to fractional bias value. . . . .  | 139 |
| A.2  | GPS positions and calculated position. . . . .  | 141 |

# List of Tables

|     |   |     |
|-----|---|-----|
| 2.1 | Overview of AUV <i>Searise</i> 's hardware and firmware. . . . .  | 12  |
| 2.2 | AUV <i>Searise</i> 's sensors and instrumentation. . . . .  | 18  |
| 2.3 | Parameters of the MOVETO command. . . . .   | 21  |
| 2.4 | Segment lengths of yo-yo dive. . . . .  | 25  |
| 3.1 | Optical Spectrometer Specifications. . . . .  | 38  |
| 3.2 | Characteristics of each spectrometer. . . . .   | 40  |
| 3.3 | Fibre-Optic Cable Specifications. . . . .   | 42  |
| 3.4 | Wavelength Ranges. . . . .  | 45  |
| 4.1 | Fitness for 12 iterations of 50 observations. . . . .   | 78  |
| 6.1 | Mission details from Cloudy Bay Lagoon. . . . .   | 106 |
| 6.2 | Coordinates of <i>Searise</i> transects for Cloudy Bay Lagoon mission.<br>Coordinates are UTM Zone 55G Easting and Northing in metres.<br>The transects were repeated four times to improve coverage. . . . . | 108 |
| 6.3 | Data collected during Cloudy Bay Lagoon mission. . . . .  | 109 |
| 6.4 | End-members detected, listed in order of abundance. . . . .   | 110 |
| 6.5 | Interpretation of end-member mix. . . . .   | 110 |
| A.1 | Statistics from GPS averaging algorithm. . . . .  | 143 |
| C.1 | MOVETO flag bit allocation. . . . .   | 148 |
| D.1 | Spectrometer Optical Efficiency. . . . .  | 155 |



*"I do not know what I may appear to the world; but to myself I seem to have been only like a boy playing on the seashore, and diverting myself in now and then finding of a smoother pebble or a prettier shell than ordinary, whilst the great ocean of truth lay all undiscovered before me."*

Sir Isaac Newton (1642-1727)

*"If people never did silly things, nothing intelligent would ever get done."*

Ludwig Wittgenstein (1889-1951)



## Introduction

Australia faces increasing environmental strain on its coastal reserves, rivers and estuaries, particularly around urban regions. Australia's marine reserves are currently over 3.1 million km<sup>2</sup> in area, and are managed primarily for biodiversity conservation. Vast marine reserves have also been proposed for the Australian Antarctic Territory, placing Australia at the forefront of protection of the marine environment. However, as discussed by Pala (2013) in his recent paper in *Science*, reserves of this magnitude create significant challenges and responsibilities.

Environmental assessments are costly and time consuming: the Australian Government has budgeted \$4.7 billion for environmental protection in 2013-2014 (Catanzariti, 2013). Increases in shipping, port and industrial activities, urban development, recreational use, exploration for oil and gas and, recently, a rapid expansion in aquaculture are all impacting our waterways and pose major challenges to effective environmental protection. These factors, and the need for better impact assessment, mean that the costs associated with marine monitoring will no doubt increase over the years to come.

Government and industry decision-makers are expected to provide immediate and appropriate responses to natural disasters and industrial accidents. Such responses include accurate assessments of the extent and impact of these events on the environment, and effective mitigation of damage.

Remote monitoring of the marine environment has typically been performed through the analysis of aerial and satellite photographs (e.g., Aswani and Lauer (2006), Cho (2007), Dierssen et al. (2003), Kutser et al. (2006a)), and some research into acoustic classification has been performed (Lundblad et al., 2006), but these methods are unable to provide accurate species-level classifications

## CHAPTER 1. INTRODUCTION

---

beyond a few metres depth. Such fine-scale classifications require labour-intensive methods such as divers conducting surveys, or manual interpretation of underwater video and images. These techniques are time-consuming, expensive, and generally undertaken infrequently. As assessments happen infrequently, and the methods for monitoring marine environments have limitations, it can be difficult to gain a reliable view of change and impacts.

Other monitoring methods include the use of manned research vessels such as Commonwealth Scientific and Industrial Research Organisation (CSIRO)'s *Southern Surveyor*, and unmanned autonomous underwater vehicles (AUVs) carrying cameras and other sensors such as side-scanning sonar. The *Slocum Glider* is one example of this type of vehicle, but it is limited to use in deeper ocean waters.

Seagrasses, in particular, are an important part of Australia's coastal ecosystem (Kirkman, 1997). They promote biological productivity and biodiversity, sequester large amounts of carbon (Fourqurean et al., 2012), and provide food and shelter for many marine organisms. Recently there has been a significant decline in seagrass meadows around the world, thought to be caused by anthropogenic factors (Orth et al., 2006). Seagrasses contribute to water quality, reduce erosion, provide food for marine organisms, and provide habitat for many important marine species. They are a significant part of the food chain. Seagrasses are delicate and respond rapidly to extremes of pollution, salinity, turbidity, and nutrients. Consequently, change in the distribution of seagrasses provides a good indicator of the health of a marine ecosystem.

Human impact and ecological stresses can cause the loss of many of these important communities, with consequent impact upon fisheries and the environment, at a significant financial cost to the country. Automated and regular monitoring and classification of these important benthic habitats would greatly assist in assessing the effects of human activity on coastal and estuarine environments.

### 1.1 Purpose of the Study

---

The purpose of this study is to develop a system to automate the classification of these benthic habitats, producing maps of the seafloor. The use of an AUV in this research allows mapping of the benthic environment at high temporal (e.g., daily basis) and spatial (e.g., 2 metre<sup>1</sup>) resolution at relatively low cost, complementing

---

<sup>1</sup>Systems described in this work are not restricted to a specific resolution, but limitations in the hardware used restrict positional accuracy to approximately this scale.

## 1.2. RESEARCH QUESTIONS

---

traditional habitat monitoring techniques.

Prior research, as presented and discussed in this thesis, has demonstrated that spectral analysis of hyper-spectral images is a feasible technique for characterising underwater habitat. Spectral matching classification of coral reef systems has demonstrated that the accuracy of these techniques is comparable to manual classifications. The spatial resolution of aerial imagery is a limiting factor in the sensitivity of mapping using these source data. The lack of spectral detail caused by attenuation of light by the water body, and varying attenuation as a result of the varying depth of the benthos being classified and turbidity of the water leads to a limit in the accuracy and classification capability of these methods.

## 1.2 Research Questions

---

The automation of benthic habitat mapping poses a number of challenges, particularly due to the attenuation of light passing through the water body. The use of an AUV also has associated issues related to measurement and processing of data and control of the vehicle. Additionally, the production of useful habitat maps from these data presents a number of difficulties. These areas of interest lead to the three primary research questions addressed in this work, listed below.

The key questions relate to the suitability of hyperspectral data for accurate identification of benthic habitat, the feasibility of automatic detection of spectrally distinct items which can be used as end-members for classification of mixed spectra, and the practicality of generating classification maps using data from optical spectrometers on-board an AUV.

### Unmixing Spectral Data

Is unmixing underwater hyper-spectral data a suitable technique for the classification and mapping of the seafloor?

- Is the use of a pre-defined spectral end-member<sup>2</sup> library a suitable technique for achieving accurate unmixing?
- Which unmixing techniques are most suitable for unmixing underwater hyper-spectral data?

---

<sup>2</sup>The definition of an 'end-member' is given in Section 3.3

## CHAPTER 1. INTRODUCTION

---

- Can underwater spectra be analysed to produce accurate estimates of the proportions of end-members?
- How can these factors be taken into account when calculating accuracy of unmixing?
- Can accuracy of unmixing be maintained in the presence of measurement noise and of end-members not in the library?

### Autonomous End-member Detection

Can an automated system detect new types of end-members and incorporate these into the unmixing and mapping process?

- Can an end-member library be constructed using automatically detected spectra that is sufficient for accurate unmixing?
- Can these automatically added end-members be converted into representations of pure end-member spectra?

### Autonomous Mapping

Is it possible to produce useful multi-layer maps representing the combined distributions of end-member types?

- Can the error in estimated vehicle position be accounted for when producing these maps?
- Is it feasible to run the unmixing and map creation software in real-time on an AUV, and allow completely autonomous underwater mapping?

## 1.3 Justification for Research

---

- Monitoring of the underwater environment is expected to be an important tool in management of marine reserves, and assist in understanding the effects of anthropogenic factors upon marine environments.
- Airborne and satellite-based classification of land vegetation has been achieved through the characteristics of green vegetation - low red-light reflectivity and high near-infrared (NIR) reflectivity. These characteristics

led to the development of vegetation indices for the classification of hyper-spectral images from these sources. However, the strong absorption of NIR by water, associated with the scattering of light by the water body, hampers the successful use of such indices in classification of submerged aquatic vegetation (SAV). Cho et al. (2008) has shown that such indices are “not a good indicator for SAV that is adapted to live underwater or other aquatic plants that are submerged during flooding even at shallow waters (0.3 m).”

- Aerial and satellite monitoring has limited discriminatory ability in mapping fine-scale substrate mixes. AUV-based classification and mapping can complement existing techniques and provide higher resolution data.
- Spectral variability including intra-species, morphological, and seasonal, creates a challenge for spectral-based classification systems.
- Automation of classification and mapping has the potential to reduce the cost and improve the accuracy of mapping underwater habitat.

## 1.4 The Approach Taken

---

### Overview

The fields of robotics, spectroscopy and mapping are robust and extensive literature is available related to benthic mapping, some dating back decades.

The method developed in this contribution involves:

- The use of small or portable optical spectrometers on board an AUV
- Recording of down-welling spectral data from light entering the water, up-welling spectral data reflected from the seafloor
- Calculation of reflectivity spectra from the up and down spectrometer data
- Unmixing of that spectral data
- Autonomous detection of new composite end-member types
- Identification of the types and proportions of benthos and habitat types
- Production of maps of the seafloor, including:
  - A layered abundance map of the seafloor
  - A coverage map that can be used to guide future mapping missions with the AUV

## CHAPTER 1. INTRODUCTION

---

- A confidence map showing how accurate these maps are likely to be
- A bathymetry map

### Control of an Autonomous Underwater Vehicle

The AUV used in this research does not have instrumentation that allows accurate localisation of underwater observations. The contribution of this work is a control behaviour for the AUV that mitigates inaccuracy of location estimates.

### Spectral Unmixing and Classification

The field of hyper-spectral classification is a mature one, and many algorithms and methods for unmixing spectral data have been proposed. The contribution of this work is to use a memetic algorithm containing a genetic algorithm (GA) with recombination/crossover normalisation and a simulated annealing (SA) algorithm. The system autonomously detects “sufficiently different” composite end-members from multiple observations and builds an end-member library which is used for unmixing the observation data.

### Layered Abundance and Coverage Maps

Unavoidable errors in vehicle position when surveying are considered when producing coverage maps representing the quality of data and identifying areas suitable for more extensive coverage. These coverage maps provide a confidence value for quality assessment of the layered classification maps.

## 1.5 Thesis Structure

---

**Chapter 2** introduces the autonomous underwater vehicle (*Searise*) used for this research. The chapter describes the vehicle’s hardware and software, communications and navigation. The command software enabling mapping transects is explained, and an algorithm for improving global positioning system (GPS) accuracy through biased averaging of signals is presented in Appendix A.

**Chapter 3** covers the use of reflected light as a mechanism to identify marine habitat. A formal description of the spectral characteristics of a benthic community is given. This is important to aid understanding of the linear mixing model (LMM), an assumption underlying much of the work related to spectral unmixing. The main instruments used for gathering data to be classified are two

optical spectrometers. The capability of these instruments, their mounting on *Searise*, and information about software and data formats is given. The chapter addresses issues related to the attenuation of light by water and instrumentation.

**Chapter 4** introduces the LMM and unmixing techniques. A memetic algorithm (a hybrid GA and SA algorithm) with a constraint preserving crossover method is explained. Results from investigation into the effect of band reduction upon unmixing accuracy (particularly in the presence of noise exhibiting Gaussian distribution) are given. The use of an end-member spectral library to facilitate unmixing is explained. This chapter also describes the process of detection of least-fit unmixing results, and the consequent addition of new composite end-members to the end-member spectral library. The final process of assisted human verification of a small number of end-member spectra leading to a re-interpretation of a classification map is explained.

**Chapter 5** describes the output data from the mapping process. The use of uncertainty in the position of the vehicle collecting data to distribute unmixed proportions over a layered map is validated through testing and examples. A confirmation of the effectiveness of this method and evaluation of accuracy of mapping is presented, using a standard computer graphics test image as a proxy for the seafloor. The final products of the mapping software are detailed, including a coverage map that can be used for optimising vehicle movement when performing subsequent measurements.

**Chapter 6** describes mapping missions performed at Cloudy Bay Lagoon (Bruny Island, south-east Tasmania). The results of processing collected spectral data using the unmixing and mapping algorithms are shown in the form of detailed maps and seafloor reconstructions.

**Chapter 7** begins with a review and discussion of the thesis, and offers conclusions based on the results of the work performed as a part of this research. Recommendations for future work are offered as a result of the experience gathered during this research.

## 1.6 Contributions

---

### Scientific Contributions

- **Unsupervised spectral unmixing:**

## CHAPTER 1. INTRODUCTION

---

- The development, implementation and validation of a new algorithm to perform spectral unmixing.
  - A recombination/crossover method with gene balancing to preserve unity and non-negativity constraints.
  - Autonomous composite end-member<sup>3</sup> detection.
  - Contextual information, including images, associated with composite end-members to assist with expert reclassification.
- **Layered Maps:**
    - The production of layered abundance maps indicating the location and proportion of end-member types using the techniques described in this thesis.
    - The production of gradient maps to indicate the coverage and confidence of classifications.
    - Reconstruction of the true-colour seafloor maps from distributed hyper-spectral observations.
    - Production of a mixture map representing the location of most abundant representation of each detected end-member type.

### Engineering Contributions

- **autonomous underwater vehicle (AUV):**
  - The use of an AUV carrying optical spectrometers close to the substrate to perform high-resolution in-situ habitat mapping.
  - A dynamic navigation system to improve accuracy of vehicle movement.
  - A weighted-averaging algorithm to improve accuracy of position estimates using multiple GPS sensors.
  - An effective profile dive control system that dynamically follows seafloor contours and divides legs into subcomponent dives.
  - Developing drivers for the AUV control systems.

---

<sup>3</sup>A composite end-member refers to a spectrum added to the end-member library (see Section 3.9.1) by the automatic end-member detection algorithm (see Section 4.9). The spectrum is most likely to be a mix of the spectra of more than one type of object, rather than a pure spectrum of a single type of object. Composite end-members and their decomposition into pure end-member spectra are discussed in Section 4.9.3.

### 1.7 Summary

---

This work details a viable method for autonomous benthic habitat mapping. The automation of the classification and mapping of marine habitat using an AUV presents a number of challenges spanning several different fields of expertise, including robotics, computational science, optics, and cartography.

This work explores the use of an AUV carrying miniature optical spectrometers to perform classification and mapping. Spectral unmixing and identification of benthic habitat end-members is presented, and the production of multi-layered benthic maps representing the distribution of end-member types is shown.

This approach has the potential to offer improved accuracy and resolution for benthic mapping, and reduce the cost of such mapping efforts.



*“The machine has no feelings, it feels no fear and no hope ... it operates according to the pure logic of probability. For this reason I assert that the robot perceives more accurately than man.”*

Max Frisch (1911-1991)

*“It gives me a headache just trying to think down to your level.”*

Marvin - The Hitch Hikers Guide to the Galaxy

# 2

## Autonomous Underwater Vehicle

### 2.1 Introduction

---

This chapter introduces the autonomous underwater vehicle (*Searise*) and discusses its use in this research. Details of the machine’s communication, navigation and sensor capabilities are discussed, and the limitations of the vehicle are also presented. A description of the software control systems and operation of the vehicle when performing transects is given.

### 2.2 System Overview

---

*Searise* (Figure 2.1) is an experimental autonomous submersible robot. It is a variant of *Starbug*, originally developed at CSIRO’s Autonomous Systems Laboratory, based at the Queensland Centre for Advanced Technologies (CSIRO’s ASLab) under the direction of Dr. Matthew Dunbabin (Dunbabin et al., 2004, 2005). *Searise* is 20% longer and has additional hardware and software capabilities. It was built by Andrew Davie, Dr. Klaas Hartmann and John McCulloch at CSIRO’s Intelligent Sensing and Systems Laboratory (ISSL) in 2008.

*Searise* was used to collect and process data for this research.

*Searise* is relatively small (see Table 2.1), and is easily transported by a single person using a small trolley. It is light enough to deploy and retrieve without using additional lifting equipment. *Searise* is capable of operating with all motors active for approximately 12 hours on a single charge of its lithium batteries.

## CHAPTER 2. AUTONOMOUS UNDERWATER VEHICLE

---

**Table 2.1:** Overview of AUV *Searise*'s hardware and firmware.

|                         |   |
|-------------------------|---|
| Length                  | 1.2 m   |
| Weight                  | 40 kg without external sensors.   |
| Buoyancy Control        | manual <sup>a</sup>   |
| Maximum Speed           | 1 knot  |
| Maximum Tested Depth    | 180 m   |
| Power                   | 4 x 25.9 V Lithium Batteries<br>@12.6 Ah each (total: 50.4 Ah)                                |
| Hotel Load <sup>b</sup> | ~ 0.75 A  |
| Mission duration        | > 12 h  |
| Motherboard             | Nano-ITX N700   |
| CPU                     | VIA C7 @ 1.5 GHz  |
| Operating System        | Ubuntu 10.04 (Linux)  |
| Storage                 | 1 GiB RAM <sup>c</sup><br>4 GiB Compact Flash card for OS<br>64 GiB USB memory stick for data |
| Propulsion              | 3 × flat thrusters (CSIRO-made)<br>2 × Seabotix SBT150 thrusters                              |

---

<sup>a</sup>*Searise* must be 'weighted' by adding or removing flotation before deployment to give the vehicle slightly positive buoyancy. It is required that the vehicle will surface on total loss of power. However, this can be difficult when operating in environments such as the Derwent Estuary, which has varying levels of salinity.

<sup>b</sup>The hotel load is the power usage of the vehicle when it is turned on but not using its motors or sensors.

<sup>c</sup>The International Electrotechnical Commission (IEC) international standard usage for binary multiples assigns the prefix 'Gi' for Gibibyte ( $2^{30}$  bytes). The prefix 'GB' is used to refer to multiples of  $10^9$  bytes.



**Figure 2.1:** ISSL’s AUV, *Searise*. The vehicle is approximately 1.2 m long and 40 kg weight. The decorative stripes are copied from a Tasmanian Tiger (*Thylacinus cynocephalus*) pelt.

### 2.2.1 Power

*Searise* is powered by four 25.9 V Li-Ion Polymer (LiPo) batteries (batteryspace.com part number PL-25.9V12.6Ah), each with capacity of 12.6 Ah. Each battery consists of seven cells, and one smart charging/discharging circuit that protects the cells from excessive overcharge and discharge by short-circuiting. The four batteries are connected in parallel, but they are charged individually with these connections disabled. As LiPo batteries do not out-gas when charging, the batteries are safely charged in-situ without opening the vehicle.

*Searise*’s internal battery charging electronics were developed by John McCulloch at CSIRO. The 25.9 V power supply is converted to 12 V and 5 V buses that are available to internal devices, and to external devices through the generic SubConn serial ports.

*Searise*’s power consumption when idle is under 20 W. Power consumption with sensors and thrusters operating is approximately 250 W. With 1305 Wh from the batteries, this allows the vehicle to be operated over a working day without recharging.

### 2.2.2 Software

The majority of *Searise* software is written in C, with the exception of ‘drivers’ that communicate with sensors and external devices that are written in Python, and the optical spectrometers’ driver that is written in Java.

### 2.2.3 Connecting Sensors

*Searise*’s instrument suite is reconfigurable, allowing the vehicle to carry many different types of sensors. Some of the sensors carried (internally and externally) are listed in Table 2.2. These sensors are discussed in Section 2.3 and Section 2.4.

#### RS-232 Serial Connectors

*Searise* can connect to and carry any external sensor small enough to be strapped to the vehicle and capable of communicating via the RS-232 standard. Four SubConn Micro (MCIL5F) 5-pin communications ports are available for external sensors using standard RS-232 serial communications. These ports all supply power (5 V and 12 V), and a RS-232 RX/TX serial channel for high-speed (up to 57,600 baud) serial communications.

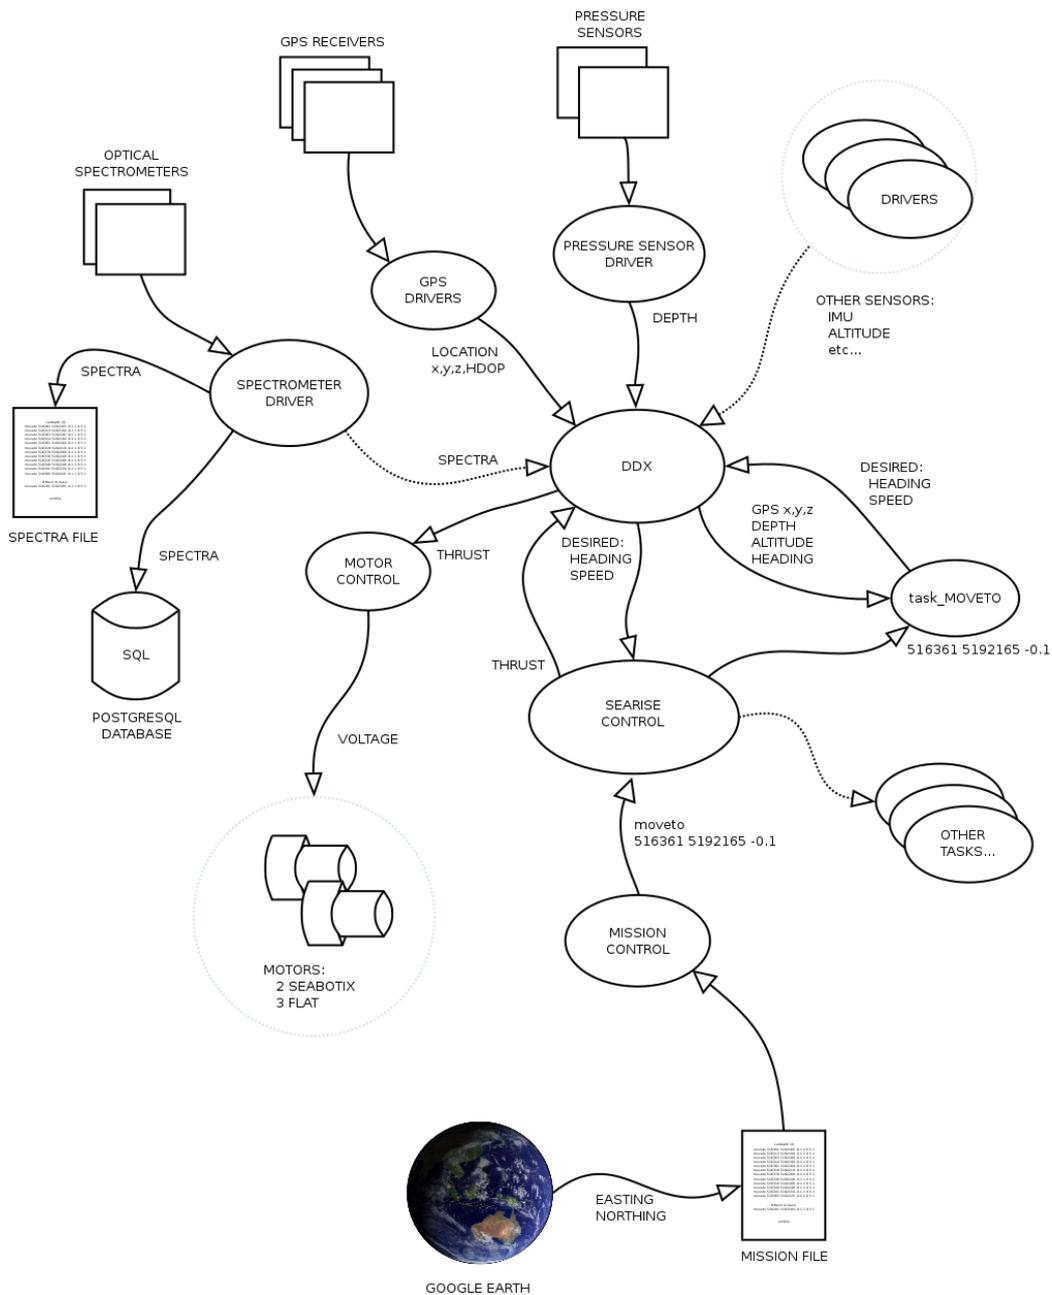
Internally, *Searise*’s devices are connected to the computer through USB. All incoming RS-232 serial ports are immediately converted to USB and connected to an internal USB hub.

*Searise* incorporates an automatic configuration script to detect which sensors are connected to the ports, and their communication speed, so the installation of sensors is essentially ‘plug-and-play’.

#### USB 2.0 Connector

One external USB 2.0 connection to an internal USB hub is also available through a SubConn Micro 6-pin port (MCIL6M). Although this has worked well, the SubConn connector’s impedance is beyond the specifications of the USB standard and so this connection may be unreliable.

## 2.2. SYSTEM OVERVIEW



**Figure 2.2:** Block diagram of software architecture on *Searise*. Coordinates from Google Earth are entered into a text file containing commands and parameters for the mission. This file is parsed by the mission control system, which passes individual commands (such as a *MOVETO*) to the *Searise* control software. Each command is handled by bespoke code. In the case of *MOVETO*, the angle to face and speed to move is calculated using the current location retrieved from the DDX store. Driver software talks to various sensors and retrieves the sensor values (such as GPS location, depth, etc.) and places these in the DDX store. The *MOVETO* code places desired heading and speed in the DDX store. The *Searise* code uses these to calculate thruster demand, and writes these to the store. The motor control code reads demand from the store, calculates voltage and applies these voltages to the individual motors.

### 2.2.4 Communication

There are two primary modes of communication with the vehicle: via Ethernet and Wi-Fi. A third hybrid method using a Wi-Fi buoy connected via Ethernet is available if required.

#### Ethernet

In the laboratory, the vehicle can be connected to the CSIRO intranet via Ethernet through a dedicated SubConn port. This allows for the vehicle to be stationed in a laboratory and diagnosis and modifications to the software and data to be made from a workstation from anywhere with an Internet connection.

#### Wi-Fi

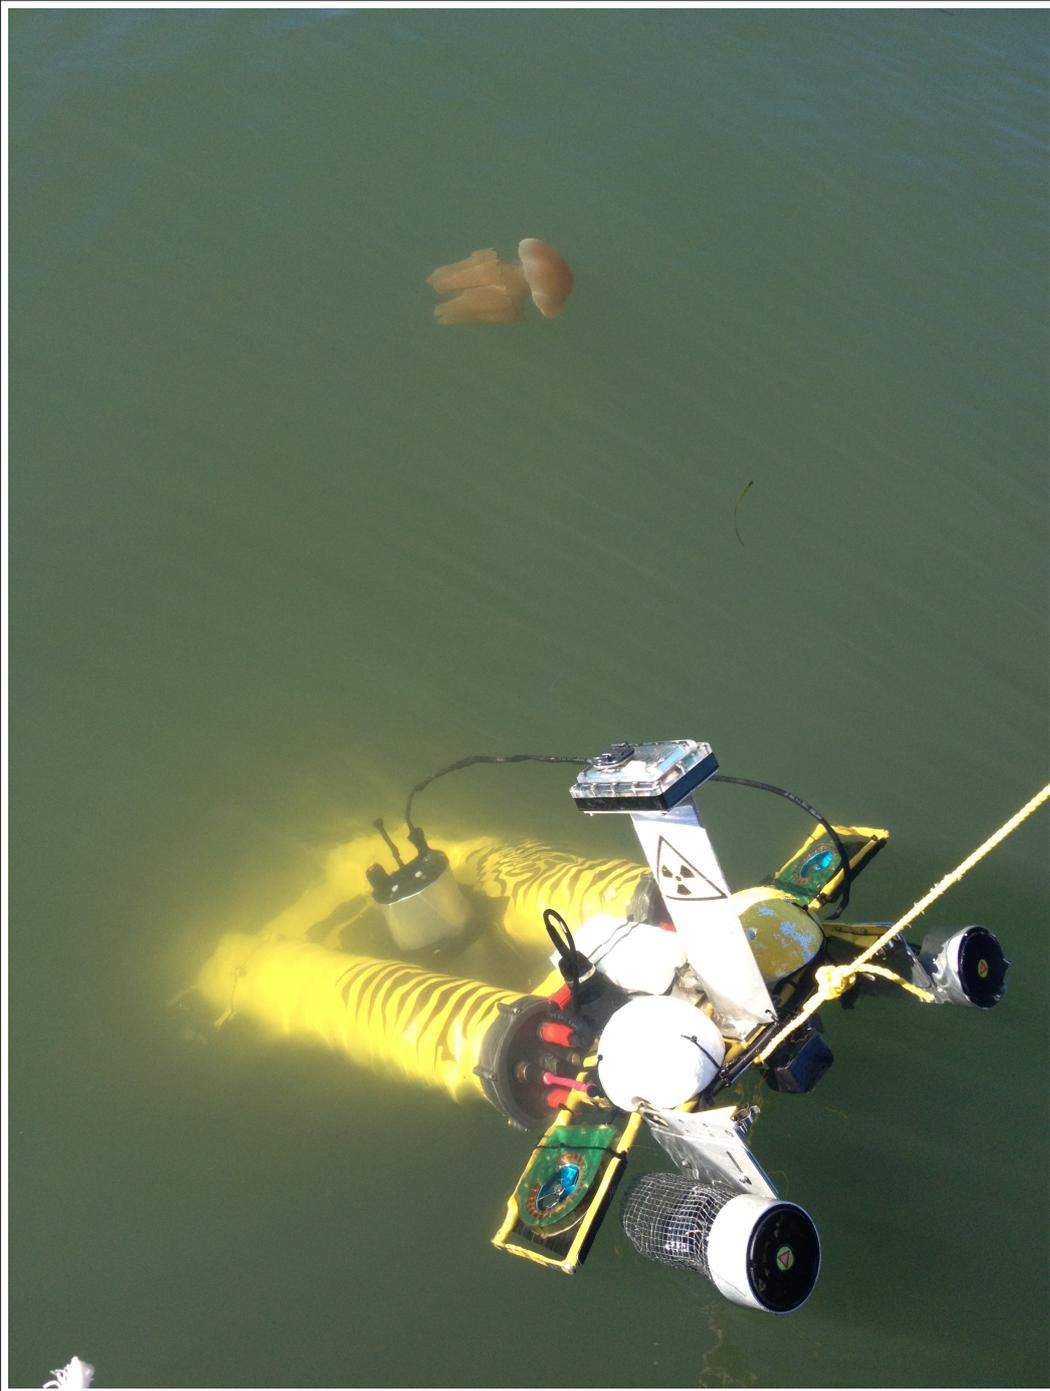
*Searise* has a generic off-the-shelf USB Wi-Fi device installed in the top of its tail housing. The vehicle establishes an ad-hoc wireless network when it boots. Most communication in the field between operator and vehicle occurs over this ad-hoc network. This is possible only when the tail is above water, and the vehicle is within range (~50 m) of the operator. *Searise*'s tail electronics were designed by Tim Warren.

To enable communication with *Searise* when it is submerged, a hybrid combination of Ethernet and Wi-Fi was developed by placing a network hub inside a buoy, and connecting this to *Searise*. The buoy with network hub floats on the surface and allows communication with *Searise* when it is submerged. John McCulloch developed this buoy at ISSL. Future use of this system could include the installation of global positioning system (GPS) receivers in the buoy to allow localisation of the submersible while it is underwater, as GPS does not work underwater.

## 2.3 Camera

---

*Searise* carries a downward-facing low-light Logitech QuickCam Pro 9000 USB camera slightly aft of the spectrometer mounts. This camera captures images at 5-second intervals (set by the driver software) to minimise storage requirements for long missions. The system camera driver software has a 4-sample image buffer, so with a frame-rate of 0.2 Hz, the sample time is offset by 20 s to match the images



**Figure 2.3:** *Searise* operating in the Hawkesbury River (NSW) near Dangar Island on 11th June 2013. The jellyfish in the image is *Catostylus mosaicus*, colloquially known as a Jelly Blubber. *Searise* is carrying an X-Ray Fluorescence spectrometer to make in-situ measurements of chemicals on the seafloor, hence the radiation symbol on the tail. The unretouched photo shows the very poor visibility (less than 2 m) of the water in this region, which has consequences for remote sensing using aerial and satellite instruments. The addition of a light to *Searise* and measuring spectra close to the substrate allows it to map habitat inaccessible to these other methods.

## CHAPTER 2. AUTONOMOUS UNDERWATER VEHICLE

---

**Table 2.2:** AUV *Searise*'s sensors and instrumentation.

| Device               | Type   | Model                                 |
|----------------------|--------|---------------------------------------|
| GPS                  | RS-232 | UsGlobalSat EM-408                    |
|                      | RS-232 | San Jose Navigation FV-M8             |
|                      | RS-232 | Locosys LS20031                       |
| Depth Sensor         | RS-232 | SSI Technologies P51-500-A-A-I36-5V-R |
|                      | RS-232 | SSI Technologies P51-100-A-A-I36-5V-R |
| Altimeter            | RS-232 | Tritech PA-500                        |
|                      | RS-232 | CruzPro ATT120BT                      |
| Compass              | RS-232 | Ocean Server OS5000                   |
| IMU                  | RS-232 | MicroStrain 3DM-GX1                   |
| Camera               | USB    | Logitech QuickCam Pro 9000            |
| Optical Spectrometer | USB    | 2 × Ocean Optics USB2000+XR-1         |
| XRF Spectrometer     | USB    | Amptek Mini-X, Amptek X-123           |

with the spectral data. This was validated by imaging a clock and checking the difference between the image file time-stamp and the captured time on the image.

One solution to remove this lag would be to call the capture subroutine four times each time a sample is required, with the latter being the correct sample, but this adds to the computational burden on the vehicle's computer and is not required; instead the offset of 20 s (4 samples) is applied when matching images to other data.

## 2.4 Navigation

---

*Searise* determines its location when surfaced from three GPS receivers in the tail assembly. The GPS units are accurate to approximately  $\pm 5$  m. Multiple GPS units are carried, allowing the accuracy to be improved to approximately  $\pm 2$  m (see Appendix A).

The GPS signal is only available when the tail is above water, so the position of *Searise* must be estimated by using the compass, IMU, thrust applied to motors, and previous known position when the vehicle is underwater. Additionally, GPS position estimates have a high degree of uncertainty when *Searise* first surfaces and the GPS units first acquire a position.

The use of multiple units from different manufacturers, with different hardware, has improved this situation. The positions from these GPS units are combined using a biased averaging method (see Section A). This averaging improves accuracy to approximately  $\pm 2$  m, and increases the speed and accuracy of position fixes when the vehicle surfaces.

### Dead Reckoning

The control software on-board the vehicle estimates the vehicle's position dynamically by using heading, thrust, and coefficient of drag of the vehicle to calculate cumulative change in position when it is underwater. This is known as dead reckoning. Additional factors influence the actual position of the vehicle, such as water currents, and wind. The consequence is that error and uncertainty in positional estimations increases the longer the vehicle is submerged. To minimise this, *Searise* regularly surfaces and re-acquires a GPS fix. This repeated diving/surfacing behaviour is called yo-yoing.

### Altitude

*Searise* carries two altimeters (a Tritech PA-500 echo-sounder with 0.3 m - 30 m range and 1 mm resolution, and a CruzPro ATT120BT echo-sounder with 1.2 m - 300 m range) that give accurate readings for the altitude of the vehicle above the seafloor. *Searise* attempts to maintain level attitude (roll and pitch zero) when performing spectral measurements.

### Depth

The vehicle's windows and seals have been pressure-tested to 180 m depth. Two pressure sensors with different range capability measure the depth, providing higher resolution at lower depths, but still retaining capability to operate at greater depths. From 0 to 40 m depth, the high resolution instrument's reading is used. From 40 m to 50 m depth, a weighted average of the low-resolution and high-resolution instruments is used. At depths greater than 50 m, the low resolution instrument's reading is used.

The sum of the altitude and the depth gives the bottom depth (see Figure 2.5).

### Heading

*Searise* does not have a rudder. Directional control is achieved by applying differential thrust to the rear motors. Although the motors can run in reverse, this capability is not currently used. The differential thrust steering method can turn *Searise* in a relatively tight circle (approximately 1 m diameter).

## 2.5 Control System

---

The software for command and control of the vehicle has been developed in-house at Commonwealth Scientific and Industrial Research Organisation (CSIRO), and varies slightly between the CSIRO's ASLab and the ISSL vehicles. The author has written many of the drivers for the sensors and vehicle control systems.

An open-source software system, Spring/DDX (Corke et al., 2004), is common to both vehicles, and was developed at CSIRO's ASLab. Dynamic data exchange (DDX) provides a centralised communications mechanism to allow various control systems to share common sensor data and communicate with each other.

### 2.5.1 Control Language

The control language for *Searise* consists of a simple command-based system implemented as a text-file (the 'mission file') with human-readable command strings. The commands can trigger complex behaviour in the vehicle, and are executed sequentially. There are no looping or repetition capabilities, but single commands can control the vehicle for many kilometres of mapping.

In particular, the MOVETO command (explained below) developed by the author provides most of the capability required for movement of the vehicle for data collection. The MOVETO command has been designed to provide a comprehensive set of behaviours, including a profile diving capability, and moving to geographical coordinates using varied steering algorithms.

### 2.5.2 The MOVETO Command

All vehicle movement for mapping manoeuvres is controlled by a single command, 'MOVETO'. The format of the command is:

```
MOVETO x y depth velocity flags radius segmentCount
```

## 2.5. CONTROL SYSTEM

---

where the parameters are described in Table 2.3. Although simple, the command can control reasonably sophisticated behaviour in the vehicle.

**Table 2.3:** Parameters of the MOVETO command.

---

---

|                           |   |
|---------------------------|---|
| <code>x</code>            | Easting <sup>1</sup> coordinate of end point (m)              |
| <code>y</code>            | Northing coordinate of end point (m)                          |
| <code>depth</code>        | Maximum depth in metres (segment C-D in profile)              |
| <code>velocity</code>     | Speed of vehicle (m/s) <sup>2</sup>                           |
| <code>flags</code>        | Behaviour control flags. See Table C.1                        |
| <code>radius</code>       | Closeness to end point for termination of command, in metres. |
| <code>segmentCount</code> | Varied use, depending on <i>flags</i> value.                  |

---

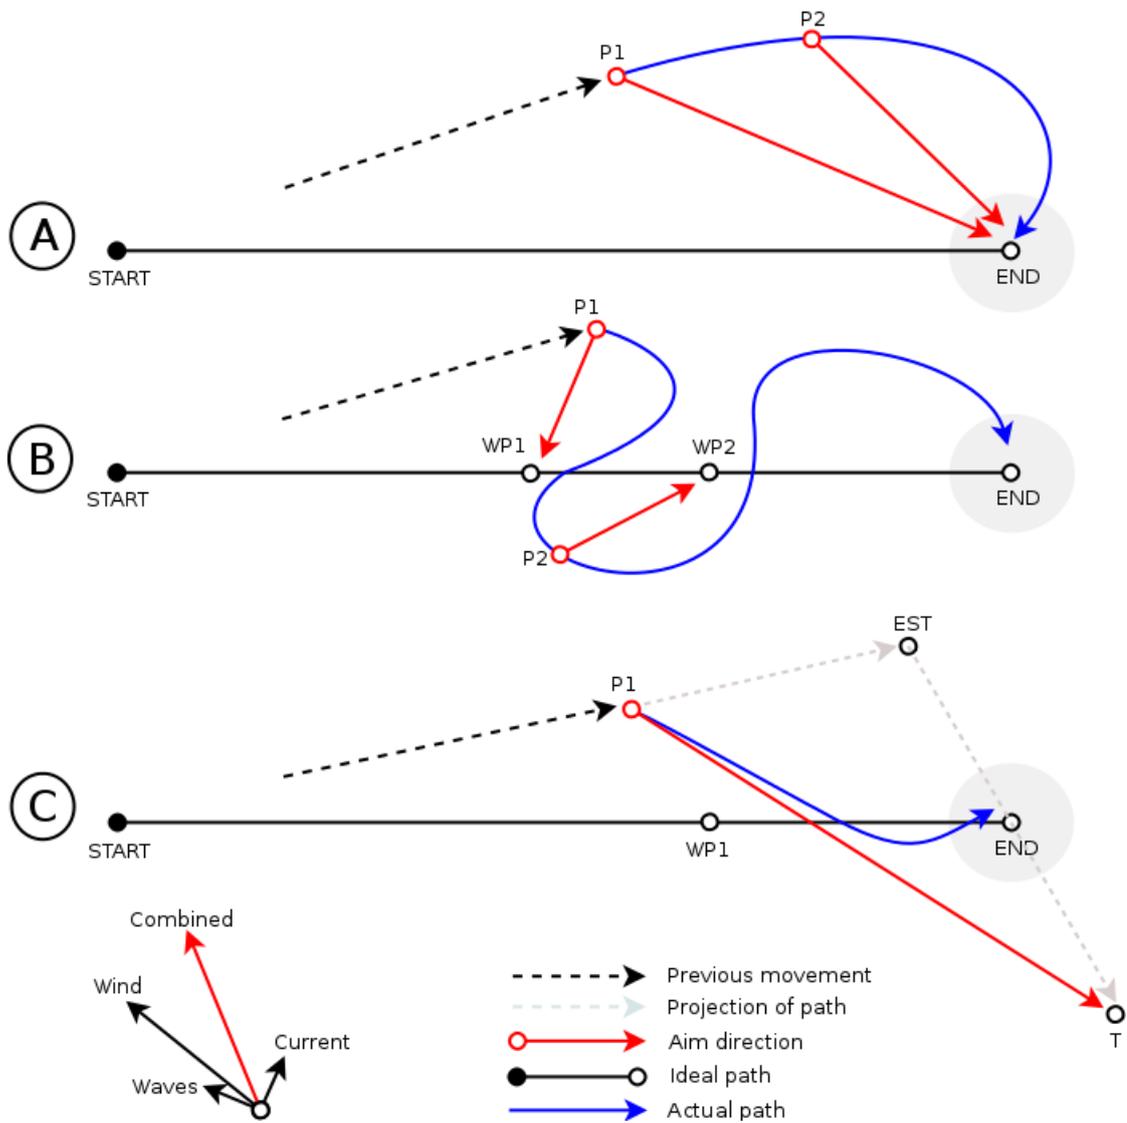
The MOVETO command performs these basic tasks:

- Moves the vehicle to a new position, specified as an easting and northing coordinate
- Manages segmentation of a long path into a series of smaller dives
- Controls the depth so the vehicle follows a dive profile
- Dynamically corrects for offset in position caused by wind, waves and current

### 2.5.3 Moving to a New Location

To get to a destination by pointing to a target and applying thrust is an unreliable method to transect on a line between two points. Any perturbation of the path (for example, wind or water current affecting the vehicle) will have a cumulative effect, displacing the vehicle from the ideal path. In this situation, as the vehicle approaches the target point this displacement can accumulate such that the path travelled between the start and the destination is curved, and the vehicle approaches the target point asymptotically.

Figure 2.4 shows typical vehicle behaviour using three different types of steering algorithms. In these diagrams, the vehicle is commanded to move from the START to END points, with the solid black line showing the ideal path. The vehicle has just established its actual position at P1 after reacquisition of a GPS fix. The position P1 has been affected by wind, waves and current (combined, red vector) and possible inaccuracies in the compass calibration and individual motor thrusts. These steering algorithms calculate a new direction (and possibly speed) to enable



**Figure 2.4:** Three types of steering behaviour shown on *Searise* under the influence of waves, wind and current; (A) Point to target, (B) Waypoint, (C) Dynamic.

## 2.5. CONTROL SYSTEM

---

the vehicle to reach the END point with minimal deviation from the path and arrival time.

All of these methods recalculate direction and speed every iteration of the control-system loop, which is active many times per second. Calculated expected position ('waypoints') move between iterations, so the position to which the vehicle steers towards during one iteration will change by the next iteration.

### Point to Target

The top diagram (A) shows 'point to target' steering. In this mode, the vehicle calculates the compass direction from the current position P1 to the target END and turns towards that direction. As the vehicle has been perturbed off the straight line from start to end, if this perturbation continues (e.g., current and wind effects causing movement to P2), then the path travelled will be asymptotic, as shown in the blue line.

### Waypoint Steering

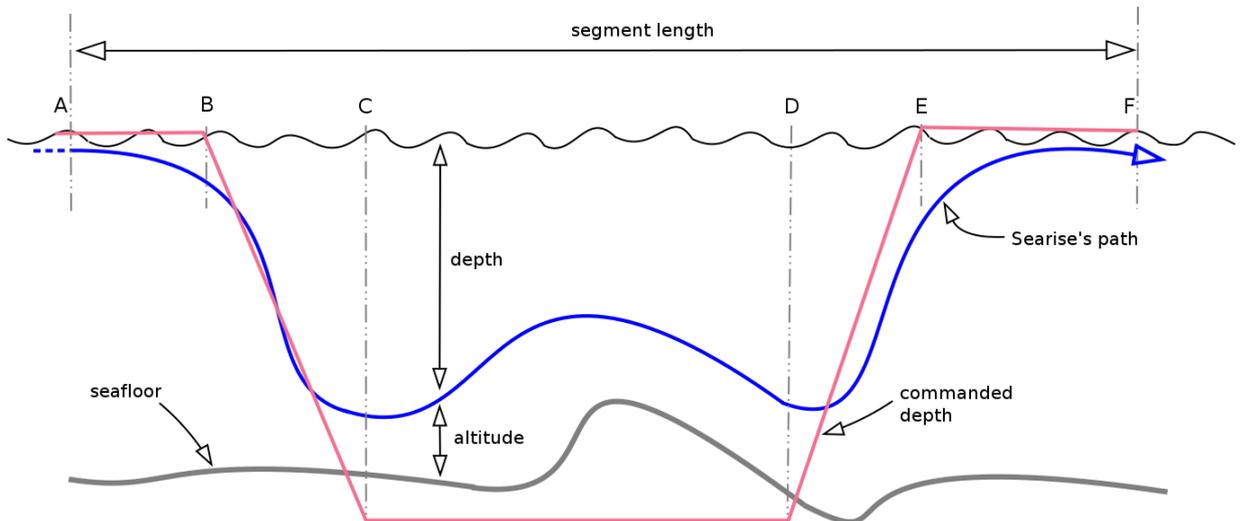
The middle diagram (B) shows waypoint steering. In this mode, the vehicle calculates where it should be along the line from start to end, using the initial parameters of speed and calculated distance and the elapsed time since starting the leg. At any given time, waypoint can be calculated (for P1, position should be WP1). The vehicle then calculates the compass direction from the current position P1 to the expected position WP1 and turns towards that direction. If the waypoint is 'ahead' of the actual position of the vehicle, this system can work reasonably well. However, if the vehicle is moving slower than expected, or has been perturbed such that the target position is 'behind' the current position, then the vehicle will tend to show wandering behaviour, even turning in a complete circle, in an attempt to reach the calculated waypoint. An example is shown where at P2 the waypoint is now WP2 and the vehicle is fishtailing. Waypoints can be effective if the vehicle is at or behind the correct distance along the path, but ineffective if the vehicle is ahead of schedule.

### Dynamic Steering

The lower diagram (C) shows 'dynamic steering'. Like method B, this method calculates a waypoint WP1 where the vehicle should be, based on initial speed and distance, and current time. The offset from that waypoint is used to

calculate an estimated final destination EST based on remaining time, assuming the perturbations accumulates at the same rate. This EST position is pivoted around the end-point to give target position T. The distance from P1 to T divided by the remaining time gives a new speed for the vehicle. By moving the aim-point to a distant location, variations in waypoint position between successive iterations of the control-system loop will have little effect on the vehicle. This method has the advantage of guaranteeing the original demand (speed/time) and remaining on the path (blue line, exaggerated).

### 2.5.4 Diving



**Figure 2.5:** Dive profile and vehicle behaviour. Transects are subdivided into one or more yo-yo dive profiles, one of which is shown here. The dive profile consists of a surface stage (AB), a diving stage (BC), a bottom transect stage (CD), a resurfacing phase (DE), and a GPS reacquisition phase (EF). The bottom transect stage automatically maintains a minimum altitude above the seafloor.

The simplest way of recording spectral data from the seafloor with *Searise* would be to transect in long single dives, maintaining the desired altitude above the bottom. However, as *Searise* does not carry instrumentation allowing it to accurately determine its location when underwater, this method cannot be used. The CSIRO's ASLab has done some work on visual odometry (Dunbabin et al., 2006) for *Starbug*, but much of the Derwent estuary is featureless mud, so this method of localisation cannot reliably be used for this area.

One option is to tow a surface-mounted GPS receiver in a buoy, which has been

## 2.5. CONTROL SYSTEM

**Table 2.4:** Segment lengths of yo-yo dive.

| Segment | Phase            | Portion | Max (m) |
|---------|------------------|---------|---------|
| A-B     | Surfaced         | 5%      | 15      |
| B-C     | Dive             | < 25%   | 25      |
| C-D     | Mapping          | > 40%   | N/A     |
| D-E     | Ascent           | < 25%   | 25      |
| E-F     | Surface, GPS fix | 5%      | 15      |

used for some missions, but this method has disadvantages such as drag on the vehicle and displacement of the buoy caused by current, wave and wind effects.

### Uncertainty in Position

Another approach is to perform short dives, and regularly re-surface to establish the position of the vehicle. Inaccuracies in estimated position accumulate as the vehicle is submerged, but the magnitude of these errors can be reduced by frequent resurfacing and re-establishing a GPS fix.

While the autonomous underwater vehicle (AUV) navigates underwater, it is operating with estimates of position calculated using direction and thrust. Water current and inaccuracy in calculations accumulate potential error the longer the vehicle is submerged. At the point just before the vehicle re-surfaces and re-establishes a GPS fix, the uncertainty in position could be considered maximum. But the re-establishment of GPS position changes the interpretation of maximum uncertainty; this can be considered the point furthest in time from any GPS fix. In the case of a dive segment, the uncertainty is therefore greatest at the halfway point in time between fixes. This uncertainty is used as a part of the mapping process, as described in Chapter 5.

Figure 2.5 shows a conceptual representation of a single segment of an automatic dive profile. The dive is divided into 5 segments (see Table 2.4); A-B, the surface movement during which GPS fix is established; B-C the dive, attaining the requested depth; C-D the profiling segment recording the seafloor spectra from a consistent altitude; D-E the ascent, returning to the surface, and E-F the final surface segment where the vehicle has an opportunity to re-establish a GPS fix and move to the correct final coordinates of the segment.

The depth is commanded in the `MOVETO` command, and is a maximum depth. The

## CHAPTER 2. AUTONOMOUS UNDERWATER VEHICLE

---

altitude is hard-wired to the minimum reliable altitude capability of the altimeter (0.5 m for the PA-500, 1.2 m for the CruzPro), and is enforced by over-riding the commanded depth. This allows bottom-hugging behaviour by simply commanding a depth greater than the seafloor depth (the actual depth achieved will then be the minimum clearance altitude above the seafloor). This is extremely useful for long-duration transects at a set-height above a variable-depth seafloor; the vehicle can simply be commanded to a depth greater than the expected depth of the seafloor, and it will bottom-hug with the hard-wired altitude. The segment-length may be specified or automatically calculated depending on requirements (see Table C.1 for details).

The segment length of C-D is maximised by the software, by minimising the time required for the descent (B-C) and ascent (D-E) legs of the transect.

### Behaviour

The diving behaviour of *Searise* is controlled by a flag byte, which is a part of the MOVETO command. Individual bits in the flag byte switch various control modes (see Table C.1). A number of movement behaviours are available. These include moving to a set position (accurate to a given radius), linear dives to a given depth, profile (yo-yo) dives, automatically calculated yo-yo dives over a given distance, and perform a specified number of yo-yo dives given a desired leg length.

### Yo-yo profile

Dives consist of a surface phase, a descent phase, a bottomed-out phase, a surfacing phase, and a final home to position phase. This up-and-down action is typically referred to as a yo-yo profile. Multiple yo-yo dives may be performed in any transect leg, through segmentation. This is achieved by indicating in the `flags` parameter (and the `segmentCount` parameter) either the number of segments to divide a leg into (`flags` bit 4), or the time for each dive segment (`flags` bit 5), or the length of individual segments (`flags` bit 6). The MOVETO command adjusts all dive parameters to best achieve the requested behaviour.

### Flat Dive and Depth-By-Pitch Dive

*Searise* control systems provide two methods of diving: depth by pitch, and flat diving. It is more efficient for the vehicle to pitch the nose down and tail up then

## 2.6. LIMITATIONS

---

thrust forward when diving, as this allows a more rapid descent. Likewise, a nose-up and tail-down attitude is more efficient for rapid ascent. This control method is not desirable for mapping, however, as the spectrometers are ideally kept in a vertical plane and measuring the same light path, and angular pitch of the vehicle invalidates this configuration. The flat diving control method maintains a pitch of 0 degrees by using the flat-thrusters all thrusting in the same direction (with the rear thrusters at half-power) for depth adjustment. This control method is less efficient and the vehicle is unable to achieve rapid descents or ascents, but for fine control over pitch attitude it is ideal.

Depth-by-pitch is used in the descent (segment B-C) and ascent (segment D-E) phases of the dive profile, and flat diving is used in the mapping phase (segment C-D). The reading from the CruzPro and PA500 altimeters are modified based on the *sin* of the pitch of the vehicle, to convert the altimeter's distance-to-seafloor to a true altitude.

## 2.6 Limitations

---

### Buoyancy

*Searise* has no active buoyancy control. Lead weights and polystyrene floats are added, before a mission, to various areas on the vehicle to configure it so that it is slightly positively buoyant and it will maintain a neutral attitude (roll and pitch zero). In the case of loss of power when underwater, the vehicle can be expected to return to the surface and float.

The lack of an active buoyancy system is problematic in an environment such as the Derwent estuary, which has a combination of fresh water (from upriver) and salt water (from the ocean). These waters of differing salinity and density form distinct layers in the Derwent River, with the fresher water above the saltier water and the mixing layer (the 'pycnocline') around the city of Hobart (Hartmann et al., 2009). This difference in water density also causes a difference in the buoyancy of the vehicle, so adjustments must be made to the flotation, depending on the area of deployment. These buoyancy differences have been a significant impediment to the easy operation of *Searise*.

### Water Currents, Wind and Waves

*Searise* is adversely affected by moderate water currents, and wind, and it does not incorporate the effects of drift when underwater into its estimate of position. Estimating the water current speed and direction based on expected and actual vehicle position after underwater legs, and applying this vector when calculating positional information while underwater could alleviate this. Some versions of the vehicle use optical tracking of features on the seafloor to improve navigational accuracy. Another option is to carry more sophisticated inertial measurement unit (IMU) devices (for example, a Doppler velocity log).

The tail of the vehicle acts somewhat like a sail when surfaced in windy conditions, and this can significantly affect the steering and speed of the vehicle. The relatively slow speed of the vehicle, and its inability to operate in moderate currents and wind is a significant issue. Currents as low as 2 knots can severely compromise the ability of *Searise* to perform a mission.

### Safety

The operation of an AUV in a busy waterway presents challenges related to safety. As the vehicle is unable to detect, yet alone avoid, other watercraft, all missions are tethered, or accompanied by either a boat or a person in a kayak. When other vehicles approached, they were warned by radio of the potential hazard, or *Searise* was removed from the water.

To improve safety, the addition of a ‘diver’ flag onto the Wi-Fi buoy was trialled, with the hope that the restrictions applying to marine craft approaching these flags will be honoured, and other marine craft will maintain a safe distance. But the reality is, the buoy/flag is quite hard to see in the absence of an accompanying boat and so this has not proven an effective solution to the problem of potential collisions.

Another hazard involves the use of a towed buoy, such as the Wi-Fi buoy. There is a danger that the rope may become entangled in the propeller of a watercraft passing over it.

For the above reasons, the autonomous operation of the vehicle has been somewhat compromised by the need to follow it closely with an accompanying vessel and human operator. Legal and safety issues related to the operation of autonomous vehicles, particularly in heavily trafficked areas, are still mostly uncertain.

## 2.7 Limitations

---

### 2.7.1 Suitability of Vehicle

*Searise* has proven to be a challenging vehicle to operate in some environments. It cannot move through the water quickly, and moderate currents, waves and wind influence it. This has been mitigated somewhat by the dynamic steering behaviour (see Section 2.5.3).

Fouling by taller seagrasses was a problem during some deployments at Bruny Island. In some cases, the seagrasses reached to the water surface and consequently the vehicle had no clearance. The ‘wing’ protrusions of *Searise* were prone to capture by these longer seagrasses and the vehicle in those situations became immobile. Consequently, mapping missions were only attempted where there was adequate clearance for the vehicle.

Another fouling issue related to the ingress of seaweeds and other material such as magnetic particles (see de Souza et al. (2010)) into the propeller assembly of all motors. Particularly near Dogshear Point in the Derwent River, the fine material often loose in the water had a tendency to wrap around the propeller shaft and seize the motors. Mesh shrouds were installed over the motors, and these greatly alleviated this problem.

In summary, *Searise* can be adversely affected by fouling if it strays too close to seagrasses and seaweeds. To some extent this is avoided by operating at altitude, however many of the areas where measurements were performed involved extremely shallow waters where this was not always possible.

## 2.8 Summary

---

This chapter has described *Searise*. The vehicle has proven to be a capable underwater robotics platform, particularly suitable for this research into benthic mapping. It is useful for mapping small areas in great detail.

*Searise* has the ability to transect across the seafloor at very low altitudes, while operating high-resolution optical spectrometers collecting both down-welling and reflected light. It provides ample storage for instrument data collected over relatively long-duration mission times. It has excellent connectivity, making the process of communication and data retrieval relatively painless.

## **CHAPTER 2. AUTONOMOUS UNDERWATER VEHICLE**

---

The vehicle is relatively slow moving, can be adversely affected by even moderate currents, and requires continuous supervision when operating in heavy-use waterways. The control software is relatively primitive, and the machine has no fail-safe capability, hazard avoidance, or ability to respond to unexpected events.

*“Indeed how might it be if things revealed their colours only when (in our terms) no light fell on them - if, for example, the sky were black? Could we not then say, only by black light do they appear to us in their full colours?”*

Ludwig Wittgenstein (1889-1951)

*“The real voyage of discovery consists not in seeking new landscapes, but in having new eyes.”*

Marcel Proust (1871-1922)

# 3

## Measuring Light

### 3.1 Introduction

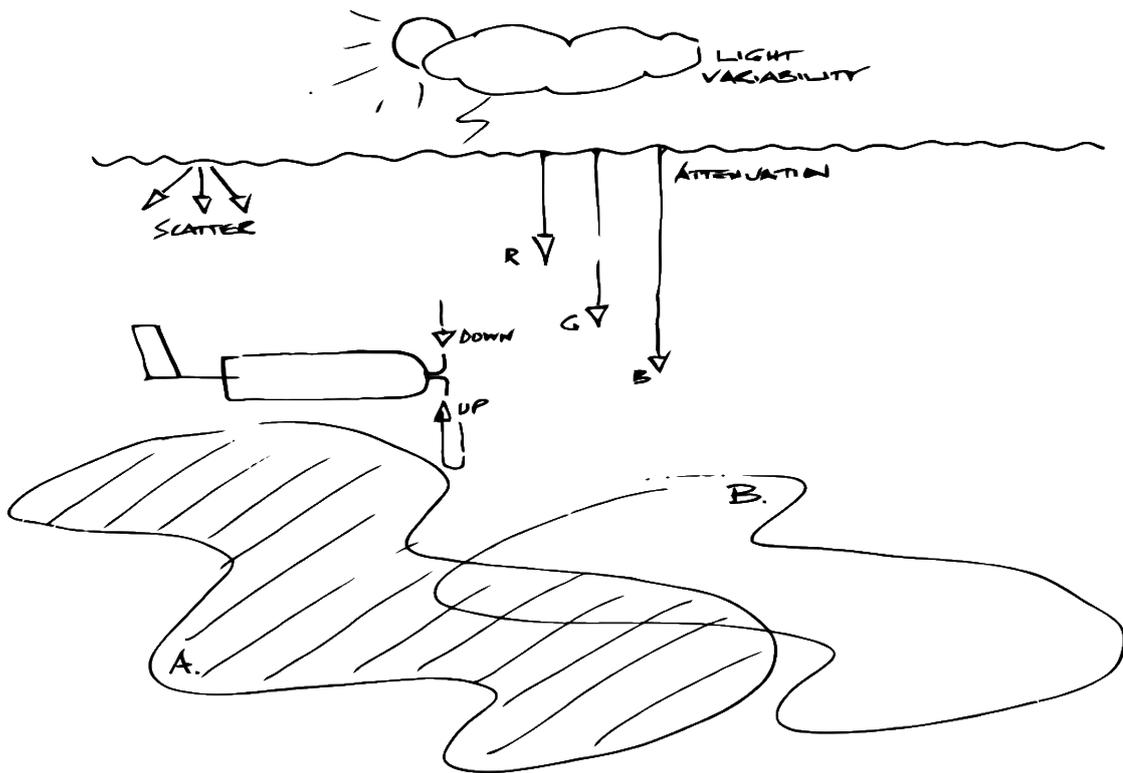
---

This chapter presents an overview of the behaviour of light passing through a body of water, and interacting with benthos and instrumentation. The capabilities of the Ocean Optics USB2000+XR-1 miniature optical spectrometers (XR-1s) and associated systems are explained, along with the consequent limitations. The format of data recorded by these instruments is presented, and the detailed operation of the interface/driver software is explained. The results of laboratory measurements of noise and optical efficiency are given, and the effects on use of the instruments for autonomous mapping are discussed.

### 3.2 Light Propagation in Water

---

Figure 3.1 shows a conceptual view of the collection of spectral data by the autonomous underwater vehicle (*Searise*). Light from the Sun (see Figure 3.9) enters the atmosphere and is attenuated by clouds, rain, airborne dust and gases. It then enters the water body. The light further decays in intensity as it passes through the water. This decay of light intensity results from absorption and scattering.



**Figure 3.1:** Conceptual view of habitat mapping. The autonomous underwater vehicle carries two optical spectrometers, one recording the light passing downwards, and the other recording the light reflecting off the bottom. The atmosphere and the water attenuate the light from the Sun. It is scattered within the water body, and then absorbed and reflected by the plants, sand and mud on the seafloor. These benthic objects all absorb and reflect the light differently. Two types of seafloor are shown here, marked A and B.

### Absorption and Scattering

Absorption is the conversion of electromagnetic waves into heat or chemical energy. The main absorbers in seawater include photosynthesis by algae, particulate matter, dissolved organic compounds, and the water itself. This process is wavelength-dependent.

Scattering occurs when electromagnetic radiation interacts with suspended particles in the water column and changes direction by reflection, which could induce “change” in colour. This process is caused by inorganic and organic particulate matter, and increases with turbidity. Change in direction of photons increases the distance they travel through the water and therefore the likelihood of them being absorbed.

### 3.3. DESCRIBING THE SEAFLOOR

---

#### Path

The down-welling light is measured by the upward-pointing spectrometer, and continues to the seafloor, where it is absorbed and reflected by the benthos. The reflected light continues upwards and is attenuated and scattered by the water until it reaches and is measured by the downward-pointing spectrometer.

After collection by the spectrometer fibre optics, the light is further attenuated at the interface between the fibre optic cable and the water, by the fibre optic cable itself, and the interface between the fibre-optic cable (FOC) and the spectrometer. Finally, only a small percentage of the light reaching the spectrometers is actually detected by the spectrometers as a result of the inefficiency of the diffraction grating used by the instruments.

#### Variability of Input Light

The variability of the input light poses a problem for measuring spectra of objects. The objects can only absorb and reflect the light that reaches them, so with variable input light, there is also variability in the reflectivity spectra. This problem is usually addressed by using a reference material in the field, and measuring the reference material's spectrum and then measuring a sample's spectrum, allowing the production of a reflectivity measurement.

### 3.3 Describing the Seafloor

---

This section describes a form of the linear mixing model (LMM), described in Section 3.3.2.

The complete benthic community,  $\overline{C}$  (equation 3.3.1), consists of a mixture of  $n$  types of 'objects'. These include species of seagrass, types of bottom-cover (e.g., sand), starfish, etc. These are referred to as *end-members*. An exemplar spectral representation of any of these objects, stored in an end-member spectral library, contains reflectance measurements at  $m$  wavelengths. The objects are assumed to be Lambertian reflectors (Lesser and Mobley, 2007)<sup>1</sup>. Each of these measurements records the mean percentage of light,  $\mu_\lambda$ , reflected by objects of the given type at wavelength  $\lambda$ . The range of values of  $\mu$  is  $0 \leq \mu \leq 1$ .

---

<sup>1</sup>Lambertian reflectance is the property that the apparent brightness of a surface is the same regardless of the angle of view.

Exemplar end-member spectra are representative spectra of the objects. Variations in spectral reflectivity of individual objects exhibit a normal distribution around  $\mu$ .

$$\bar{C} = \begin{bmatrix} \overline{Obj}_1 \\ \overline{Obj}_2 \\ \vdots \\ \overline{Obj}_n \end{bmatrix}, \quad \overline{Obj}_i = [ \mu_1 \ \mu_2 \ \dots \ \mu_m ] \quad (3.3.1)$$

### 3.3.1 Community of Light

The mix of objects in any location (specifically, within the field of view (FOV) of the spectrometer) is described by equations 3.3.2, where  $\omega_i$  is the abundance of object type  $i$ . The range for abundance is specified in equation 3.3.4. The abundance,  $\omega$ , of any object may be zero (i.e., the object is not present in the FOV), but the total sum of all abundance values at any location is equal to one (equation 3.3.3).

$$\bar{C}_{(x,y)} = \begin{bmatrix} \omega_1 \cdot \overline{Obj}_1 \\ \omega_2 \cdot \overline{Obj}_2 \\ \vdots \\ \omega_n \cdot \overline{Obj}_n \end{bmatrix} \quad (3.3.2)$$

$$\sum_{i=1}^n \omega_i = 1 \quad (3.3.3)$$

$$: \omega_i \geq 0 \quad (3.3.4)$$

The area of seafloor recorded by the FOC ‘lens’ of the spectrometer varies based on the relative geometry of the autonomous underwater vehicle (AUV) and the seafloor, and the light sensitivity cone of the FOC (see Table 3.3). The spectrometer measures a single spectral pixel independent of AUV/seafloor geometry. This sampled spectra (equation 3.3.5) is a mix of spectra of all visible objects within the spectrometer’s FOV.

$$\bar{S}_{(x,y)} = \sum_{i=1}^n \omega_{(x,y)_i} \cdot \overline{Obj}_i \quad (3.3.5)$$

### 3.3. DESCRIBING THE SEAFLOOR

---

The objects within the FOV contribute to the recorded spectra in the same proportion as they are visible in the FOV. The mixing of spectra is linear.

The classification task involves recovering  $\bar{\omega}_{(x,y)}$ , given  $\bar{S}_{(x,y)}$ . This may be assisted by an end-member spectra library  $\bar{T}$  of  $t$  exemplar spectra (equation 3.3.6) where each member of  $\bar{C}$  is ideally represented in exemplar by a single member of  $\bar{T}$ . It is possible for an object to not be represented in  $\bar{T}$ , or for variation in spectra within an object class to be so different that an object is represented multiple times in  $\bar{T}$ . The spectra of  $\bar{T}$  may also represent composite objects. Equation 3.3.7 shows the reconstructed classification from the end-member library, an estimate of Equation 3.3.5.

$$\bar{T} = \begin{bmatrix} \overline{Template_1} \\ \overline{Template_2} \\ \vdots \\ \overline{Template_t} \end{bmatrix} \quad (3.3.6)$$

$$\widehat{R}_{(x,y)} = \sum_{t=1}^n \omega_{(x,y)t} \cdot \bar{T}_t \quad (3.3.7)$$

#### 3.3.2 Formal Description

An observation must deal with sensor noise, end-member variability, the presence of unknown object types, and other inadequacies. Keshava (2003) describes an observed spectrum thus; “If we have  $K$  spectral bands, and we denote the  $i^{th}$  end-member spectrum as  $\mathbf{s}_i$  and the abundance of the  $i^{th}$  end-member as  $a_i$ , the observed spectrum  $\mathbf{x}$  ... can be expressed as

$$\begin{aligned} \mathbf{x} &= a_1 \mathbf{s}_1 + a_2 \mathbf{s}_2 + \dots + a_M \mathbf{s}_M + \mathbf{w} \\ &= \sum_{i=1}^M a_i \mathbf{s}_i + \mathbf{w} \\ &= \mathbf{S} \mathbf{a} + \mathbf{w} \end{aligned} \quad (3.3.8)$$

where  $M$  is the number of end-members,  $\mathbf{S}$  is the matrix of end-members, and  $\mathbf{w}$  is an error term accounting for additive noise.”

### 3.4 Optical Spectrometers

---

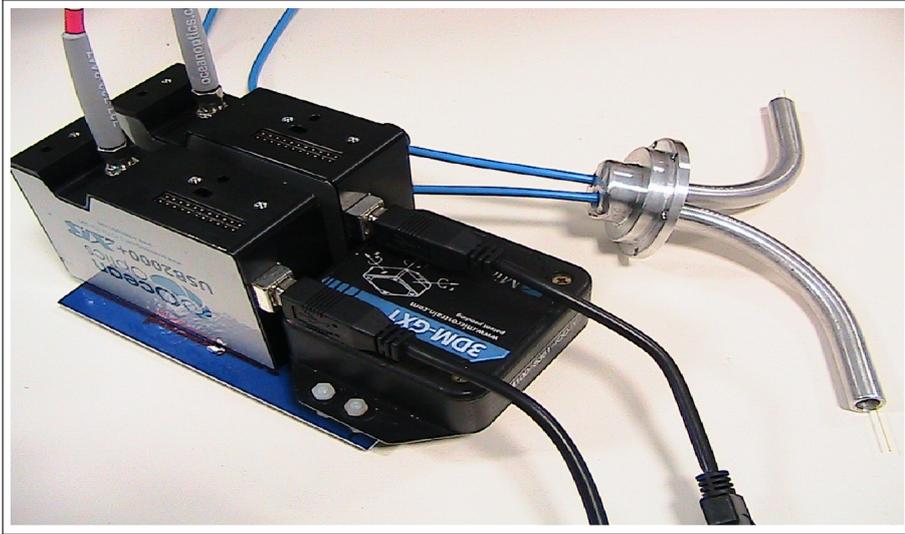


**Figure 3.2:** Ocean Optics USB2000+XR1 miniature fibre-optic spectrometer. Size: 89 mm x 62 mm x 34 mm. *Searise* carries two of these devices internally, each connected to an external fibre-optic guide arm (seen in Figure 3.3).

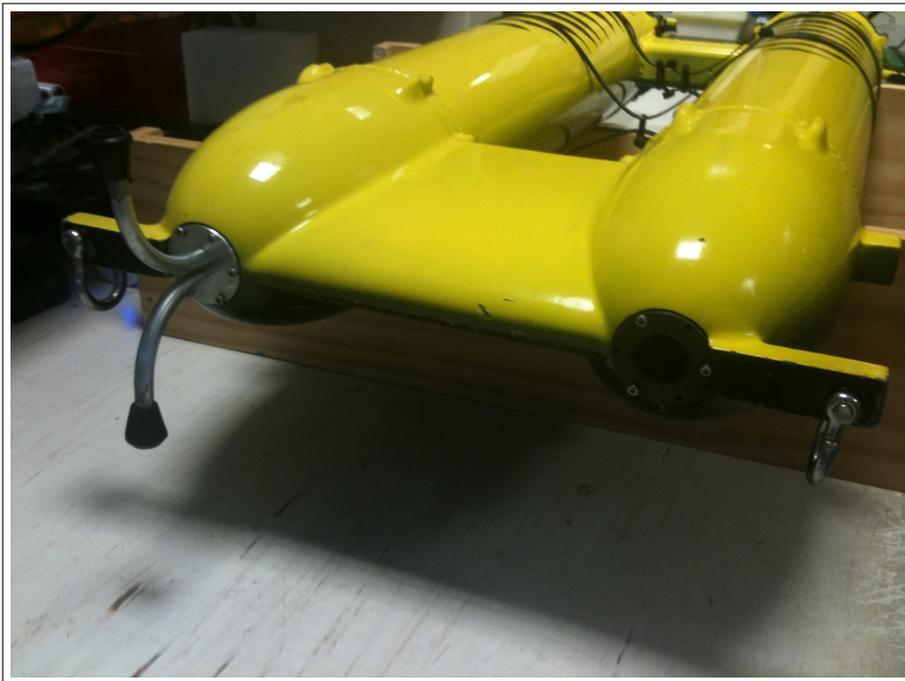
In keeping with the low-cost design of the vehicle, the author modified *Searise* to carry two Ocean Optics USB2000+XR-1 miniature optical spectrometers (XR-1s). The spectrometers (see Figure 3.2 and Table 3.1) are carried inside the vehicle and connected via fibre-optic cables (FOCs) to an external guide with opposed arms (designed for this research; see Figure 3.3 and Figure 3.4), attached to the starboard nose cone. One spectrometer detects down-welling light from above, and the other detects light reflecting from the benthos below. Both also capture incidental light scattered through the water.

Each spectrometer has slightly different specifications (see Table 3.2). These spectrometers also show different spectra in response to darkness. This base spectrum for each device has been measured by enclosing the collection arms in a lightproof enclosure and recording an averaged spectra.

### 3.4. OPTICAL SPECTROMETERS



**Figure 3.3:** Internal mount holding two spectrometers. Also shown, associated fibre-optics, external guide before potting, and the AUV's inertial measurement unit.



**Figure 3.4:** Optical fibre guide installed in place of *Searise's* starboard front window. The rubber end-caps protect the light-collecting surface, and are removed during missions. The curve of the guide arms respects the physical resistance constraints of the optical fibres. The vehicle is approximately 50 cm in width.

## CHAPTER 3. MEASURING LIGHT

---

**Table 3.1:** Optical Spectrometer Specifications.

|                           |   |
|---------------------------|---|
| Spectrometer Type         | Ocean Optics USB2000+XR-1                                     |
| Detector                  | Toshiba TCD1304AP Linear CCD array                            |
| Dimensions                | 89.1 mm × 63.3 mm × 34.4 mm                                   |
| Weight                    | 190 g   |
| Grating                   | Grating #31, 500 lines/mm<br>blazed at 250 nm                 |
| Spectral Range            | 200 - 1025 nm   |
| Entrance Slit             | 25 $\mu\text{m}$  |
| Optical Resolution (FWHM) | 1.7 - 2.1 nm  |
| Channels                  | 2048  |
| Integration time          | 3.8 ms - 10 s   |
| Dynamic range             | $3.4 \times 10^6$ (system)<br>1300:1 for a single acquisition |
| Sensitivity               | 60 photons/count at 600 nm<br>130 photons/count at 400 nm     |
| Signal-to-noise ratio     | 300:1 (at full signal)  |

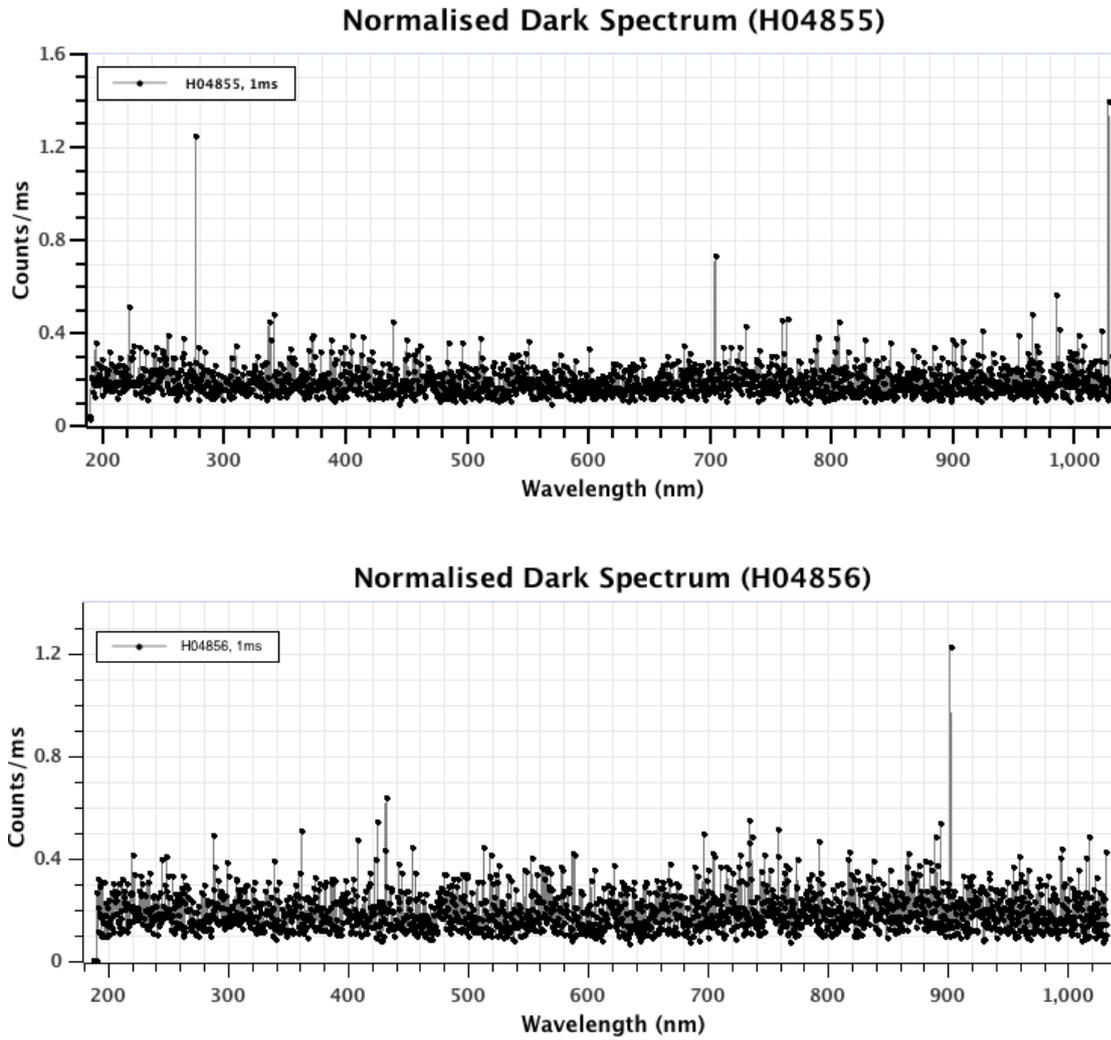
### 3.4.1 Dark Spectra

The dark spectra are a measure of the inherent noise in each spectrometer. The dark spectrum for each device is scaled based on sample time and subtracted from the measured spectra to effectively remove instrument noise. If this subtraction results in a negative count at any wavelength, then zero is used as a count for that wavelength. This process also caters for the anomalous peaks evident in the black body spectra graphs.

Although the spectrometers have an internal dark spectra measurement, these dark spectra were explicitly measured so that the spectra incorporated the entire physical light conduit was included (fibre-optic cable, interface connections, and internal to spectrometers).

For device H04855 (see Figure 3.5), the dark spectrum was integrated for 10 s and converted to a 1 ms equivalent. The peaks at 275.94 nm, 703.92 nm and 1028.94 nm are anomalous but consistent for this unit. For device H04856, the dark spectrum was integrated for 45 s and converted to a 1 ms equivalent. The peak at 901.73 nm is an anomaly but consistent for this unit. These spectra represent average noise

### 3.4. OPTICAL SPECTROMETERS

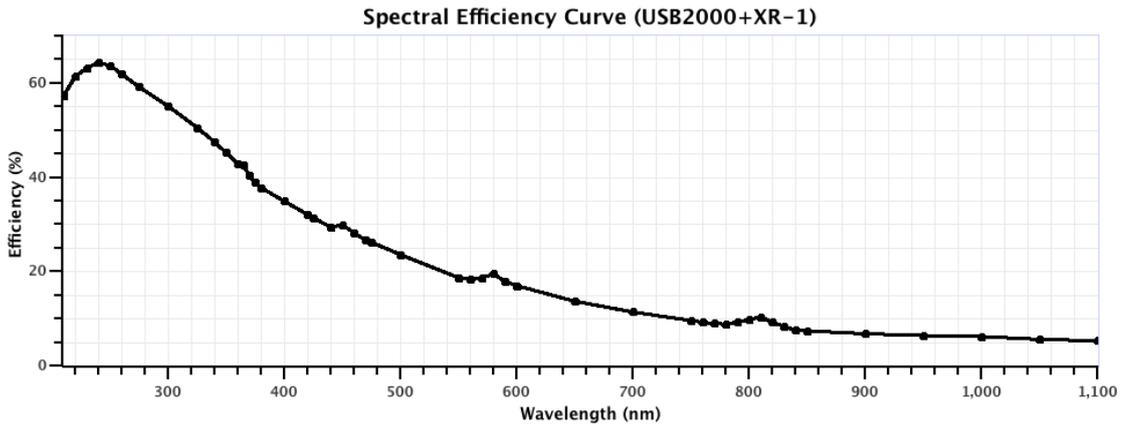


**Figure 3.5:** Normalised spectra of instrument noise in spectrometers H04855 and H04856. The fibre connecting the spectrometers was divided into two halves, and these were glued into the external guide arms. The assembly was then placed in a dark enclosure and the spectra for each instrument was measured. The spectra shown here represent instrument noise.

## CHAPTER 3. MEASURING LIGHT

**Table 3.2:** Characteristics of each spectrometer.

| Serial Number | Firmware | Min (nm) | Max (nm) | Band Width (nm) |
|---------------|----------|----------|----------|-----------------|
| USB2+H04855   | 1.04.0   | 187.7    | 1033.3   | 0.4             |
| USB2+H04856   | 1.04.0   | 188.1    | 1031.4   | 0.4             |



**Figure 3.6:** Typical grating optical efficiency curve.  
*from data kindly provided by Ocean Optics.*

per millisecond, and are scaled and subtracted from the associated spectrometer's readings to obtain spectra with minimised instrument-related noise.

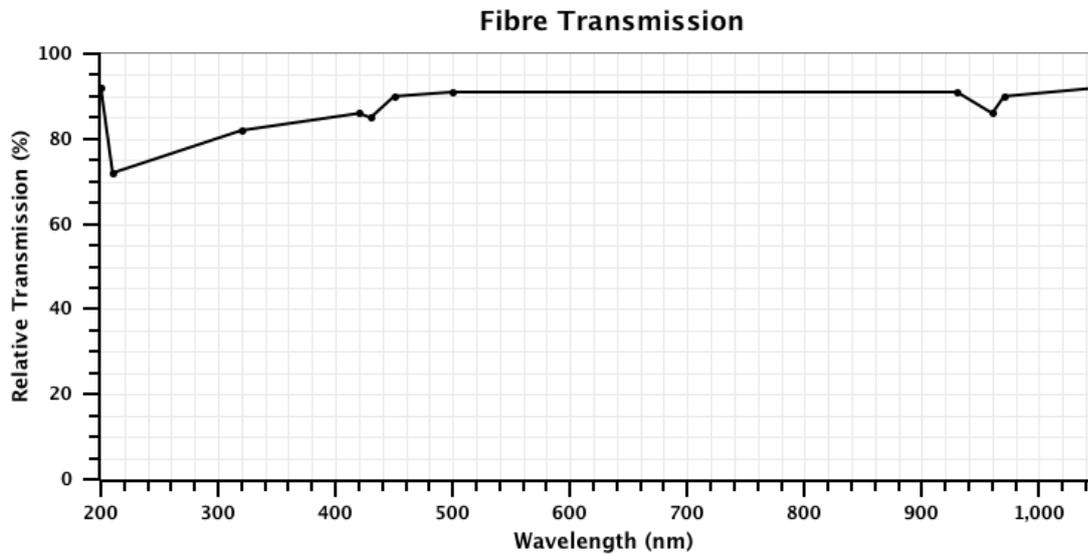
### 3.4.2 Spectrometer Efficiency

Not all of the light present at the FOC interface is detected by the spectrometers. The FOC attenuates the light it carries (see Figure 3.7), and the spectrometers only detect a proportion of the light they receive.

A typical efficiency curve for the spectrometers is shown in Figure 3.6, which was generated from data in Table D.1. Individual spectrometers vary in wavelengths detected, and in the optical efficiency of the instrument at any given wavelength. As the efficiency of each spectrometer is unique, the efficiency curve for a particular spectrometer must be measured against a calibrated light source. Ocean Optics supplies spectrometers with a response guaranteed to be within 10% of a defined reference spectrometer, but these units are expensive and not carried on board the AUV. Consequently, the efficiency curve for each spectrometer used in this

### 3.5. FIBRE-OPTIC CABLE

---



**Figure 3.7:** Fibre-optic cable efficiency curve.  
*from data kindly provided by Ocean Optics (see Appendix D)*

research is only roughly comparable to the example curve shown, and the exact efficiency of each is not measured.

## 3.5 Fibre-Optic Cable

---

The spectrometers use an optimal (in terms of ability to carry light in the wavelength range detectable by the XR-1) fibre-optic cable (FOC), purchased from Ocean Optics (see Table 3.3). The FOC was cut into two halves, and the outer sheath of each half stripped from the last few centimetres before the cut point. The fibres were mounted in the arm extension, one fibre in each arm, and sealed in place with epoxy resin. The ends of each arm were sanded and polished to a smooth finish. Thus, a single FOC is used to provide a light pathway from the outside world to the two spectrometers inside *Searise*.

The FOC attenuates light differently at different wavelengths, as shown in Figure 3.7. The calibration sheet for the FOC used on *Searise* is presented in Appendix D. The light that enters the spectrometers is attenuated at the interface between the FOC and the water, when travelling through the FOC itself, and further attenuated by the XR-1's efficiency at detection as a function of wavelength (see Figure 3.6).

## CHAPTER 3. MEASURING LIGHT

---

**Table 3.3:** Fibre-Optic Cable Specifications.

|                        |                          |
|------------------------|--------------------------|
| Fibre-Optic Cable Type | Ocean Optics P400-025-SR |
| Core Diameter          | 400 $\mu\text{m}$        |
| Wavelength Range       | 200-1100 nm              |
| Light sensitivity cone | 19° in water, 25° in air |
| Buffer / Coating       | polyimide                |
| Long term bend radius  | 16 cm                    |
| Short term bend radius | 8 cm                     |

To recover the original light spectra entering the FOC, a reverse function could be applied to firstly adjust for the spectrometer efficiency curve, and secondly adjust for the FOCs attenuation. It would be difficult to quantify these parameters, as they are dependent on the physical characteristics of the instruments and equipment. For example, roughness on the surface of the extension arms may significantly affect the light entering the FOC.

Figure 3.8 combines the attenuation caused by the fibre optic transmission and the spectrometer efficiency. The bandwidth has been restricted to the common wavelengths available in both data sets. It is clear that a small fraction of the available light at higher wavelengths is actually detected.

### 3.6 End-Member Spectra Collection

---

Light reflected from a sample is a function of the input light. Figure 3.9 shows a reference solar spectrum. The spectra of a sample lit by different light sources are different.

The presence of so many potential barriers to perfect collection and detection of incoming light (the FOC light interface, the long-term bend radius of the FOC exceeded, the connection between the FOC and the XR-1 being under tension having being discussed) makes it difficult to measure with confidence the absolute intensity of incoming light. More to the point, these issues to some extent being present in both the upward and downward spectrometer instrumentation means that consistency between the units measuring exactly the same input is difficult to achieve.

### 3.6. END-MEMBER SPECTRA COLLECTION

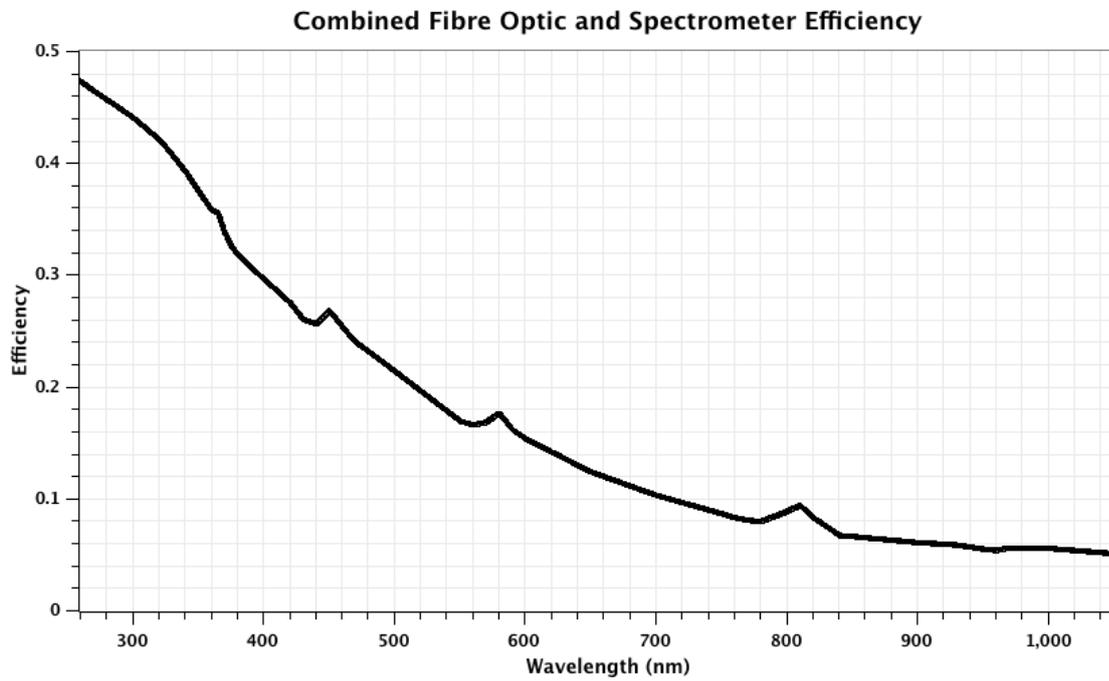


Figure 3.8: Combined fibre-optic and spectrometer efficiency curve.

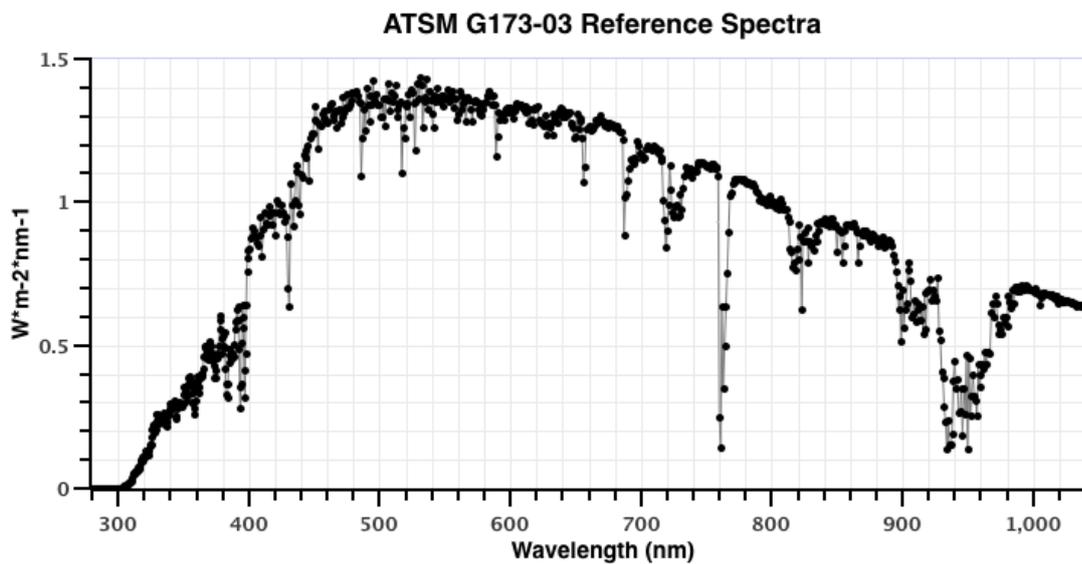


Figure 3.9: ASTM G173-03 direct + circumsolar reference spectrum.

### 3.7 Dependency on Particular Hardware

---

An end-member spectral library should ideally represent the actual reflectivity of the end-members without effects from the light source or measuring equipment. Calculation of reflectivity uses a normalisation process where measurements are made on a white reflectance standard surface, followed by measurement of the target end-member. This allows normalisation of the end-member spectrum through an analysis of the differences in the white reflectance standard spectrum against expectations.

There are several factors that affect the quality of the measured spectra:

- Efficiency of the spectrometer.
- Attenuation by the optical fibre.
- Variability of the light.
- Attenuation of light by the water.
- Variability of the end-members.

A drawback with the normalisation process using a white reflectance standard is the need to measure the white reflectance standard along with the end-member. This physical placement of the white reflectance standard makes this approach infeasible for use on an AUV. The normalisation of spectra is critical for matching measured spectra to predefined end-member libraries.

By using reflectivity spectra, the effects of light variability are mostly removed.

#### 3.7.1 Efficiency of the Spectrometer

The proportion of light detected by each XR-1 is not consistent across the range of wavelengths measured. The spectra gathered with these instruments are not the true spectra of light reflected off objects. The use of these spectra to represent end-member library objects effectively couples the end-member spectra to the specific spectrometer brand/model. With differences between individual units, the coupling extends to the units themselves.

The end-member spectra therefore represent the light detected by the instrument, as attenuated by the FOC and the efficiency of the spectrometer, and this is specific to the particular instrument configuration. This is a proportion of the light reflected by the object. This coupling enforces the use of this model spectrometer,

## 3.8. FACTORS AFFECTING LIGHT

---

or spectrometers with similar efficiency characteristics to the XR-1 unit, if desiring to use the raw spectra of end-member library data elsewhere.

Previously, work with *Searise* on calculating the attenuation of light underwater was performed by Hartmann et al. (2009).

### 3.8 Factors Affecting Light

---

Figure 3.1 gives a basic overview of the operation of the spectrometers in collecting light. As light passes through the water it is attenuated, with the longer wavelengths (reds) more highly absorbed by the water. Light is also scattered through the water body. Light entering the upward-pointing spectrometer arm represents a combination of direct light from the surface, and scattered light from the water body.

The light that reflects off the seafloor and benthos is recorded by the downward-pointing spectrometer. This light has undergone attenuation by the water body (relative to the light recorded at the upward-pointing spectrometer) of twice the altitude of the vehicle (once on the downward path, and once on the upward path).

#### 3.8.1 Format

**Table 3.4:** Wavelength Ranges.

|                      | Min (nm) | Max (nm) |
|----------------------|----------|----------|
| XR1 #4855            | 187.69   | 1033.33  |
| XR1 #4856            | 188.06   | 1031.43  |
| Buiteveld absorption | 250      | 1300     |
| Fibre Efficiency     | 200      | 1050     |

The range of wavelengths at which available spectra have been recorded is dependent on the range and resolution of the instruments used (see Table 3.4), and availability of suitable data. End-member data are reflectivity percentage values (from 0 to 1). Each value represents the percentage of light at the given wavelength that is reflected from the object.

The XR-1 specifications list a range of wavelengths from 200 nm to 1025 nm, but each unit varies slightly from these values because of the physical characteristics

of the diffraction grating each uses. The characteristics of the two spectrometers used in this work are shown in Table 3.4.

Spectra are stored as comma-separated-value ASCII text files, or as Excel-compatible spreadsheet files.

### 3.9 Factors Affecting Unmixing Accuracy

---

The use of an end-member library as components for unmixing spectra presupposes that the spectrum of an observed end-member is the same as the spectrum of that end-member in the library. As discussed in this chapter, a number of factors contribute to the attenuation and distortion of light forming an observation, and the resultant spectrum could bear little similarity to the library spectrum. These attenuation and distortion processes are wavelength dependent. Consequently, unmixing can fail to provide accurate solutions because of the difference between observed spectra and library spectra of end-members.

One technique to correct this issue is to attempt to reconstruct an un-attenuated, undistorted spectrum of each observation. In principle, this is possible because the attenuation caused by water can be reasonably estimated and an inversion of that attenuation can be applied to the observed spectra. Likewise, the attenuation caused by the measuring instruments (fibre-optic attenuation, spectrometer attenuation) is well defined, as described in this work. These factors can also have inversions applied to the observed spectra, and in this way an estimated reconstruction of the observation sans attenuation and distortion can be made. This reconstruction can then be passed to the unmixing algorithm and be more likely to accurately unmix given the end-member library.

#### Absorption by Water

Cho and Lu (2010) give a water-depth correction algorithm for submerged vegetation spectra.

Pegau et al. (1997) have reported on the effect of temperature and salinity on absorption by water. Their results show a linear relationship between absorption, and salinity and temperature. They report “linear slopes can be used to correct the absorption coefficient for changes in both temperature and salinity encountered in natural waters.”

The absorption coefficient of water varies with the wavelength of the incident

### 3.9. FACTORS AFFECTING UNMIXING ACCURACY

---

light. Available sources of absorption coefficients were very consistent in their measurements.

#### Beer-Lambert Law

The Beer-Lambert law (also know as Beer's law) describes the relationship between absorbency and concentration of a material through which light is travelling. This law can be used to predict the attenuation of light travelling through a body of water. It describes a logarithmic dependence between the transmission,  $T$ , of light through a substance and the product of the absorption coefficient of the substance,  $\alpha$ , and the distance the light travels through the material,  $\ell$ .

#### Scattering

The water itself, and suspended matter in the water can also scatter light, and the combined absorbency and scattering describes the attenuation of the light passing through the water. Beer-Lambert's law does not account for such scattering.

In pure waters, visible light is absorbed more at the longest wavelengths. Reds, oranges and yellows are absorbed at shallow depths, and only blue and violet wavelengths reach deeper waters. This characteristic causes open waters to appear blue to the human eye. In coastal waters with more phytoplankton, Chlorophyll-a pigments absorb light, and the plants scatter light. Thus, coastal waters are less clear than open ocean waters. As Chlorophyll-a absorbs strongly in blue and violet of the visible spectrum, the green wavelengths reach the deepest and thus coastal waters often appear green-blue or green to the human eye.

#### Shadowing

There is some potential for the shadow of *Searise* to affect the spectral data. However, the location of the XR-1 inputs ahead of the front of the vehicle, and the general lack of shadows evident when gathering data, suggest that this was not a problem for data gathered during this research.

The addition of a light source on the vehicle, illuminating the area measured by the downward-spectrometer, alleviates any shadowing issue. The spectrum of the light source is known in advance. This allows mapping in dark and turbid environments, such as shown in Figure 2.3. The contribution of the upward-spectrometer measurement is minimised.

### 3.9.1 End-Member Library

A library of spectral recordings for objects in Australian coastal waters has been built by CSIRO Land and Water (Dekker et al. (2003) and references therein). The library consists of spectral data recorded at ground level, measuring surface reflectivity of end-members. This was calculated by measuring the radiant flux from each object and normalising this using the radiant flux from a white reflectance standard (representing the light falling on the target).

Spectral irradiance reflectance (also known as irradiance ratio and reflection coefficient)  $R(z, \lambda)$  is defined as the ratio of spectral up-welling to down-welling plane irradiances at depth  $z$  and wavelength  $\lambda$  (Mobley, 1994).

$$R(z, \lambda) = \frac{E_u(z, \lambda)}{E_d(z, \lambda)} \quad (3.9.1)$$

alternately,

$$\bar{R} = \frac{\bar{E}_u}{\bar{E}_d} \quad (3.9.2)$$

where  $\bar{R}$  is the fraction of down-welling light that reflects upward at a given depth,  $z$ .

An end-member spectrum is considered representative of a species; variability of reflectance within a species due to morphological differences and seasonal changes can be problematic for the use of end-member libraries for classification. This work presents an algorithm for automatically adding new spectra to the end-member library as a way of catering for these variations (see Section 4.9).

Randomly generated reflectivity spectra representing end-members were used for much of this research.

#### First Order Derivative

Kutser et al. (2006b) reports that shape of the spectra of the macro-algae they studied was consistent within each group (red, green, brown algae), with most of the variability in reflectance values, i.e. the brightness of the samples. First-order derivatives of the spectra are commonly used to analyse the shape of the spectral curve and remove the differences caused by brightness. This technique is not necessary, however, if reflectivity spectra are used.

## 3.10. LIMITATIONS OF END-MEMBER LIBRARIES

---

### 3.9.2 Calibration

Successful unmixing of laboratory-mixed spectra (with added noise) was successfully performed with software using these templates (and similar spectra gathered from Cloudy Bay Lagoon using *Searise*'s spectrometer) using a variety of techniques, including genetic algorithm (GA), simulated annealing (SA), linear algebra (LA), and Non-Negative Least Squares (NNLS). There are a large number of techniques for spectral unmixing in the literature.

As did Hochberg and Atkinson (2000), to convert the radiance measurements to irradiance, I have assumed that all benthic surfaces reflect radiance equally into all directions (i.e., they are Lambertian reflectors).

### 3.9.3 Reflective Spectra

Representation of end-member spectra in the library is in a form that is independent of the lighting of the sample. By representing the spectra as the proportion of reflected light, the end-member library becomes independent of input light.

## 3.10 Limitations of End-Member Libraries

---

The use of a single end-member spectrum to represent a species of seagrass, or type of object has inherent limitations. Variability of light absorption within individuals of a species because of morphological differences, different environmental conditions, health of individual plants, and other factors are not accounted for. The brown decaying leaf is spectrally dis-similar to the newly formed shoot. Single end-member spectra cannot account for this variability.

On a larger scale, single end-member spectra can be generated from an average of many observations of typical examples of the end-member, or from averaged observation of large areas containing the end-member. Where the 'pixel-size' of an observation covers a large physical area and includes large amounts of the end-member (for example, with satellite and aerial remote-sensing) this assumption of a single end-member spectra is reasonable. At finer-scales, particularly at the level approaching individual organisms, this assumption breaks down and poses a problem for accurate classification using end-member libraries.

## CHAPTER 3. MEASURING LIGHT

---

This work avoids the end-member selection problem by proposing an automatic end-member detection algorithm (see Section 4.9).

### 3.11 Summary

---

This chapter has discussed the detection of light reflecting off benthos. The spectrometers' limitations, and issues related to differences between individual units have been addressed. The factors affecting the attenuation of light create a dependency between the measuring instruments and the end-member library. Commonly, measurements are corrected to try and remove these attenuating factors before matching to end-member libraries.

*“In general, the more one augments the number of divisions of the productions of nature, the more one approaches the truth, since in nature only individuals exist, while genera, orders, and classes only exist in our imagination.”*

Comte Georges-Louis Leclerc de Buffon (1707-1788)

*“We see only what we know.”*

Johann Wolfgang von Goethe (1749-1832)

# 4

## Unmixing

### 4.1 Introduction

---

This chapter examines the process of unmixing spectral data to recover the type and proportion of spectrally distinct objects that have contributed to observed spectra. These spectrally distinct objects are referred to as ‘end-members’.

There have been many methods proposed for solving the unmixing problem, ranging from algebraic techniques such as linear algebra and quadratic programming, non-negative least squares optimisation, through to nature-inspired algorithms such as genetic algorithms (GAs) and simulated annealing (SA). Both quadratic programming and non-negative least squares offer techniques for rapid solution of the constrained unmixing problem.

This work uses a memetic algorithm, a hybrid GA and SA technique, to unmix spectral data. This memetic algorithm is presented here, along with a description of the genome representation and a constraint-preserving crossover method. The algorithm uses a fitness function for comparing reconstructed spectra against observed spectra that allows for domain-knowledge to be incorporated (specifically, expert knowledge about the likely presence of particular end-member types can be used as a modifier of the fitness evaluation of a solution).

The effect of band reduction and noise on the effectiveness of these unmixing algorithms is analysed.

The use of a fitness function as a measure of unmixing accuracy offers a mechanism to autonomously identify the presence of new composite end-member types. An algorithm to detect new end-members is presented with examples of its results.

The detected composite end-members are classified by an expert system (in this work, the memetic algorithm also acts as a domain-expert) to produce a final classification of observations.

The automatic composite end-member detection algorithm allows classification of observations into an arbitrary number of composite end-member groups. The data produced by these unmixing methods feeds directly into the mapping algorithm, described in Chapter 5.

## 4.2 End-Member Extraction

---

End-member extraction is a process to identify pure source signals from a mixture. The process of unmixing spectral data to identify the types and proportion of end-member constituents has been an active field of research for many years. As higher-resolution spectral data have become more readily available, accurate unmixing and substrate classification has become more feasible.

There have been many different techniques proposed to perform unmixing. Keshava (2003) provides a taxonomy of unmixing algorithms. Typically, there are three stages common to the process; dimension reduction, end-member determination, and inversion. The three criteria for categorising unmixing algorithms are; how mixed-pixel spectra are interpreted, how randomness is incorporated by the algorithm, and the objective function being maximised by the algorithm. The most common optimisation function is minimisation of squared error.

### 4.2.1 Linear Mixing Model

Some unmixing techniques rely on the presence of ‘pure pixels’, which are detected to serve as end-members. When no such pure pixels are available, the data is called highly mixed, and the problem may be known as an ‘unsupervised hyper-spectral data unmixing problem’. Pure pixels are often not present in spectral data. A common assumption behind many of the unmixing techniques is the linear mixing model (LMM), described in Section 3.3.2.

## 4.3 Formal Description

---

Given a sampled spectrum  $\bar{S}_{(x,y)}$  (see Equation 3.3.5), calculate the proportions,  $\bar{\omega}$ , of objects in benthic community  $\bar{C}$  contributing to  $\bar{S}_{(x,y)}$ . This requires a-priori knowledge of all possible objects and their spectral signatures (and variants). An equivalent solvable problem is to calculate the proportions,  $\bar{\omega}$ , of objects in an end-member spectral library  $\bar{T}$  (see Equation 3.3.6) which, when mixed as  $\hat{R}_{(x,y)}$ , give a most-similar spectrum to  $\bar{S}_{(x,y)}$ .

### 4.3.1 Reflectance and Reflectivity

‘Reflection coefficient’ refers to the fraction of an electric field reflected (see Equation 3.9.1). Reflectivity is a value that applies to thick reflecting objects, and is the square of the magnitude of the reflection coefficient. The unmixing process requires spectra to be represented as reflective spectra.

$$R_\lambda = (E_{u\lambda}/E_{d\lambda})^2, \quad \boxed{4.3.1}$$

where  $E_u$  is the downward-spectrometer measurement, and  $E_d$  is the upward-spectrometer measurement.

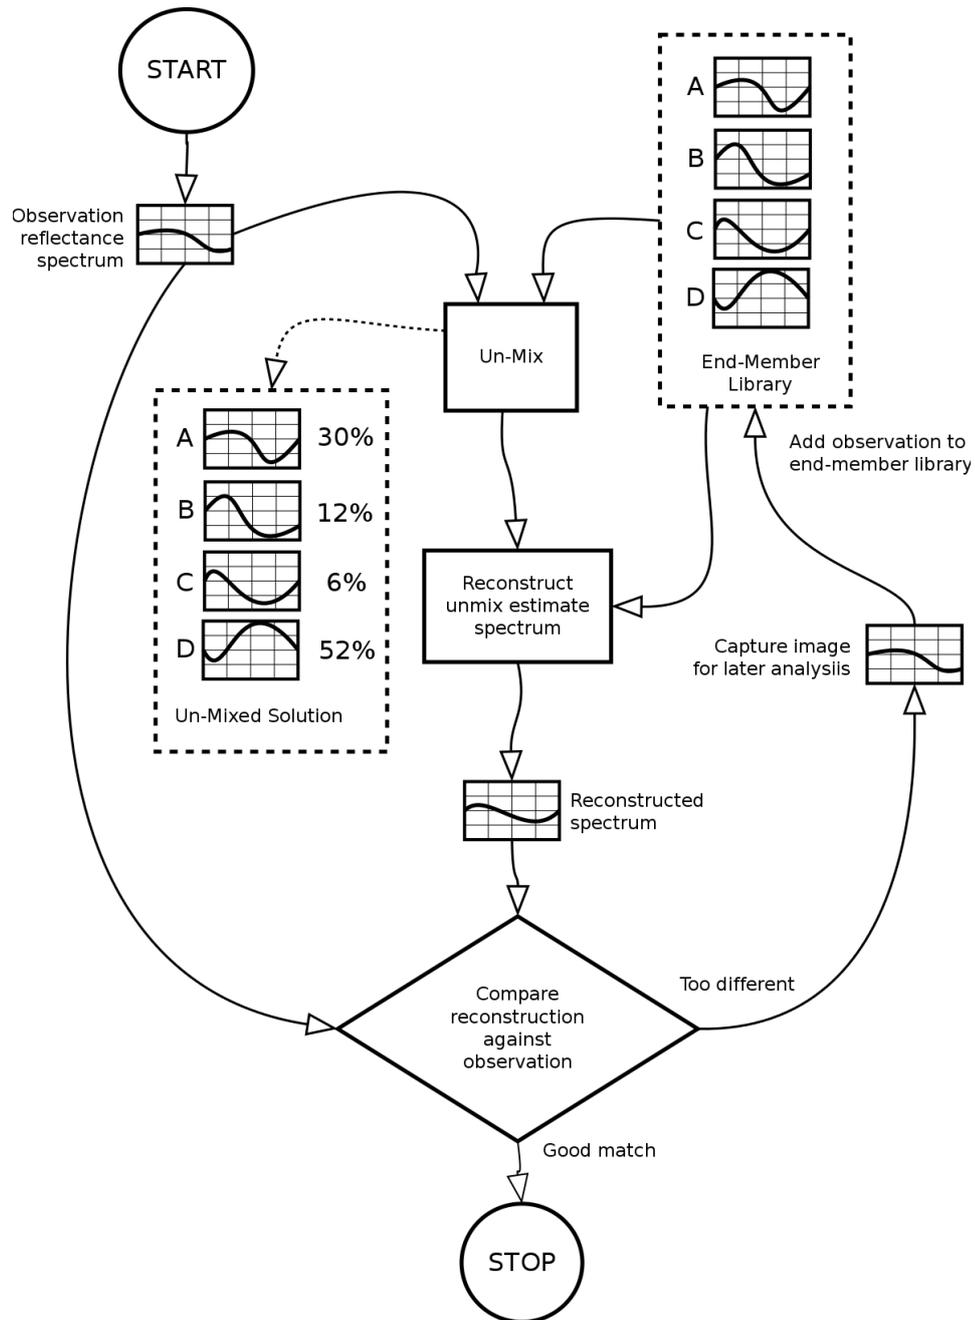
## 4.4 Techniques

---

A brute-force approach to unmixing, examining all values of  $\bar{\omega}$  (see Equation 3.3.7) is computationally prohibitive and rapidly becomes infeasible as the number of end-members increase. This method also cannot account for variation of spectra within an end-member type (for example, morphological or seasonal differences).

A mathematically exact solution using matrix algebra is straightforward, given low noise and a complete end-member library (i.e.,  $\bar{C} \subseteq \bar{T}$ ), but this technique suffers from inability to enforce the sum and unity constraints, particularly in the presence of noise, and also cannot account for end-member variability.

Quadratic programming techniques have frequently been used for spectral unmixing (e.g., Ambikapathi et al. (2010)). These techniques appear to work well, but the cited example only uses six end-members.



**Figure 4.1:** Unmixing and evaluating the result. The reflectance spectra arrive at ② and pass through the unmixing algorithm. An end-member library is used by the unmixing algorithm to produce unmixed estimates for the observations. These estimates are re-mixed using the end-member library and the resultant spectra are compared against the original reflectance spectra. The least-fit observation is added to the end-member library and the process repeats until the halting condition is met.

A discussion of some spectral unmixing techniques using library end-member templates is given in Meng and Ramsden (2000). Other curve fitting techniques have been attempted Jackson et al. (1991).

### 4.4.1 Reconstruction of Spectra

Given an unmixing estimate  $\bar{\omega}$ , Equation 3.3.7 shows the reconstructed mix of  $\bar{T}$ . Although this reconstructed spectrum,  $\widehat{R}_{(x,y)}$  may be close to the sampled spectrum, the calculated proportions of objects may differ from the actual values. The end-member library  $\bar{T}$  may be a subset or a superset of  $\bar{C}$ ; end-members may be combinations of objects. The variation of the objects' spectra, and measurement noise, may be such that the optimum proportions calculated differ from the actual proportions.

If  $\bar{S}_{(x,y)} = \widehat{R}_{(x,y)}$  then we have an optimal, but not necessarily correct, unmix. There is the possibility that combinations of end-members in  $\bar{T}$  can produce a spectrum most similar to  $\bar{S}_{(x,y)}$ , and thus indistinguishable from an exact unmixing. This is more likely when the number of bands is small.

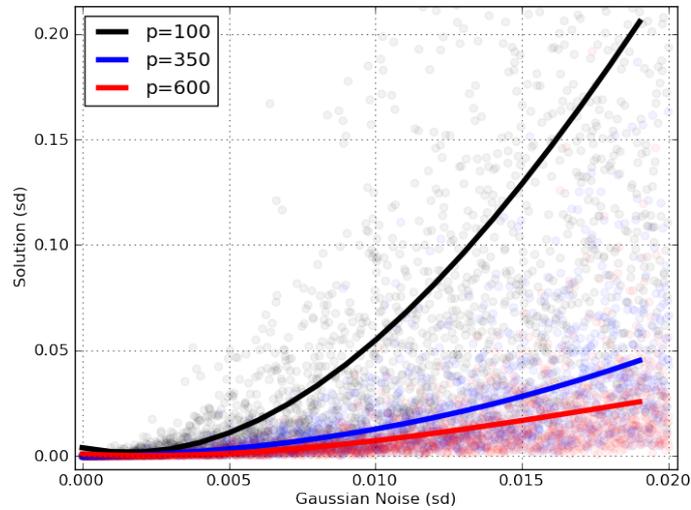
## 4.5 Fitness Criteria

---

These algorithmic approaches to unmixing use a measure of 'fitness' to compare the quality of successive solutions. The fitness is a measure of the similarity of two spectra, calculated by summing the squares of the differences between each spectrum at each wavelength and normalising by dividing by the number of bands, so that the fitness is independent of the bands in the observation (see Equation 4.5.1).

$$F_{(x,y)} = \sum_{\lambda=min}^{max} \left( \widehat{R}_{(x,y)_\lambda} - \bar{S}_{(x,y)_\lambda} \right)^n / B, \quad (4.5.1)$$

where  $min, max$  represent the range of wavelengths over which the fitness is evaluated, and  $B$  is the number of bands. Typically,  $n$  is set to 2, but this work explores higher cardinality to prioritise fits to peaks in spectra. Fitness approaches zero as the mixed solution more closely matches the sampled spectrum. A normalised fitness value, independent of the number of bands, is shown in Equation 4.5.2



**Figure 4.2:** Separate runs of the linear-algebra unmixing algorithm using varied number of bands ( $p$ ) with values (100,350,600). The lines are quadratic best fit for each  $p$ -value. This clearly shows the dependence of a linear-algebra solution on large  $p$ -values to achieve good results.

$$normalF_{(x,y)} = F_{(x,y)} / count(\lambda) \quad (4.5.2)$$

## 4.6 Unmixing Software

---

The author developed Python software for the evaluation of genetic algorithms (GAs), simulated annealing (SA), linear algebra (LA) and hybrid solutions to unmixing. The results are displayed graphically with plots generated using `matplotlib` (see <http://www.matplotlib.org>).

### 4.6.1 Methods

The software can read end-member spectral data from spreadsheet input, but for discussion purposes, random spectra are used here. The system supports an unlimited end-member library size. The number of spectral bands is configurable, and the wavelength range variable. The effect of band size reduction, end-member library size, run-time and noise are discussed below.

An ‘observation’ is computed by randomly combining end-member library spectra as per the linear mixing model (LMM), where the proportions of all contributing

end-members sum to 1, and all have a non-negative contribution. These proportions are then used to generate a spectral ‘observation’, by summing at each sampled wavelength the proportional amount of each contributing end-member spectrum, thus producing an observation spectrum. This spectrum may then be further modified by the addition of Gaussian noise, giving a mixed spectrum ready for testing the unmixing algorithms.

The software passes this observation spectrum to the unmixing algorithm and attempts to recover the original proportions of end-members used to generate it. Only the observation spectrum and end-member spectra are available to the unmixing algorithm; there is no information about the original mix proportions of the end-members. The presence, or not, of any individual end-member in the observation is unknown, as is the contribution of noise to the test spectrum.

### 4.6.2 Presentation of Results

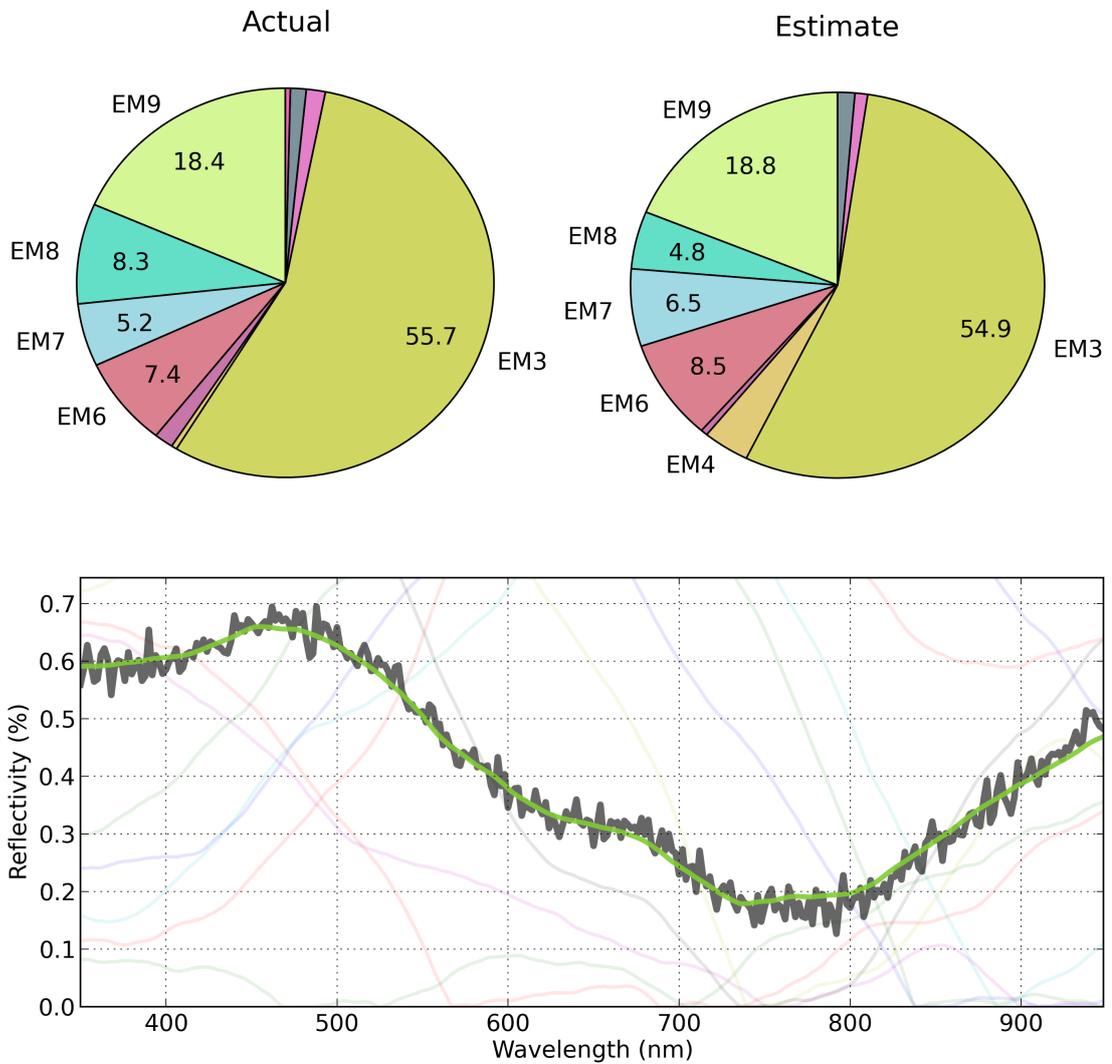
Results of the unmixing calculations are presented in two forms: pie charts and spectral graphs (see Figure 4.3). The pie charts show the proportion of each of the end-members (labelled around the perimeter  $EM_n$ , with the numerical proportion displayed inside each significant slice). Slices of negligible size are not labelled. The contributions of all proportions fill the chart, as expected. The colours allow comparison between pie charts to evaluate the quality of results.

The chart labelled ‘Actual’ is the randomly generated observation. The observation spectrum is plotted as the black or red line in the bottom graph in the figure (in this example, there is significant noise present). The original end-members’ spectra are also shown as faint lines in the background of the graph.

The unmixing algorithm attempts to find an optimal mixture of end-member proportions, such that expression of those proportions as a spectrum (the green line) is minimally different from the observation spectrum (the black line).

The pie chart labelled ‘Estimate’ is a solution generated by the unmixing algorithm. A perfect solution would have identical proportions of end-members to the ‘Actual’ observation, and in the absence of noise the green line would be coincident with the black line. The sizes of the pie slices give a good indication of the closeness of the estimate to the observation, as does the accuracy with which the green line overlaps the black line. Internally, a fitness value - the sum of the squares of the distances between observation and unmixed solution over all wavelengths (see Equation 4.5.1) - is used as an indication of the match.

In general, an observation in the real world will probably have a reasonably



**Figure 4.3:** Outputs from the unmixing software. The estimate shown here is early, but has already identified the major constituents. The green reconstructed spectrum is a reasonable match for the noisy ( $\sigma = 0.02$ ) observation spectrum.

low number of end-members present because of the small area covered by the spectrometer visibility cone. Consequently, the major components of a solution are considered the most important indicator of unmixing accuracy.

### 4.7 The Unmixing Algorithm

---

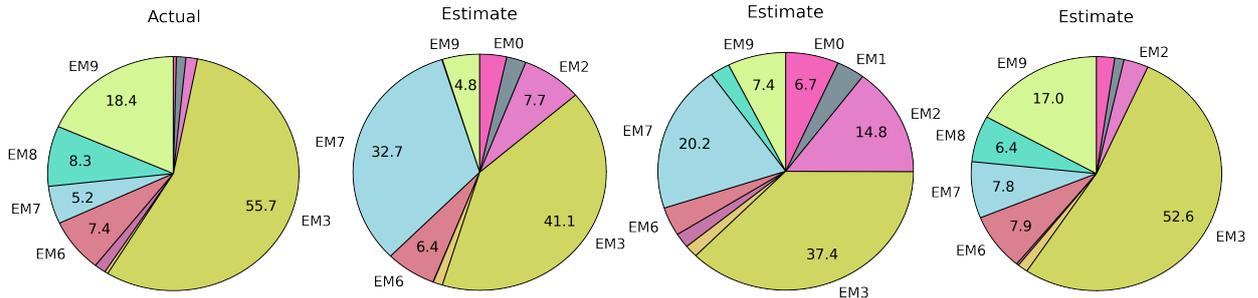
Genetic algorithms (GAs) are a well-known class of search heuristics that are inspired by the process of natural evolution. They are capable of generating solutions to high-dimensional optimisation and search problems that are intractable by exhaustive searches. Whitley (1994) provides a comprehensive overview of GAs and their operation. Farzam et al. (2008) have previously presented work on the use of GAs for unmixing hyper-spectral image data, and addition of new end-member types using residual analysis of unmixing.

Simulated annealing (SA) algorithms are also well-known heuristics for optimisation problems (Kirkpatrick, 1984). The method involves controlled reduction of randomness in the selection of candidate solutions in a large problem-space, akin to the cooling of a material to increase the size of its crystals and reduce defects. Bateson et al. (2000) have used simulated annealing algorithms to detect end-member bundles in spectral mixture analysis.

In this work, a hybrid GA+SA algorithm has been implemented to simultaneously evolve and anneal solutions to the unmixing problem. This algorithm combines elements of GA and SA into a single heuristic, incorporating unusual implementations of mutation and crossover, that has proven practical for signal mixture analysis (converging to good solutions in reasonable time) and robust in the presence of noise.

Alves and Carbonari (2008) have previously explored an alternate technique using a hybrid combination of GA and SA algorithms. They claim this method offers “considerable savings in analysis time, and allows the determination of good fits without well-defined initial values”. Their work involved fitting perturbed angular correlation spectra. Memetic algorithms are an active field of research in evolutionary computing, and the algorithm used in this work operates in a similar fashion.

A population of candidate solutions (phenotypes) is created. Population size is small: typically 20-50 individuals. Each phenotype represents a solution as a group of chromosomes (genotype) that are altered by recombination/crossover and mutation. Traditional implementations of GA represent the genotype as a binary string, but in this implementation, the genotype consists of genes expressed as



**Figure 4.4:** Progressive solutions from the unmixing algorithm performed on a random observation with 10 end-members using 256 bands and  $\sigma = 0.02$ . Shown are optimal solutions from generations 0, 1 and 181. Original end-member spectra were available for the unmixing estimates. Note the similarity in proportions of end-members between the observation and the final estimate.

floating-point numbers, one gene for each end-member type. A gene represents the expression (proportion) of a single end-member.

### Recombination/Crossover

Another unusual addition in this GA implementation is that Equation 3.3.3 (all phenotypes' genotypes sum to one) and Equation 3.3.4 (no genotypes are negative) are enforced. This occurs during recombination (crossover) through a normalisation process.

Each member of the population contains a random genome. The genome of each candidate solution consists of multiple genes, one for each end-member, representing the proportion of contribution of the associated end-member. Each gene consists of a floating-point value representing the fractional contribution of that gene towards the spectrum encoded by the genome. The sum of the contributions of all genes in the genome is equal to 1.

Several crossover/recombination methods are used to select genes from parents when creating children. Parent pairs are selected randomly from the population, and every individual in the population has children. Multiple children are created from each parent pair, using a variety of methods. Firstly a uniform crossover where each child contains a random mix of the parents' genes (0.5 probability for each gene) is used. Secondly, a randomised mix crossover, where for all genes a random proportion is used from one parent's gene, and the rest from the other parent's gene. In both cases, mutation is applied to the genes of the children by

## 4.7. THE UNMIXING ALGORITHM

---

adding or subtracting a small amount to the gene's value. The final genome is then normalised to ensure that the total contribution of gene proportions is equal to 1.

Multiple children are created for each crossover/recombination method. An additional 'walk-in' child with random genes is added, akin to the entrance of a member from an isolated population.

Finally, SA is implemented by iterating each member of the population and generating a 'child' with a random offset applied to each gene. The maximum offset is 0.1, and this is designed to allow the system to escape from local minima. The implementation of the SA is very simple. The use of a (hybrid) SA as an adjunct to the GA may be questioned, but at worst it does not adversely affect the search. In fact, it is likely to improve the search, as it provides a finer detail search of the problem space. Both GA and SA were tested in isolation, and both provided effective search algorithms for optimizing the unmixing. Combination of the two (GA+SA) has previously been done by other researchers, as referenced in the work by Alves and Carbonari (2008), and was found to be very effective.

As referred to by Bioucas-Dias et al. (2012)

The population is then sorted according to fitness by expressing the genome as a spectrum and calculating the fitness by comparison with the observation. Finally, the population is culled to the constant population size by removing (discarding) the least-fit individuals.

This technique produces a reasonable unmix solution for most observations in a short run-time, given a large number of end-members.

### Terminating Condition

There are two terminating conditions. One is to halt the algorithm after a set number of iterations. Good solutions appear in under 100 generations. The second terminating condition is to continue the algorithm until a certain number of generations have elapsed since finding an improved solution. This tends to find better solutions but takes longer to run, particularly when there are large numbers of end-members or in the presence of significant noise. In both cases, the fittest member of the population is the solution to the unmix.

### Factors Affecting Accuracy

Reduction in the number of bands in each observation decreases the ability to differentiate between similar end-member types, and increases the likelihood that multiple end-member combinations can appear similar. Reduction in the number of generations, or time available for performing the memetic algorithm search can result in sub-optimal solutions.

## 4.8 Software Testing

---

The software can read end-member spectra from an input spreadsheet, and augment these with randomly generated spectra. For the discussion here, only randomly generated spectra are used. The software handles any number of end-members, and has been tested with a range from 1 to 100. An inordinately large number of end-members is detrimental to the algorithm's speed. Most of the tests and validation performed in this section use from 8 to 12 end-members, as that range appears to be a reasonable size for real-world use. These end-member spectra represent pure instances of each end-member.

A large number of spectral samples representing observation spectra are generated by randomly mixing end-member spectra from this library. Gaussian noise is added to the observation spectra to simulate variation within end-members and instrument noise during observation.

## 4.9 Autonomous End-Member Detection

---

Without a complete end-member library providing spectra of all end-member types and variations, the ability of unmixing algorithms to calculate accurate solutions is hampered. Representative libraries of spectra of end-members are available, but these generally provide just a single spectrum for each end-member type, and tend to be instrument-dependent in wavelength range.

Blind unmixing is the technique of recovering end-members from a mixture, without using a library of end-member spectra. This approach is attractive as it allows previously unknown end-member types and variations to be classified through the process of automatic addition of newly-discovered spectral classes to the end-member library.

This section describes such an autonomous end-member detection algorithm.

## 4.9. AUTONOMOUS END-MEMBER DETECTION

---

As noted in Section 4.5, the fitness function (see Equation 4.5.1) provides a measure of the accuracy of an unmixed solution through comparison of the similarity between an observed spectrum and a reconstructed spectrum generated from an end-member library of spectra. This fitness value is independent of the number of bands in the observation. It is typically calculated as the sum of the squares of the difference between the two spectra at all wavelength bands, divided by the number of bands, and it approaches zero as the compared spectra become more similar. Alternate versions of the fitness algorithm can increase the cardinality of the exponent, which has the effect of treating deviations from transient spikes in spectral variation as being more costly in fitness.

### Quality of Unmixing

The end-member library  $\bar{T}$  contains representative spectra of the end-members in the community  $\bar{C}$ . The unmixing of a sample,  $\bar{S}_{(x,y)}$ , produces proportions of end-members in  $\bar{T}$ . If there are end-members in  $\bar{S}_{(x,y)}$  which are not present in  $\bar{T}$  then the proportional contribution of spectral data from these end-members will be included in  $\bar{S}_{(x,y)}$  and the unmixing algorithm(s) will be less able to calculate accurate unmix solutions. In this case, the fitness of the solution,  $F_{(x,y)}$ , will be worse than would be the case if all end-members were present in the end-member library.

### Spectral Variability

The representation of an end-member type by a single spectrum in the end-member library does not account for spectral variance within that end-member type. The spectrum of a young plant can vary significantly from an old or decaying plant. Seasonal differences in plant morphology and reflectance may be significant. Consequently, the unmixing of these variations using the single end-member spectrum will be less likely to find a good fit. Traditionally, these variations have not been problematic because of the large area encompassed by observations, and the effective averaging of all variations of each end-member type when observing a larger area. In these situations, a typical representative spectrum can be used.

Mapping on a fine-scale, close to the substrate, means that the effects of variation within an end-member type are likely to affect the unmixing accuracy significantly. To avoid the problem of variance within an end-member type, this system allows 'sufficiently different' spectra to be added to the end-member library. Detection of

these new composite end-member ‘types’ can be achieved by analysis of the fitness of unmixed solutions over a range of observations.

When the fitness is poor, the end-member library may be considered a poor representation of the types of end-members present in the contributing observation, and new composite end-member types can be added to the library to improve unmixing accuracy.

### 4.9.1 Validation of Methodology

The algorithm was tested by generating a large number of random mixes of an end-member library, and unmixing the resultant spectra. Randomised end-member spectra were used for the tests. These end-members were randomly mixed, and noise was added, to produce observation spectra. Three of these observations from a test-run are shown in Figure 4.8. The algorithm attempts to recover the original end-member spectra by classifying observations into a mix of composite end-members, detected by the algorithm. The small subset of composite end-members is presented to an expert system for analysis and representation as mixes of original end-members.

The random mixing iterates through the end-members in random order, assigning a randomised amount of the 100% mix to each. The process is biased, assigning 90 % of a randomised amount of the remaining proportion to each end-member. On average, this effectively gives each end-member 45 % of the remaining proportion of the mix. The earlier assigned end-members tend to have larger proportions, and the mix has (on average) a small number of larger-proportioned end-members. The leftover portion after this process is added to a random end-member, so that the unity constraint is maintained. This mixing procedure mirrors the real world where one or two end-members are likely to be dominant in most observations.

Up to 32 end-member spectra were tested, and up to 4096 randomised observations using those end-members. The number of bands tested ranged from 16 to 2048. Generally, band reduction had little significant effect on the accuracy of unmixing (see Figure 4.9), and 64 bands became the default choice because it offers good algorithm-speed and classification results.

#### Seeding the Library

The algorithm requires a non-empty end-member library against which unmixing can proceed. Several techniques have been tested for generating the initial seed

## 4.9. AUTONOMOUS END-MEMBER DETECTION

---

spectrum:

- Generate a random spectrum but discard it once an alternate end-member is found by algorithm.
- Choose a random observation's spectrum, and optionally discard this first-generation end-member once a worst-fit candidate has been chosen.
- Calculate an average spectrum over all observations, and then find the observation that is least different (using the fitness function) from the average. Use this as the first end-member.

All of these techniques produce excellent results, but the last (use the observation closest to the averaged spectrum as the seeding end-member) is the default, because it does not require a complete iteration of unmixing, only to discard the end-member used for the unmixing.

### Single End-Member

Given a single end-member in the library, all unmixed solutions using that end-member will be identical. The single end-member will have 100% expression in every solution. The fitness of these solutions will vary; indeed, they may all be poor. However, if the least-fit solution's spectrum is added to the end-member library as a new composite end-member, then a subsequent unmixing of that least-fit observation will contain 100% of the newly-added composite end-member, and the fitness of the unmix will be optimal (zero).

The addition of the new composite end-member to the library provides opportunity for fitter unmix solutions to be detected for the other observations. At worst, existing observations will have the same fitness. The previously iteration's least-fit observation will now have a fitness of zero and no longer participates in unmixing due to *alpha pruning* (see Section 4.9.1). In practise, the addition of the new composite end-member contributes to an improvement in fitness for many observations. By iterating this process (unmixing, adding the least-fit solution as a new composite end-member and repeating), the fitness of all observations is improved (see Table 4.1).

### Iterating to find End-Members

An iteration consists of unmixing all observation samples, reconstructing the spectra these unmixed solutions represent, comparing the reconstructed spectra

to the observed spectra, finding the least-fit solution and adding it to the end-member library as a new composite end-member. The iteration is then complete, and another may be performed.

The newly detected composite end-member will most likely be a combination of one or more currently unrepresented object types. A camera image of the area of the observation is associated with the composite end-member to assist in later expert re-classification of the composite end-member in terms of pure end-member types.

### Threshold

When the fitness function for an unmix estimate returns a value beyond a designated threshold, then the contribution of one or more new end-member types, or a significant variation within an end-member is inferred. Even though the observation is almost certainly not representative of a pure instance of a single new end-member type, but instead a composite, it may be added to the end-member library to improve the unmixing accuracy. The process of adding the new composite end-member type consists of adding the observation spectrum to the library, and recording a concurrent camera-image of the seafloor at which point the spectral observation is collected. The composite end-member is of currently unknown type, and probably, as noted, a mix of known and unknown end-member types, but is considered sufficiently distinct from all other end-members or combinations thereof. The constituent types of these new composite end-members can be determined post-mission by manual analysis of the associated image by an expert.

The newly added end-member may represent a composite object, a sufficiently different variation of an already existing end-member, or a completely new object. The location with the worst  $F_{(x,y)}$  is considered the purest example of the new object(s) or combination of objects, or more to the point the observation which is least explained by the existing end-member library.

### Halting Conditions

The automatic end-member detection algorithm can be halted after any iteration has completed. The halting-condition(s) of the algorithm may be varied; each has benefits and disadvantages...

- When a pre-determined number of end-members have been added. There is no disadvantage to adding too many composite end-members, as the final

## 4.9. AUTONOMOUS END-MEMBER DETECTION

---

stage of the algorithm reduces automatically added composite end-members to a minimal true end-member set.

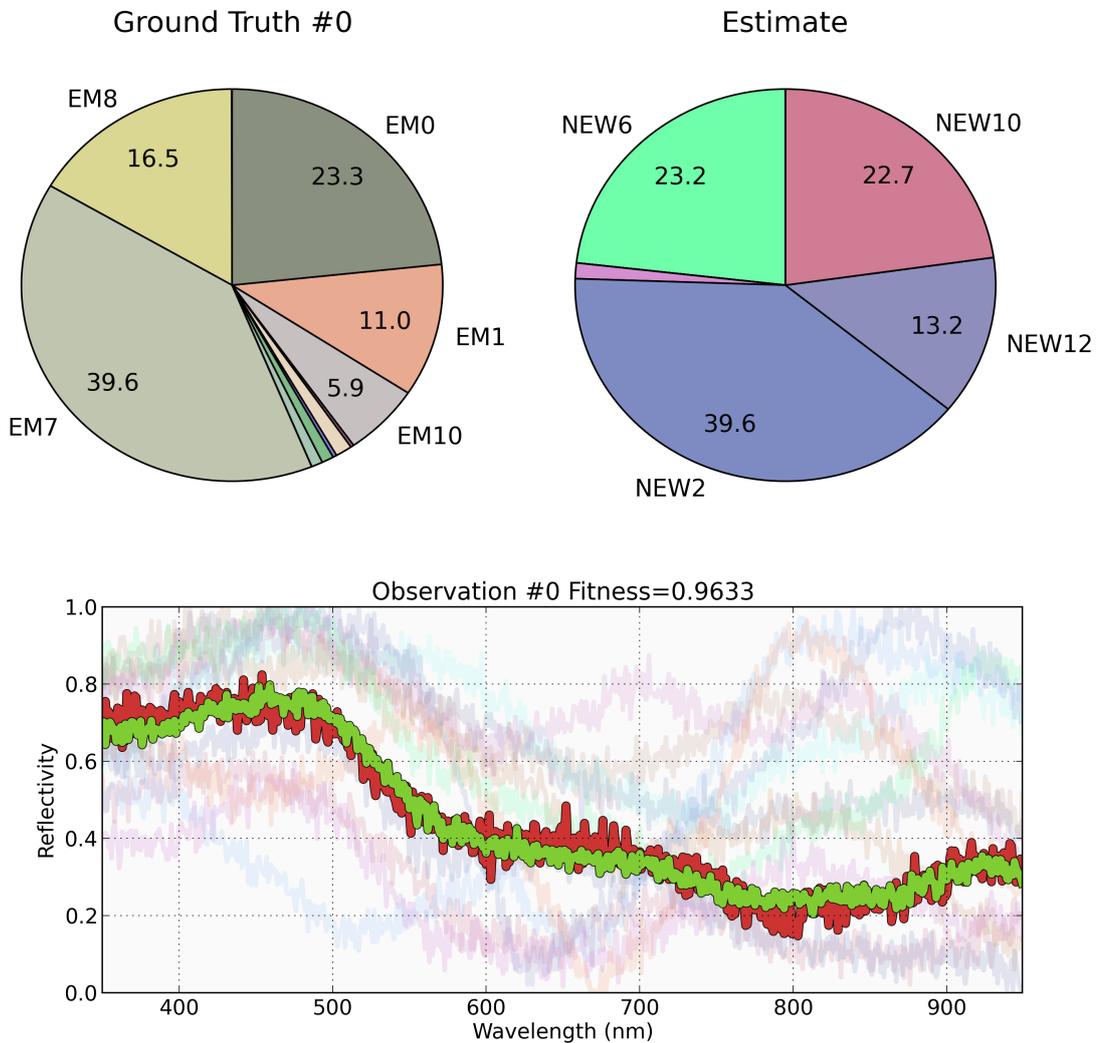
- When the fitness values for all of the unmixed solutions are below a threshold. This technique can cause many single-case composite end-members to be added which offer little improvement to the other observations' unmixing. This method is not very good for noisy data, as the noise increases the fitness values, and therefore the number of iterations.
- When the average fitness value of all unmixed solutions is below a threshold. This can allow some unmix solutions to be reasonably poor, but on average the majority of solutions are good. This allows an early termination with good solutions. This method also suffers from noise increasing iteration count.
- When the relative improvement in average fitness of all solutions between iterations falls below a threshold.
- When the addition of a new composite end-member does not improve the fitness of any unmix solution other than its own.

### Unmixing the End-members

The composite end-members are mixes of pure end-members, whereas the pure end-members themselves are each a representation of a single object type. For validation only, conversion of composite end-member mixes to pure end-member fractional abundances is straightforward; each composite end-member's spectrum was unmixed by the algorithm using the pure end-member library spectra. This technique effectively uses the unmixing algorithm in the place of a human. This step produces a composite end-member genome, representing the fraction of each pure end-member present in the composite end-member (see Figure 4.5).

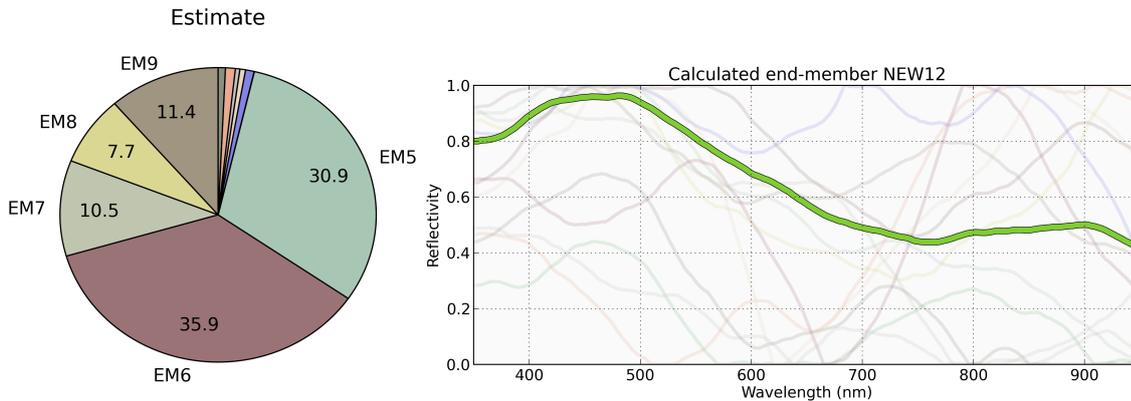
Table 4.1 shows the improvement in fitness as least-fit observations are added as new composite end-members and the observations are unmixed with the new end-member library. In this table, the rows represent observations, and the columns are iterations of the algorithm. Observations that have been added as new composite end-members (the least-fit after an iteration) are shown in bold. The first observation was used as a seed (this is discarded after the first iteration). The fitness of the observations generally improves from column to column, as does the average fitness of all observations.

The algorithm performs an unmix on every observation. With just a single end-member (EM0) in the end-member library, every unmix estimate consists of 100%



**Figure 4.5:** Ground truth and unmixed estimate. The pie chart proportions do not match because they reference different end-member concepts; the ‘Ground Truth’ expresses proportions of pure end-members, and the ‘Estimate’ those of composite end-members. However, the consequent spectra (red/green lines on graph) are coincident, indicating an excellent unmix solution, and that the end-member libraries are analogous. The automatically detected composite end-members are each a mix of pure end-members, so they must first be unmixed in terms of the pure end-member library to represent them in proportions of those spectra. An example unmix, showing a single composite end-member in terms of pure end-members, is shown in Figure 4.6.

## 4.9. AUTONOMOUS END-MEMBER DETECTION



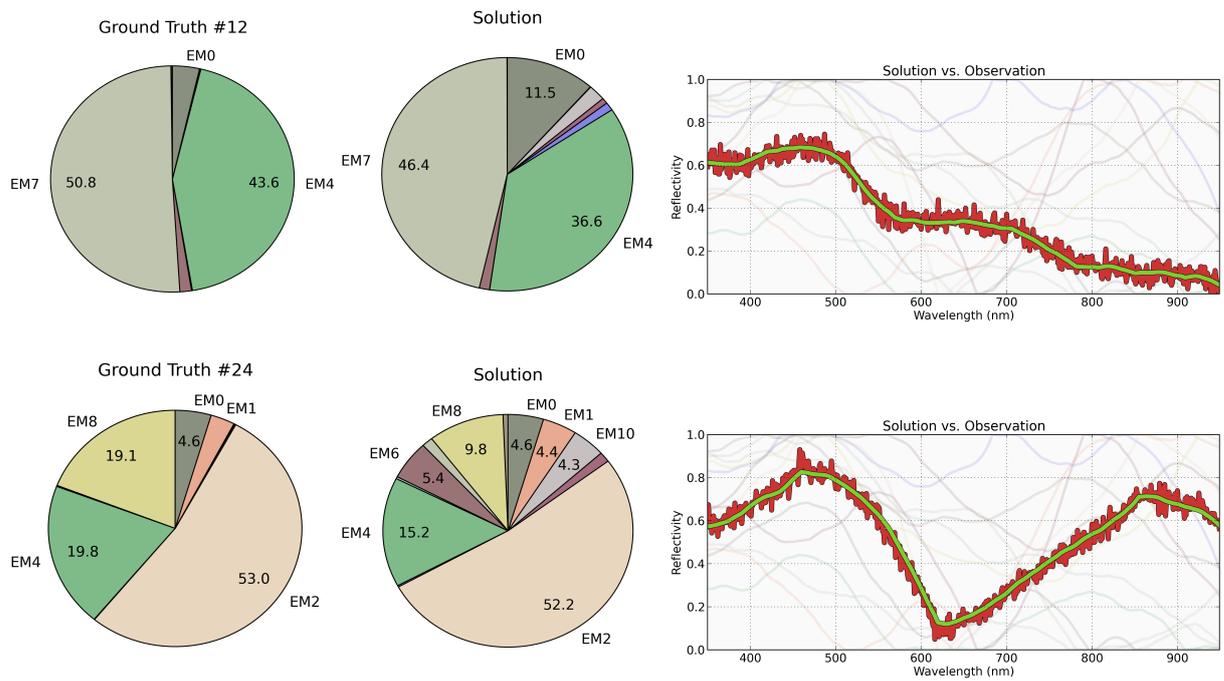
**Figure 4.6:** Composition of an automatic composite end-member in terms of original end-member library. The ‘expert’ used was the unmixing algorithm itself. In real usage, a human expert would do this analysis. The original composite end-member spectrum (red) is obscured by the unmixed spectrum (green). In this case, a near-perfect fit. The original composite end-member is a representation of pure end-members as shown in the pie chart.

of that end-member. The fitness of each unmix is determined (see 4.5.1) by reconstructing the spectrum represented by the unmix proportions, and comparing this reconstruction with the observation. If the least-fit observation’s fitness is worse than a threshold value, chosen in advance, then observation’s spectrum is added to the end-member library ( $NEW_n$ ), being a representation of the most-different spectrum off all of the observations. The observation which contributed its spectrum to the end-member library will unmix to 100% of this new end-member type, but other observations will incorporate fractional amounts of this new composite end-member in their unmix results and hopefully show improved fitness as a result. The algorithm is repeated, until a halting condition is reached (see Section 4.9.1).

### Speed of the Algorithm

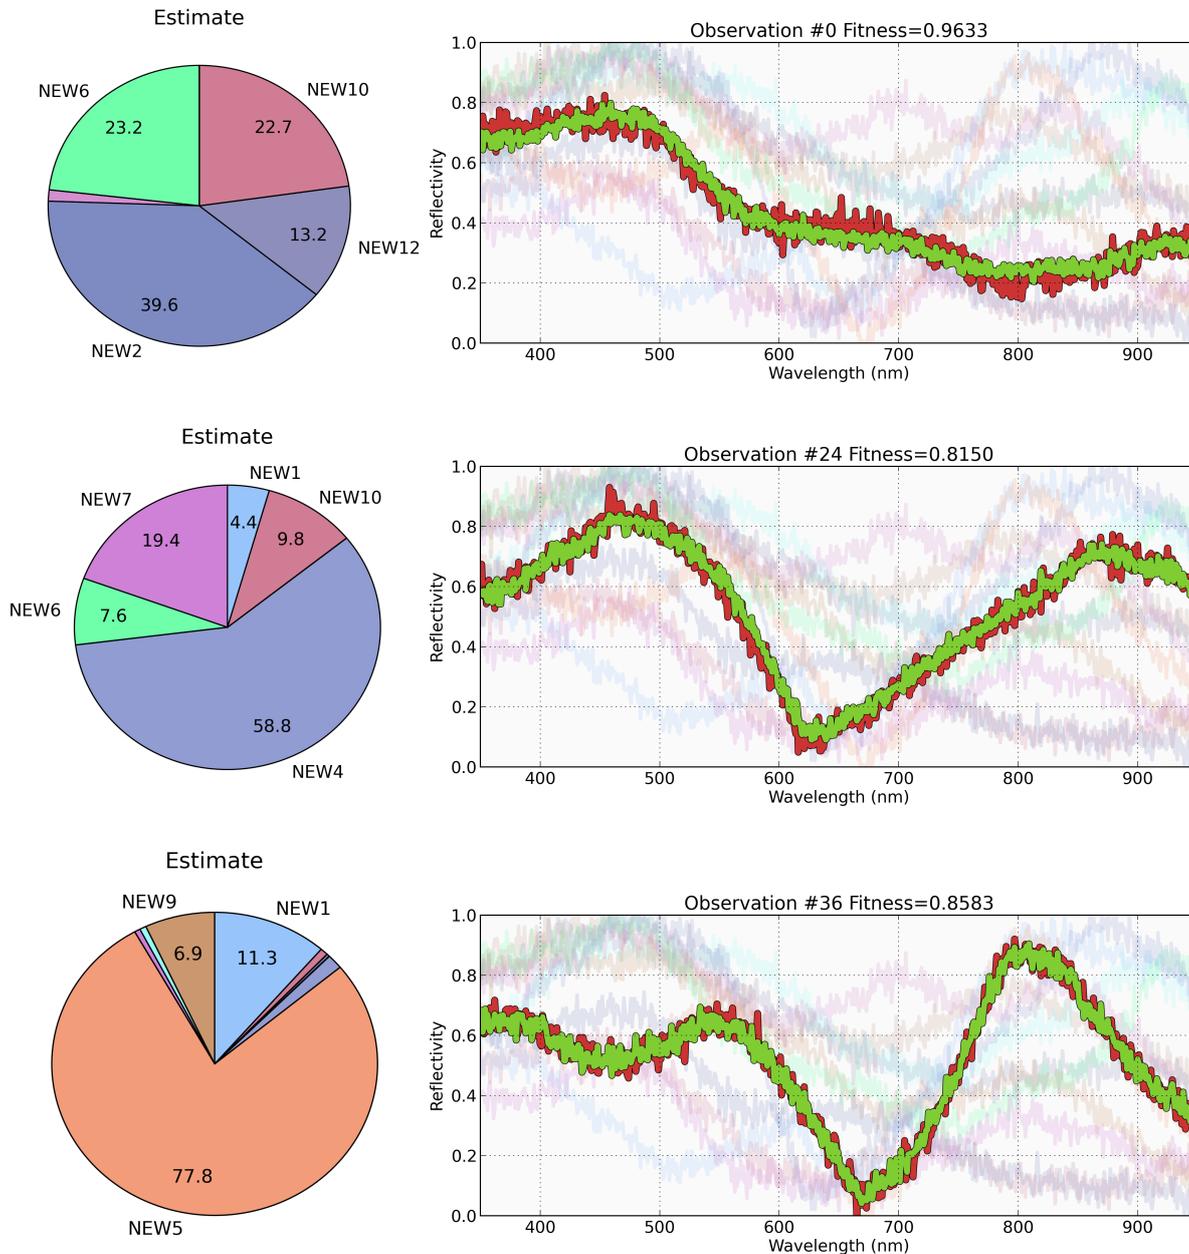
The speed of this unmix varies with the number of end-members, the number of bands, the number of observations, and the halting conditions of the algorithm (see Section 4.9.1). As increases in the number of end-members exponentially increase the size of the problem-space that the memetic algorithm searches for a best-fit solution, the addition of end-members significantly slows the algorithm. Increases in the other parameters have a much less dramatic effect on speed.

However, with unmixing computation taking many seconds per observation when



**Figure 4.7:** Pure representation of solution in terms of end-member spectra. The left pie chart shows the original observation’s mix of end-members (the ‘ground truth’). The right pie chart shows the observation’s unmix solution. The graph shows the solution expressed as a reconstructed spectrum. Note the lack of noise in the fit (green line), as a result of the use of original end-member spectra for creating the solution spectrum. Compare to the noisy fits in Figure 4.8.

## 4.9. AUTONOMOUS END-MEMBER DETECTION



**Figure 4.8:** Three observations (red) and solution (green) after the algorithm has added twelve composite end-members to its end-member library. The composite end-members are the faint lines in each graph, and represent composites of the actual end-members. The pie charts show the proportions of each composite end-member (NEW $n$ ) in the unmixed solution. The composite end-members are noisy, as they have been created using the noisy observations. This noise is removed when the composite end-members are themselves unmixed in terms of the pure end-member spectra (see Figure 4.7).

the number of end-members is 10, optimisations in the algorithm are required.

### 4.9.2 Pruning the Search

As described above, the algorithm recursively iterates through the observations, performing unmixing, to find the least-fit unmix solution of all the observations. The spectrum of the least-fit observation (i.e., corresponding to the least-fit solution) is added to the end-member library as a new composite end-member. This iteration through the observations requires a complete memetic unmix per-observation.

The memetic algorithm is computationally intensive. A full unmix solution with 10 end-members can take half a minute or more per observation on a 2.5 GHz MacBook Pro laptop. Recursively unmixing thousands of observations becomes impractical.

#### Alpha-Pruning

However, the goal of the recursive iteration is, as noted, to find the least-fit unmix solution of all the observations. After calculating the fitness of the first observation's unmix solution, if at any point during the memetic calculation the most fit member of the population is fitter than the least-fit solution previously found, there is no point completing the current unmix – for we already know that the current solution is fitter and therefore the final unmix of the current observation will be fitter, too. In this situation, the current unmix can be abandoned as it would not be chosen as a new 'worst' solution anyway.

This optimisation, termed *alpha pruning*, provides an order-of-magnitude speed increases in the iterative unmixing algorithm itself, as it bypasses the most computationally intensive part of the search – the memetic optimisation of the unmixing.

#### Beta-Pruning

There is a second level of fitness-based optimisation available.

Given an end-member library and a complete unmix of all observations, each observation's fitness is recorded with the observation. During the next iteration, unmixing with an added end-member, each observation would at worst unmix to the fitness obtained with the previous end-member library. This can be used to

## 4.9. AUTONOMOUS END-MEMBER DETECTION

---

improve the efficiency of the unmix, as each observation's previous fitness can be used as a cut-off to early-terminate the search for worst-fitness observations. In short, if any observation's fitness from the prior iteration of unmixing is better than the current iteration's worst fitness, then do not consider this observation at all.

This optimisation, termed *beta pruning* provides similar performance improvements as the *alpha pruning* described above.

### 4.9.3 Expert Review

At the conclusion of a mapping mission, a number of composite end-members may have been added to the end-member library. These composite end-members are in likelihood not pure representations of a single object type, but instead a mix of end-member types.

As each addition to the end-member library is accompanied with a camera image taken at the same time/location of the sample considered to be the most 'sufficiently different' to existing end-members, it is possible for a human expert to re-classify each member of the end-member library as a mix of defined actual end-member types.

Prior knowledge about the physical distribution of plants and other distinguishing features can be assist in classification. Information about the depth and location at which particular species of plants are found may be useful in differentiating between species with otherwise similar spectral signatures. Such information can also be used to exclude end-members from consideration.

This expert reclassification of the composite end-members is then used as an authoritative representation of the composition of each end-member. This reclassification, applied to all of the end-member library objects, produces a list of associations. These indicate the actual proportions and types of end-members that compose each end-member library type. The observations' proportions in terms of composite end-members are converted into proportions in terms of pure end-members, and the final classification maps are produced.

### 4.9.4 Visual Representation of Fitness

The unmixing fitness is an input into the mapping algorithm (see Chapter 5). The fitness contributes to the confidence layer produced by the algorithm, and provides a visible map of the areas that have been poorly unmixed.

## 4.10 Analysis of Unmixing Results

---

Various factors influence the effectiveness of an unmixing algorithm. Optimal unmixing can be achieved using linear algebra, but the presence of noise and unknown end-member types makes this an impractical technique for real-world classification because of the non-enforceability of non-negativity and unity constraints.

### 4.10.1 Noise and Sampling

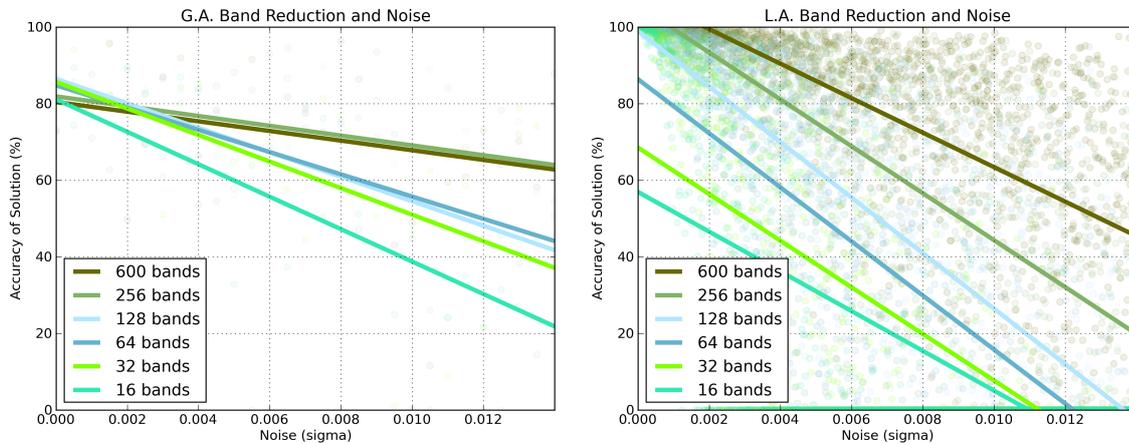
Gaussian noise ( $\sigma$ ) was added to each sample in  $\widehat{S}$  before performing the unmixing, to test the effect of measurement noise on sample data. In these tests,  $n$  represents the number of end-members,  $m$  represents the maximum number of spectral bands available, and  $p$  represents the number of spectral bands in the observations. The algorithm is exact where  $\sigma = 0$  and  $p \geq n$ . Increasing  $p$  improves the accuracy of the result when noise is present.

The effect of noise was determined by plotting the standard deviation of the error in proportions of solutions for noise  $0 \leq \sigma \leq 0.01$ . These data were plotted for  $p = 100$ ,  $p = 350$  and  $p = 600$  (see Figure 4.2). In the presence of noise, increasing  $p$  improves the accuracy of the unmixing. The addition of noise can lead to solutions where  $\widehat{w}_i < 0$ . In the context of benthic classification, such values are meaningless and replaced with  $\widehat{w}_i = 0$ .

Figure 4.9 shows test-bed software output showing sample unmixing for  $p = m$  with noisy data. In the graph, the sampled spectrum with added noise is shown in brown, and the reconstructed spectrum is green, with the end-members in grey. The histogram shows  $w_i$  (brown) against  $\widehat{w}_i$  (green). The negative value(s) for  $\widehat{w}_i$  are in this context meaningless, and replaced with 0

### Band Reduction

This analysis compares the accuracy of genetic algorithm unmixing against linear algebra unmixing. It varies the number of bands (wavelength groupings), using 16, 32, 64, 128, 256 and 600 bands. 4 and 8-band analysis was also performed, but results were generally poor and are omitted. A sub-sampling of the original 600-band end-member spectra was used to generate end-members with the required number of bands. The analysis also varies the noise present in the observed spectral signal, using Gaussian noise with  $0 \leq \sigma \leq 0.015$ .



**Figure 4.9:** Analysis of the effect of band reduction and noise on the accuracy of genetic and linear algebra unmixing. Tests were performed with six different band sizes (600, 256, 128, 64, 32, and 16), and a range of noise values. The algorithms (genetic algorithm on the left, and linear algebra on the right) were used to unmix spectra produced from randomly mixed sample end-members with added noise, and the solution proportions were used to generate a reconstructed spectrum. The accuracy of each solution was calculated and plotted as single coloured point on the graph. Linear best-fit lines are plotted for all bands. The response of the algorithms in the presence of noise and reduction of bands is apparent; the genetic algorithm is more robust in the presence of high noise levels and with low band resolution.

Another issue to consider is that light may not be present at some of the wavelengths recorded by the spectrometers. The proportional spectrum calculation will fail at these wavelengths. The fitness/comparison functions should therefore ignore these zero-light wavelengths.

To allow easy comparison, a hard-wired genome has been chosen.

The sum of proportions of all end-members is unity.

## 4.11 Summary

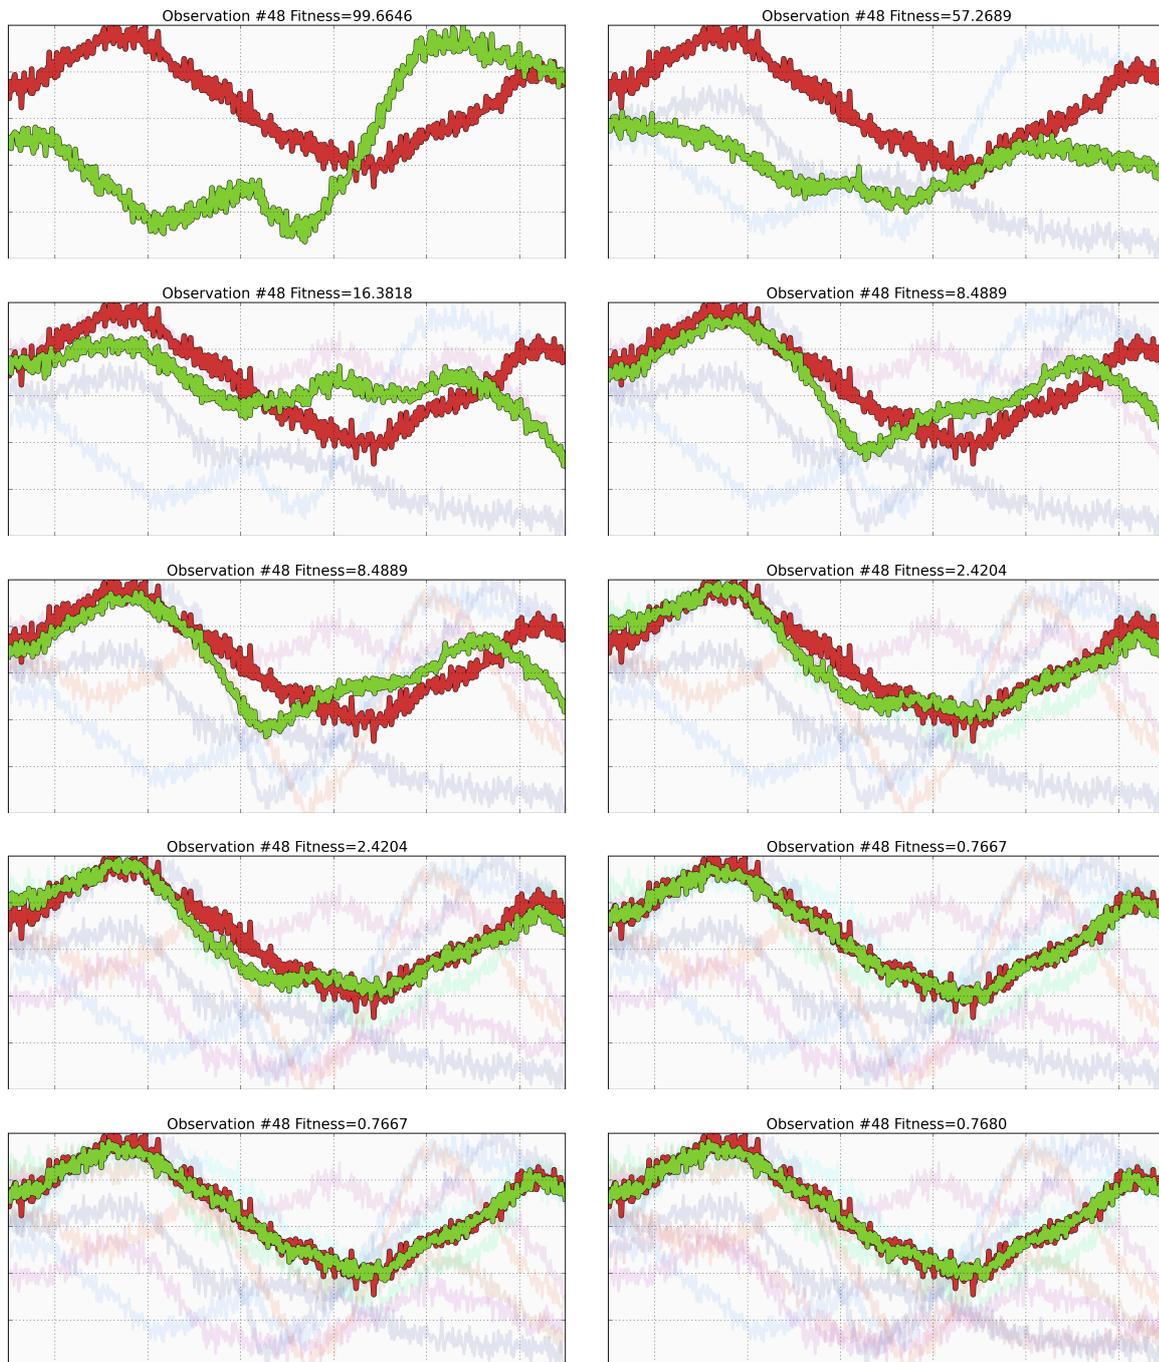
Many techniques have been developed for solving unmixing problems, some of which were explored in this chapter. The memetic (hybrid-genetic) algorithm presented here works well with spectral data, and allows for domain-knowledge to be incorporated to improve unmixing accuracy. The fitness function used for evaluating the memetic solutions' fitness can also be used to determine when

## CHAPTER 4. UNMIXING

---

solution-quality is poor. An iterative process can autonomously detect and add new composite end-members to the end-member library, and improve the fitness of unmixed solutions. The final end-member library (along with associated images) can be evaluated by experts to determine the end-member mix of the composite end-members, producing a final end-member library and complete unmixing of observations.

## 4.11. SUMMARY



**Figure 4.10:** Memetic best-fit solutions for example observation (#48) over successive iterations of the automatic composite end-member detection algorithm. The red line is the observation spectrum, and the green line is the reconstructed best-fit spectrum. Each iteration adds a new composite end-member (faint, in background). This improvement in fit happens for all observations at the same time. The minor worsening in fitness in the last iteration is because of the approximate solution returned by the memetic algorithm. More time spent analysing would return a better solution. See Table 4.1 for numerical fitness values.

## CHAPTER 4. UNMIXING

**Table 4.1:** Fitness for 12 iterations of 50 observations ( $\sigma = 0.03$ ). Bold values are those in each iteration with the worst fitness, and these spectra were added as new composite end-members for the next iteration. Observation 9 was the seeding end-member. Observation 48 is graphed in Figure 4.10.

| Obs.     | Iteration     |              |              |              |              |              |             |             |             |             |             |             |
|----------|---------------|--------------|--------------|--------------|--------------|--------------|-------------|-------------|-------------|-------------|-------------|-------------|
|          | 0             | 1            | 2            | 3            | 4            | 5            | 6           | 7           | 8           | 9           | 10          | 11          |
| 0        | 107.32        | 5.97         | 3.13         | 2.43         | 2.43         | 1.11         | 1.11        | 1.11        | 1.11        | 0.99        | 0.99        | 0.96        |
| 1        | 63.34         | 31.90        | 8.20         | 0.86         | 0.81         | 0.81         | 0.77        | 0.77        | 0.77        | 0.77        | 0.75        | 0.75        |
| 2        | 103.69        | 47.87        | 22.84        | 12.95        | <b>12.95</b> | 0.00         | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        |
| 3        | 59.18         | 39.42        | 10.80        | 1.37         | 1.35         | 1.33         | 1.32        | 1.32        | 1.28        | 1.29        | 0.91        | 0.91        |
| 4        | 60.46         | 29.67        | 11.43        | 6.33         | 1.15         | 1.07         | 1.07        | 1.07        | 0.75        | 0.75        | 0.69        | 0.70        |
| 5        | 74.99         | 5.17         | 5.17         | 5.17         | 5.17         | 5.17         | 1.58        | 1.58        | 1.58        | 0.93        | 0.93        | 0.93        |
| 6        | 126.04        | 12.06        | 2.44         | 2.44         | 2.44         | 2.21         | 2.21        | 2.21        | 1.83        | 1.83        | 1.83        | 1.49        |
| 7        | 89.27         | 17.37        | 2.39         | 1.32         | 1.30         | 0.86         | 0.86        | <b>0.86</b> | 0.71        | 0.65        | 0.64        | 0.63        |
| 8        | 124.80        | 27.79        | 7.15         | 6.99         | 6.73         | 6.47         | 6.47        | <b>6.47</b> | 0.00        | 0.00        | 0.00        | 0.00        |
| 9        | 0.00          | 0.00         | 0.00         | 0.00         | 0.00         | 0.00         | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        |
| 10       | 111.59        | 9.72         | 3.07         | 2.10         | 2.10         | 1.96         | 1.96        | 1.96        | 0.93        | 0.89        | 0.88        | 0.86        |
| 11       | 80.04         | 63.69        | <b>34.96</b> | 0.00         | 0.00         | 0.00         | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        |
| 12       | <b>136.23</b> | 0.00         | 0.00         | 0.00         | 0.00         | 0.00         | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        |
| 13       | 43.39         | 23.80        | 18.57        | <b>16.33</b> | 0.00         | 0.00         | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        |
| 14       | 43.41         | 3.06         | 1.72         | 1.72         | 1.72         | 1.40         | 1.40        | 1.40        | 1.40        | 1.44        | 1.42        | 1.41        |
| 15       | 2.28          | 1.84         | 1.79         | 1.78         | 1.78         | 1.68         | 1.67        | 1.67        | 1.67        | 1.66        | 1.66        | 1.66        |
| 16       | 115.30        | <b>75.87</b> | 0.00         | 0.00         | 0.00         | 0.00         | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        | 0.01        |
| 17       | 40.25         | 7.06         | 7.06         | 7.06         | 7.06         | 7.06         | 1.84        | 1.84        | 1.84        | 1.59        | 1.59        | 1.59        |
| 18       | 99.39         | 53.44        | 2.77         | 2.74         | 2.65         | 2.64         | 2.64        | 2.64        | 2.38        | 2.37        | 1.45        | 1.50        |
| 19       | 53.61         | 4.22         | 4.22         | 4.22         | 4.22         | 4.22         | 1.72        | 1.72        | 1.72        | 1.06        | 1.06        | 1.07        |
| 20       | 82.98         | 11.19        | 11.19        | 11.19        | 11.19        | <b>11.19</b> | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        |
| 21       | 15.63         | 10.22        | 2.34         | 1.14         | 1.12         | 1.10         | 1.10        | 1.10        | 1.09        | 1.09        | 0.89        | 0.89        |
| 22       | 124.03        | 9.22         | 3.38         | 3.28         | 3.27         | 3.07         | 3.07        | 3.07        | 1.23        | 1.18        | 1.18        | 1.17        |
| 23       | 83.73         | 34.93        | 20.98        | 10.34        | 10.34        | 1.57         | 1.40        | 1.40        | 1.40        | 1.39        | 1.39        | 1.39        |
| 24       | 55.79         | 24.31        | 17.05        | 1.51         | 1.51         | 1.51         | 0.82        | 0.82        | 0.82        | 0.81        | 0.81        | 0.82        |
| 25       | 110.41        | 66.80        | 12.30        | 8.15         | 8.15         | 5.10         | 5.10        | 1.36        | 1.22        | 1.22        | 1.22        | 0.97        |
| 26       | 32.73         | 3.90         | 3.21         | 1.07         | 1.07         | 1.07         | 0.95        | 0.95        | 0.95        | 0.77        | 0.76        | 0.77        |
| 27       | 105.32        | 63.72        | 4.46         | 3.87         | 3.65         | 3.02         | 3.02        | 3.02        | 2.89        | <b>2.89</b> | 0.00        | 0.00        |
| 28       | 110.65        | 1.89         | 1.63         | 1.52         | 1.52         | 1.39         | 1.36        | 1.36        | 1.36        | 0.94        | 0.94        | 0.94        |
| 29       | 127.42        | 51.32        | 8.01         | 7.25         | 7.25         | 4.28         | 4.28        | 4.03        | 2.09        | 2.09        | <b>2.09</b> | 0.00        |
| 30       | 119.29        | 7.59         | 7.59         | 7.59         | 7.59         | 7.59         | 3.17        | 3.17        | <b>3.17</b> | 0.00        | 0.00        | 0.00        |
| 31       | 75.06         | 48.29        | 8.77         | 1.07         | 0.95         | 0.77         | 0.77        | 0.74        | 0.74        | 0.74        | 0.67        | 0.66        |
| 32       | 77.14         | 42.86        | 10.93        | 1.06         | 0.99         | 0.99         | 0.99        | 0.99        | 0.88        | 0.88        | 0.81        | 0.80        |
| 33       | 77.20         | 8.29         | 2.68         | 1.49         | 1.49         | 1.46         | 1.20        | 1.20        | 1.20        | 0.93        | 0.93        | 0.93        |
| 34       | 76.40         | 33.00        | 5.67         | 3.10         | 3.10         | 1.79         | 1.79        | 0.74        | 0.73        | 0.73        | 0.71        | 0.69        |
| 35       | 45.85         | 6.63         | 6.57         | 4.55         | 2.41         | 2.41         | 1.38        | 1.38        | 1.38        | 1.10        | 1.13        | 1.10        |
| 36       | 37.29         | 18.16        | 13.26        | 11.43        | 0.89         | 0.89         | 0.89        | 0.89        | 0.85        | 0.86        | 0.86        | 0.86        |
| 37       | 90.98         | 43.98        | 23.65        | 10.81        | 10.81        | 1.25         | 1.21        | 1.21        | 1.21        | 1.22        | 1.21        | 1.21        |
| 38       | 89.26         | 32.61        | 18.57        | 8.14         | 8.14         | 1.39         | 1.13        | 1.13        | 1.13        | 1.09        | 1.09        | 1.09        |
| 39       | 17.38         | 7.09         | 3.64         | 1.26         | 0.77         | 0.77         | 0.76        | 0.75        | 0.75        | 0.75        | 0.75        | 0.75        |
| 40       | 73.48         | 51.02        | 23.01        | 1.33         | 1.21         | 1.07         | 1.08        | 1.07        | 1.06        | 1.06        | 0.98        | 0.98        |
| 41       | 36.90         | 23.15        | 6.13         | 4.34         | 4.34         | 2.88         | 2.88        | 0.88        | 0.88        | 0.88        | 0.88        | 0.88        |
| 42       | 71.97         | 4.26         | 3.90         | 3.60         | 1.72         | 1.72         | 1.12        | 1.12        | 1.00        | 0.98        | 0.97        | 0.98        |
| 43       | 92.97         | 32.02        | 2.58         | 1.95         | 1.88         | 1.45         | 1.45        | 1.45        | 1.11        | 1.10        | 0.73        | 0.72        |
| 44       | 3.04          | 2.23         | 2.23         | 2.23         | 2.23         | 2.23         | 1.84        | 1.84        | 1.84        | 1.86        | 1.84        | <b>1.84</b> |
| 45       | 72.65         | 26.63        | 3.96         | 1.16         | 1.05         | 1.05         | 0.93        | 0.93        | 0.86        | 0.85        | 0.78        | 0.78        |
| 46       | 9.53          | 2.47         | 2.47         | 2.45         | 2.45         | 2.45         | 1.73        | 1.73        | 1.73        | 1.56        | 1.62        | 1.56        |
| 47       | 100.73        | 67.46        | 14.99        | 9.93         | 9.93         | 7.17         | <b>7.17</b> | 0.00        | 0.00        | 0.00        | 0.00        | 0.00        |
| 48       | 99.66         | 57.27        | 16.38        | 8.49         | 8.49         | 2.42         | 2.42        | 0.77        | 0.77        | 0.77        | 0.77        | 0.75        |
| 49       | 64.03         | 18.82        | 3.11         | 2.64         | 1.77         | 1.77         | 1.69        | 1.69        | 1.46        | 1.43        | 0.87        | 0.87        |
| Average: | 74.28         | 25.43        | 8.25         | 4.27         | 3.50         | 2.30         | 1.67        | 1.35        | 1.08        | 0.95        | 0.83        | 0.78        |

*“The foremost cartographers of the land have prepared this for you; it’s a map of the area that you’ll be traversing.” Blackadder opens it up and sees it is blank. “They’ll be very grateful if you could just fill it in as you go along.”*

Black Adder

*“The small part of ignorance that we arrange and classify we give the name of knowledge.”*

Ambrose Bierce (1842-1914)

# 5

## Mapping

### 5.1 Introduction

---

This chapter covers the production of maps using the observations made by the spectrometers, global positioning system (GPS) and depth and altitude instruments. The distribution of observations onto abundance layers representing the location and proportions of end-members is described. The method by which the uncertainty in the vehicle’s position is used as a factor for distribution of observations onto the map is also explained. A confidence factor accounting for the estimated accuracy of an observation is factored into the mapping of the observation into the layers.

The production and format of these layers is described in detail, and an evaluation of the effect of variation of parameters on the accuracy of the results is shown through visual representation of reconstruction mapping of a standard computer graphics image.

The use of the term ‘layer’ refers to a single feature (e.g., one end-member, or depth), and the term ‘map’ represents a combined group of similar features (e.g., an abundance map consists of multiple abundance layers). A map may also consist of a single layer (e.g., the confidence map) and in these cases the term layer and map are synonymous.

### 5.2 Habitat Maps

---

The author developed the mapping software used in this work. It consists of Python code, and it was used to create the sample maps and test image reconstructions included in this chapter. The first program, `unknownDetector.py`, performs blind unmixing of the spectral measurements and produces a proportion estimate for each measurement, and a library of end-member spectra. The proportion estimate represents each measurement in terms of proportions of the detected end-members. The second program, `cblMap.py`, takes each of these unmixed proportion estimates and distributes them over the area of uncertainty of the position of the vehicle. The distribution ensures that a single unit of classification is added to each map layer for each measurement. This program also produces the graphical display of the resultant maps as shown in this work.

The desired outputs of the mapping process are a layered habitat map showing the location and abundance of types of benthos in the mapping area, a confidence map showing how accurate the habitat map is likely to be, a coverage map showing how well mapped the area is, and a bathymetry map showing the depth of the seafloor in the mapped area.

#### Observation

An measurement is a group of related data recorded at a single place and time. Observations are derived from these data, and contain a list of the proportions of each end-member (obtained through the unmixing of spectral data), a value representing the confidence in the observation, the seafloor depth, the physical location, and a measure of the accuracy of the location information. Each observation is distributed onto all of the map layers by the method described in Section 5.4.

#### 5.2.1 Format

This system produces a map consisting of multiple layers, each a grid of cells. Each map cell consists of all the layer cells that are in the same location, and represents the description of the physical area covered by the cell.

## 5.3. DESCRIPTION OF MAPS

---

### Layer Cells

Each layer cell corresponds to a physical location in the area being mapped. The physical size of the area represented by a cell is determined by the lower limit of the vehicle position estimates, but is independent of observations, so maps with different cell resolution may be produced by this system.

Layer cells indicate the characteristics of some feature in the mapped area at that cell's location. A layer cell accumulates the total amount of observation fractions of the feature represented by the layer to which it belongs, and a count of the proportional allocation of observations made to that cell. The sum of observation fractions divided by the count of proportional allocation gives the cell's measurement of the feature.

### Cell Resolution

The practical precision limit of a map is restricted by the accuracy to which location of the measurements made by the vehicle can be determined. The position of the autonomous underwater vehicle (*Searise*) is limited in accuracy by the GPS accuracy, and time between GPS fixes (during which dead-reckoning positioning calculations are employed). GPS accuracy with the equipment carried on board *Searise* approaches approximately  $\pm 2$  m (see Section A). The grid-size of a generated map is determined by weighting the memory requirements of using a fine-scaled resolution against the size of the area being mapped and the accuracy of the vehicle's positional estimates. For any grid-cell size, uncertainty in the position of the autonomous underwater vehicle (AUV) increases during underwater transects, to an extent where the positional uncertainty is greater than the size of the grid cells themselves. At this point, any classification made is in danger of being allocated to a completely wrong grid cell (map location).

## 5.3 Description of Maps

---

The map outputs consist of a coverage map, multiple layers forming an abundance map, a confidence map, and a bathymetry map. These maps are described below.

### 5.3.1 Coverage Map

The coverage map is a representation of the relative amount of observations distributed to each map cell. Those cells which have had more observations inserted into them (taking into account the distribution method; see Section 5.4) have higher coverage values. This map is useful for understanding how well mapped any location is, compared to the other mapped locations.

Figure 5.1 (right) shows a sample contour plot of the coverage map. The colour represents a relative coverage between cells, not an absolute coverage value, and identifies those cells most and least visited.

The coverage map may accumulate between successive mappings in an area, giving a lifetime representation of the relative quality of coverage of the mapped area.

### 5.3.2 Bathymetry Map

The bathymetry map (see Figure 5.1) represents the depth of the seafloor. The seafloor depth is the sum of the vehicle's altitude (determined via echolocation) and vehicle depth (determined by pressure sensors). See Figure 2.5 for a representation of the depth calculations.

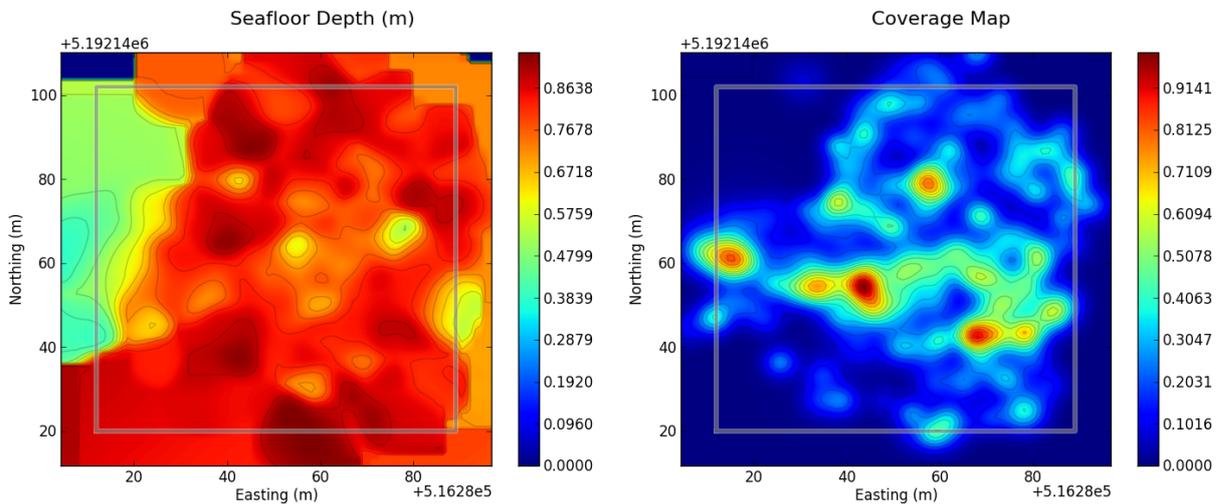


Figure 5.1: Bathymetry map from Cloudy Bay Lagoon (left), and coverage map (right).

### 5.3.3 Abundance Map

The abundance layers (see Figure 5.2) are an associated group of map layers representing the distribution of end-members in the mapped area. There is one layer for each end-member type. Each layer records the proportion of the associated end-member type at each cell location. The proportions for any cell across all abundance layers preserve the sum-to-one constraint (see Equation 3.3.3).

The group of corresponding cells in the abundance layers represents the classification at any cell location. The coverage of each cell (a representation of the number of observations distributed to the cell) is given by the count. The proportion of the object type for each cell is calculated by the sum of proportions divided by the coverage.

The unmixing of spectral data (see Chapter 4) produces a fractional proportion for each end-member type. This fractional amount is distributed over the abundance layers, based on positional uncertainty. The distribution process is described in Section 5.4.

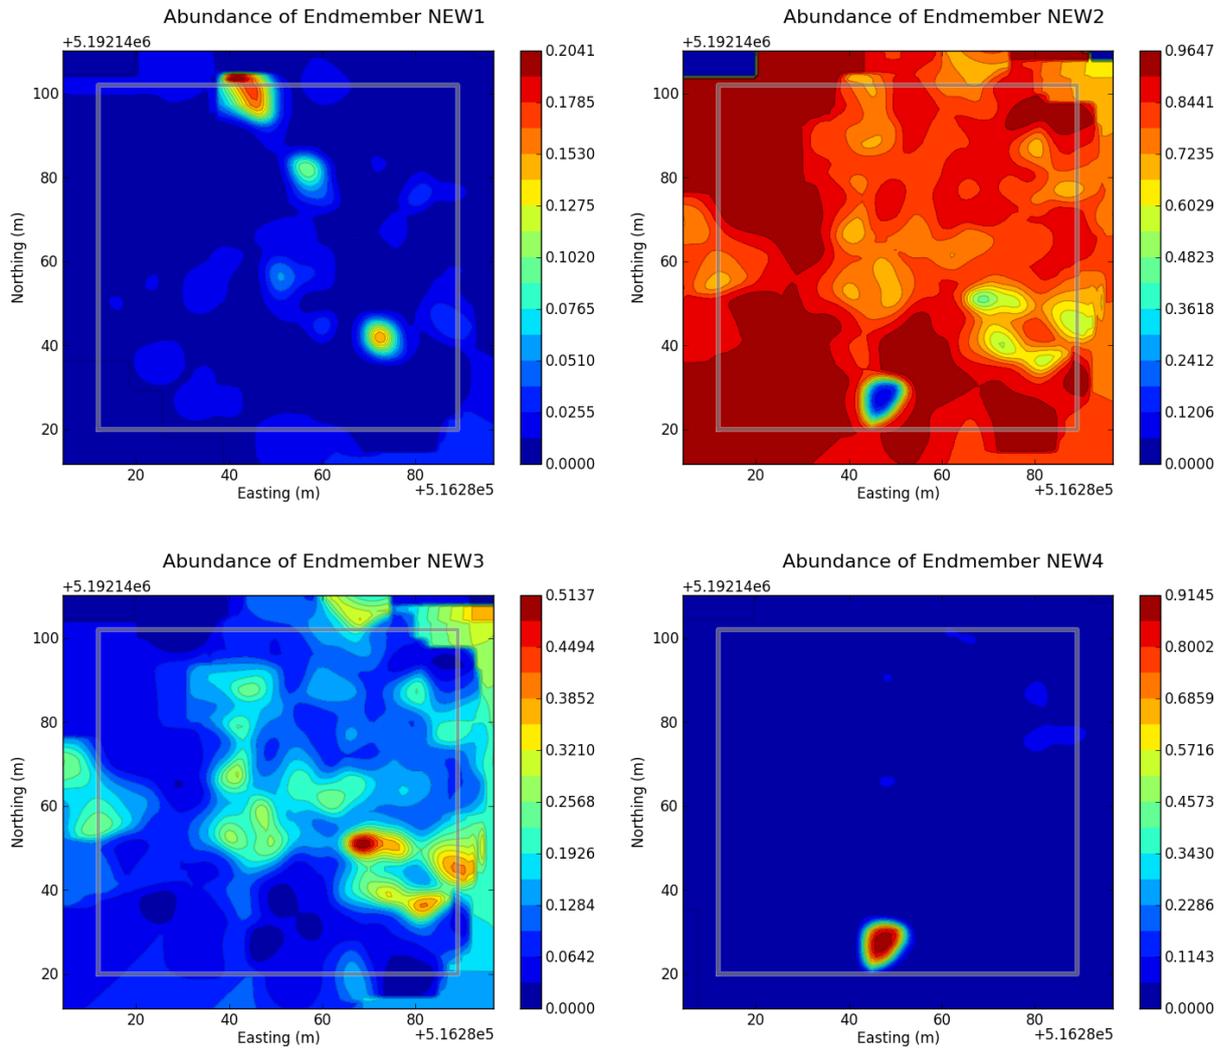
### 5.3.4 Confidence Map

The genetic algorithm (GA) fitness function (Equation 4.5.1) provides a measure of the quality of unmixing results, as it compares the observed spectra with reconstructed spectra based on the mix estimate. This metric contributes to a map representing confidence in the quality of the mapping at each location. The quality of an unmixed solution is also dependent upon the altitude at which an observation is made. As altitude above the seafloor increases, attenuation of the reflected spectra increases twofold (with respect to the upward spectrometer's measurement of down-welling light), as the light path travels from that spectrometer's measurement point to the seafloor and back again before being measured by the downward spectrometer. Although a correction process performs a de-attenuation, the attenuation factor is estimated, not measured, so the higher altitude readings are less accurate. Other factors that can affect the accuracy of measurement include turbidity and uncertainty in vehicle position.

These factors are combined to form a confidence value for each observation, which is translated into the layered map by attenuating the observation's contribution to the map. This attenuation is achieved by multiplying the confidence value by the height at each of the Gaussian sub-sampled points.

The confidence map can be used in conjunction with the coverage map when evaluating future mapping paths for maximising scientific return. The confidence

## CHAPTER 5. MAPPING



**Figure 5.2:** Contour plots of abundance layers, forming a combined abundance map. The layers are complementary, showing the same physical area, and the abundances at any location sum to unity. Each graph represents a single end-member, and the colours represent the magnitude of expression of that end-member. The usage of colour is different for each layer, as the scales are different ranges. Two types of abundance layers are created; a global abundance and a relative distribution (see Section 6.3.5). The layers shown here are the result of a 4 end-member classification of the data from a field-test on Bruny Island (see Chapter 6).

## 5.4. OBSERVATION DISTRIBUTION

---

map can also identify areas of consistently poor-quality unmixing, which may be a consequence of the presence of new end-member types or high variability in spectra for a given end-member.

### 5.4 Observation Distribution

---

To counteract the problem of not knowing the precise location of the vehicle, and hence exactly where observations have been made, this system distributes observations as a ‘unit of classification’ across the map based on a two dimensional Gaussian function (see Equation 5.4.1, and Figure 5.3). The centre-point of the distribution represents the vehicle’s assumed position. The  $z$  axis of the Gaussian at any point in the  $(x, y)$  plane represents the probability of the vehicle being at location  $(x, y)$ , and the sum of the  $z$  values over all  $(x, y)$  locations represents a single observation. The dimensions of the Gaussian are determined by the uncertainty in vehicle position ( $\sigma$ ). Each single observation is distributed across the area in which it is likely to have been made, as a series of fractional observations.

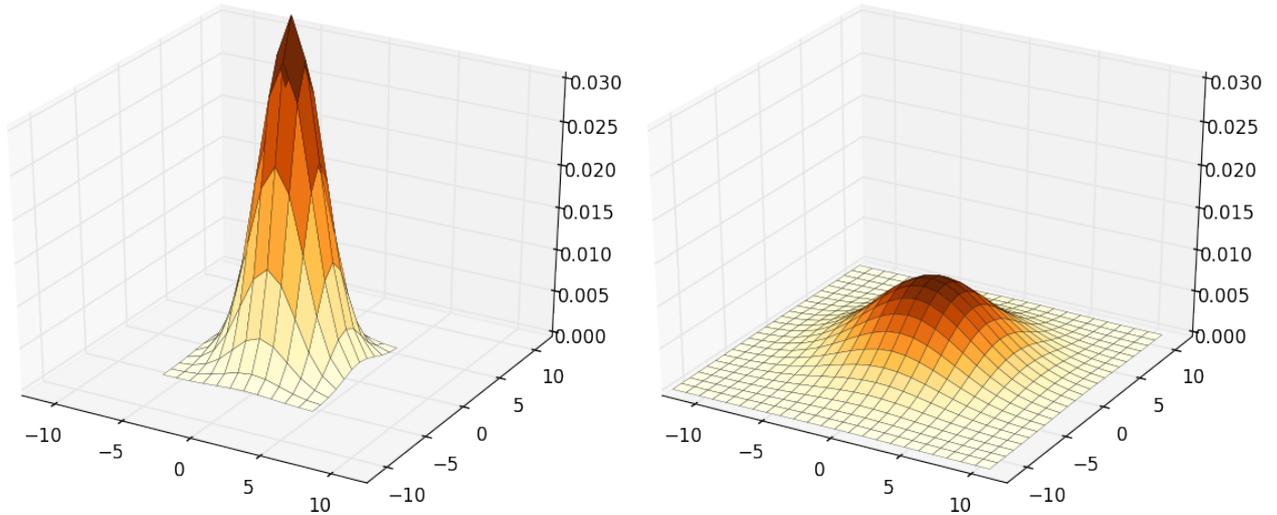
Such problems are the province of the field of spatial statistics. Spatial statistics involves the use of statistical methods in the analysis of spatial data. These statistical methods are used to characterise and make inferences about spatial patterns, particularly in ecology and geology. For an overview of the state of the art in this field, see Gelfand et al. (2010).

The area of the seafloor represented by an observation varies with the altitude of the vehicle. The fibre-optic cable (FOC) collects light with a  $19^\circ$  cone field of view (FOV). This relationship between the FOV and altitude is negligible compared to grid size and uncertainty in vehicle position, but could be factored into the observation distribution by increasing the  $\sigma$  based on altitude.

$$G(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \quad \boxed{5.4.1}$$

All maps incorporate observations using the same distribution mechanism.

The normal distribution is assumed to be a valid representation of the vehicle positional uncertainty, but other distributions could be used (for example, a skewed function including the effects of water currents). The normalisation process allows for any spread function to be used in the same manner as the Gaussian; the important factor being that the quantity of observation distributed over the map layers is constant, regardless of the  $\sigma$  value, sub-sampling resolution, or grid size.



**Figure 5.3:** Normal distributions with  $\sigma$  of 2 and 4, plotted at the same scale. The  $x$  and  $y$  axes are offsets (m) from the vehicle’s estimated position. The extents of the  $(x, y)$  axes are  $\pm 3\sigma$  for  $\sigma = 4$ .

### 5.4.1 Normalising the Sub-sampled Distribution

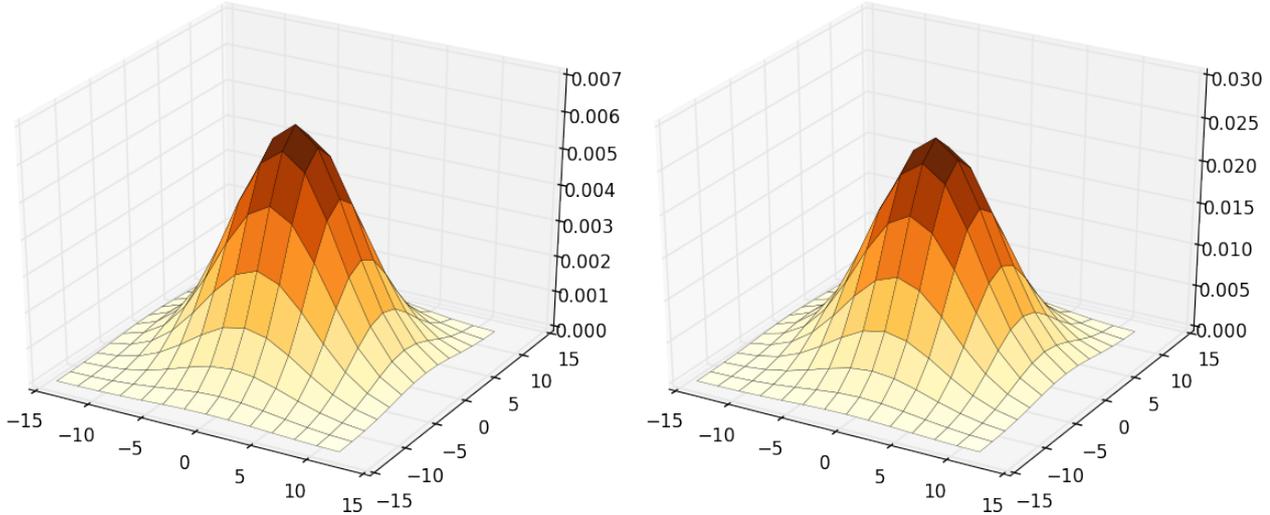
The unmixed proportion of each end-member must be distributed over the map area covered by the Gaussian curve, in proportion at each  $(x, y)$  position dictated by the Gaussian height at that position. The area under the curve represents the entire contribution of the observation to the map. To achieve this distribution, an iteration of  $(x, y)$  position over the Gaussian curve area is performed between  $\pm 3\sigma$ . This range limits the processing time required for distribution of the classification over each layer. Beyond the  $3\sigma$  extremes, the height of the Gaussian function (hence, the probability of the AUV being at that position) is close to zero. Approximately 99.73% of the values lie within 3 standard deviations of the mean.

The sum of the probability distribution of the classifications to the map layers must be constant (all observations contribute an equal amount to the map – see Equation 3.3.3), representing a single observation, regardless of  $\sigma$ .

The  $(x, y)$  iteration over the Gaussian between  $\pm 3\sigma$  is discrete, sub-sampling the Gaussian curve, and  $z$ -values (probability) are point samples on the curve. Although the Gaussian function creates a surface enclosing a constrained area, the discrete sub-sampling of this curve when distributing onto the map layers means that unequal contributions of observations would be made, based on the value of  $\sigma$ .

## 5.4. OBSERVATION DISTRIBUTION

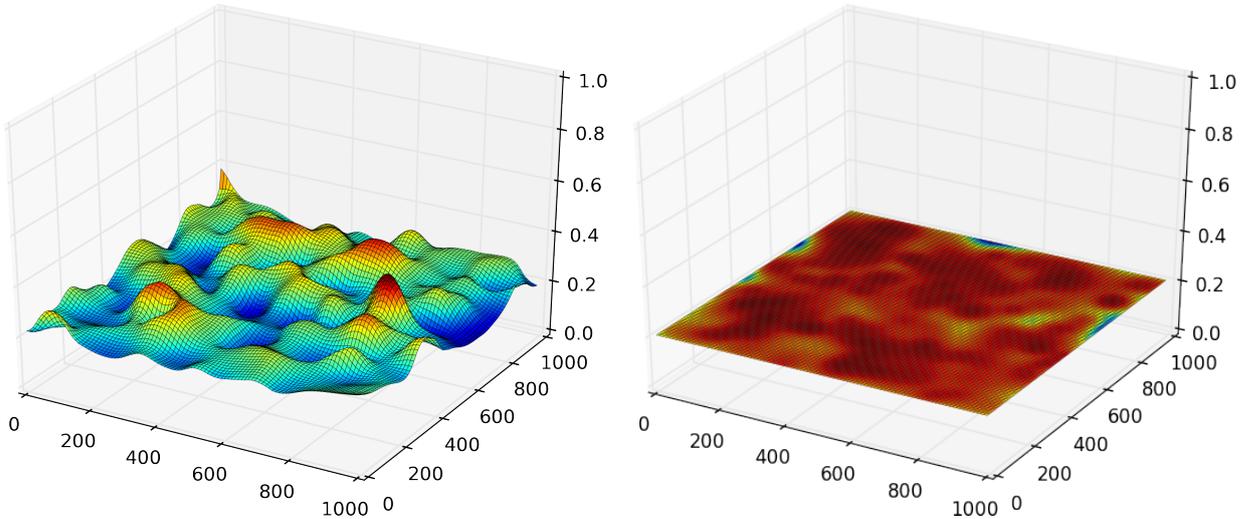
This issue is corrected by first normalising the Gaussian, as shown in Figure 5.4. In the graph at left, shown with a  $24 \times 24$  mesh, the contribution is only 46.8% of expectation. To enforce a 1-unit distribution, the sum of all  $(x, y)$  probabilities at the grid intersections is pre-calculated and the individual values are normalised so that the total sum of probabilities is unity. This provides a resolution-independent and  $\sigma$ -independent method for distributing the classification onto the map.



**Figure 5.4:** Adjustment of 2D-Gaussian function ( $\sigma = 4$ ) to ensure constant contribution of sub-sampled data to the map. The  $(x, y)$  axes are offsets in metres, and the  $z$  axis represents probability that the vehicle is at the  $(x, y)$  location. Because of the discrete sampling of the curve at left, the area under the curve calculated as the sum of the heights at grid intersections is approximately 0.247. The same curve (normalised at right; note the difference in  $z$  scale) is adjusted so that the calculated sum of the discrete sample points is 1.

The  $(x, y)$  iteration over the Gaussian area uses the height ( $z$ ) of the Gaussian curve at offset  $(x, y)$  as a representation of the fraction of the unmix proportion for the end-member at the  $(x, y)$  offset from the estimated observation position. The sum of the  $z$  values over all  $(x, y)$  positions should be the same as the area under the curve, but because of the discrete point sampling for each axis and varying  $\sigma$ , this sum is variable. Without adjustment, the unmix proportion will not *totally* be allocated to the map.

To remove this variation caused by sub-sampling and differing  $\sigma$ , the Gaussian heights ( $z$  values) are first summed and each value is normalised (divided by the sum) so that the total contributions of all points is constant (100%) regardless of the  $\sigma$  value, and the value at each  $(x, y)$  point is a valid representation of the proportional fractional contribution to the whole.



**Figure 5.5:** Sample layer before and after normalisation. The normalised graph shows that the original observation ( $a=0.1676$ ) has been correctly distributed over the layer. The slight inconsistencies in the normalised graph are caused by precision limitations in the floating-point mathematical calculations.

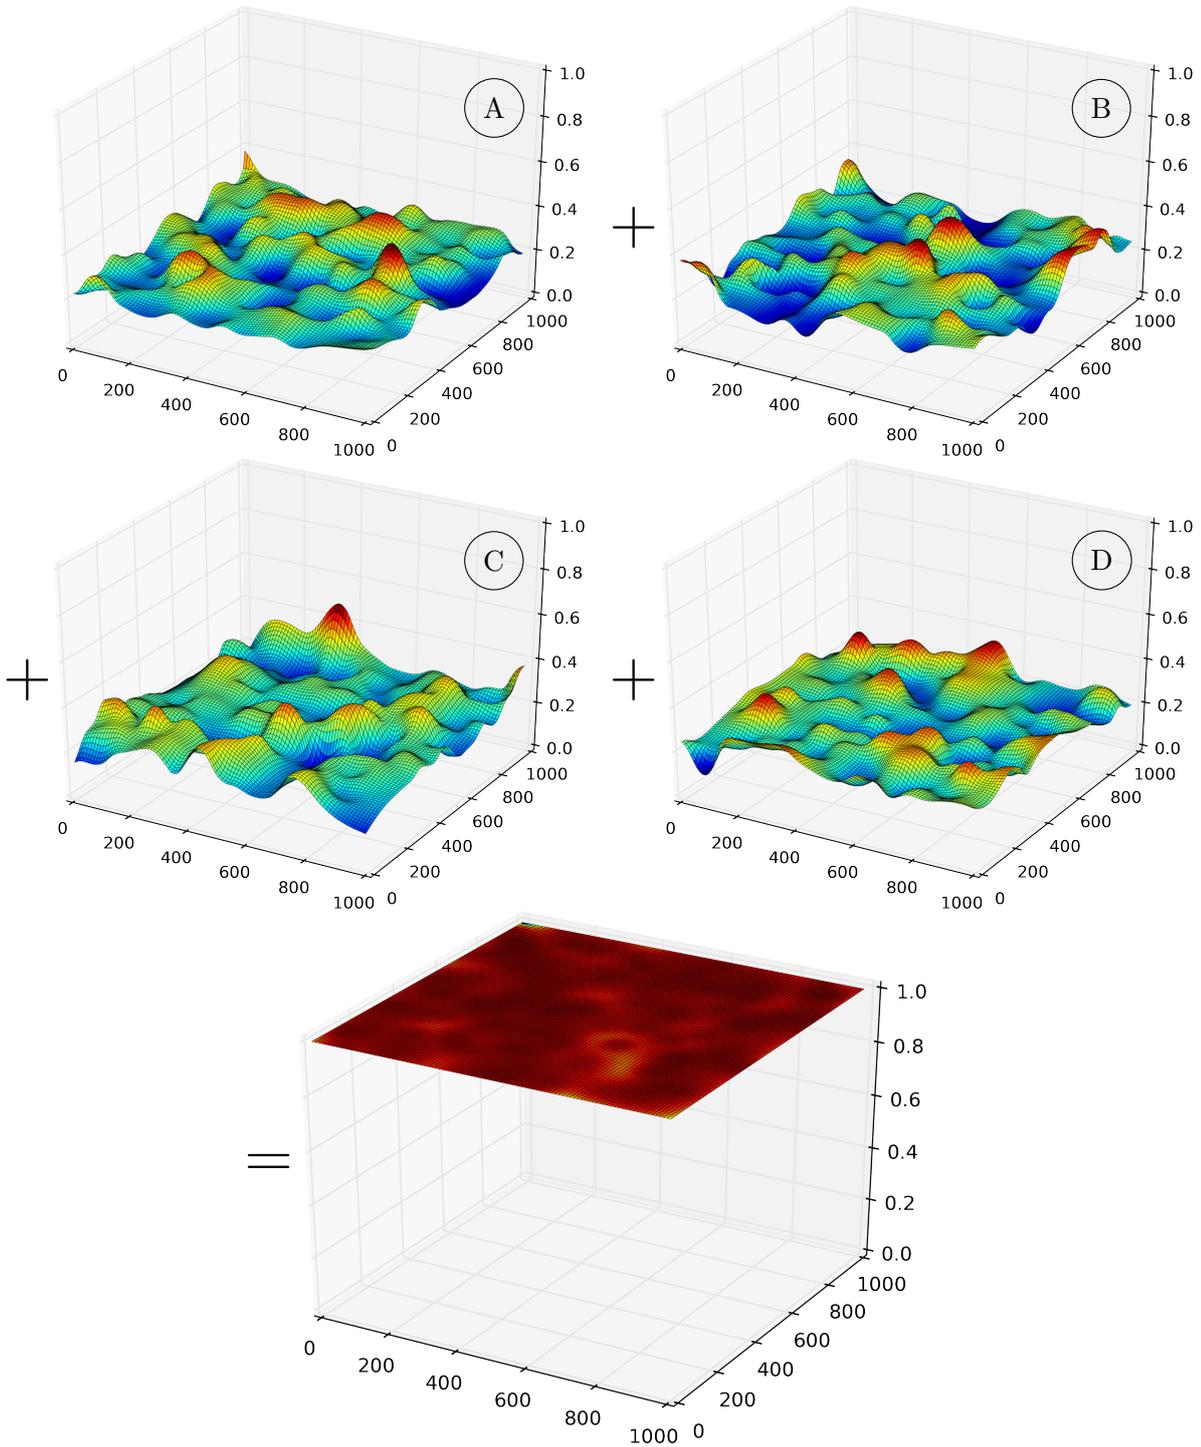
The constant area under the discretely sampled Gaussian curve, now independent of  $\sigma$  or mesh size, represents a unit of classification added to the map. This ‘unit’ is subsequently modified by a confidence factor (see Section 5.3.4).

### 5.4.2 Distribution to Cells

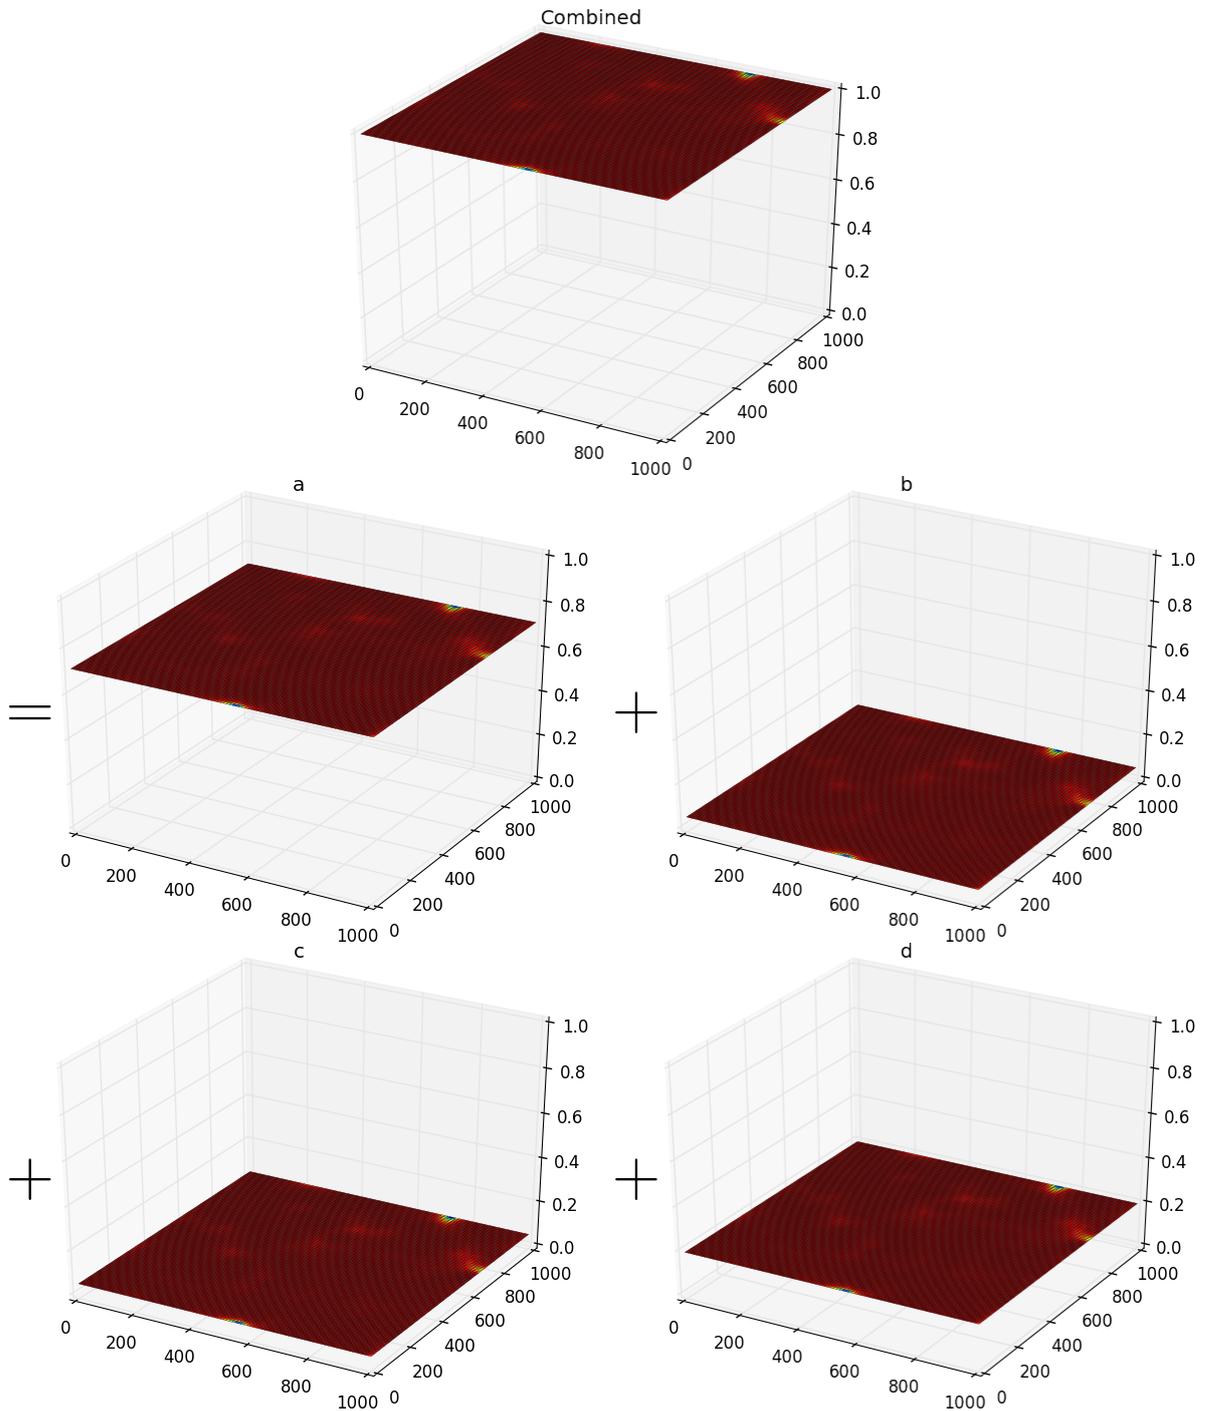
Once the normalisation factor has been calculated for the Gaussian, the observation is added to the map.

The normalised Gaussian is in the form of an  $(x, y)$  array, with dimensions determined by  $3\sigma$  and the mean at 0. The process of adding the observation consists of iterating over the array, and attenuating the observation readings by the  $z$  value at each  $(x, y)$  location. The normalisation has converted the array such that the sum of  $z$  values is now 1, so the complete iteration over the curve effectively adds 1 unit of an observation to the map, but distributes it over the area occupied by the normalised Gaussian.

## 5.5. VALIDATING THE ALGORITHM



**Figure 5.6:** The end-member abundance layers sum to one. Graphs (A) to (D) represent end-member abundance layers. The horizontal axes give location, and the vertical axis represents abundance. The layers sum to one, as shown in the lower graph.



**Figure 5.7:** Validation of Gaussian Distributions on a single observation repeated in 250 locations. The horizontal axes are offsets in metres, and the vertical axis represents proportional contribution at the offset location. Observation was a: 0.71, b: 0.05, c: 0.05, d: 0.19.

### 5.5 Validating the Algorithm

---

The Gaussian distribution of observations into the classification abundance layers must have the property that the sum of abundances of the layers at any location should equal 100%, as described in Equation 3.3.3. The generated maps should be an accurate representation of the seafloor, even in the presence of inaccuracy in vehicle position.

Confirmation that the software satisfies these requirements is performed below through a number of tests. These were performed by software written in Python by the author.

#### 5.5.1 Procedure

A sample grid 1000 m<sup>2</sup> with cell-size 10 m<sup>2</sup> was generated, with one abundance layer for each end-member. The validation was tested with a varying number of end-members ranging from 4 to 16, with similar results for all. Two hundred and fifty random observations were generated, consisting of a mix of end-member abundance fractions. These were located at random positions within the grid. Each observation consisted of a normalised random mix of the end-members, such that the sum of the proportions was unity (Equation 3.3.3). Each observation was added to the grid using a random standard deviation ranging from 20 m to 50 m, representing uncertainty in position. A random confidence value between 0 and 1 was also used, representing confidence in quality of the observation. At the end of this process, each abundance layer consisted of a grid of cells containing a sum and a count. Each abundance layer was normalised by dividing the sum by the count (see Figure 5.5), for each cell. These normalised layers were plotted, along with a summing layer showing the sum of all the normalised layers. The summing layer is observed to be uniform value of 1 (representing a mix of 100%), within the limits of the mathematical accuracy of the software and accounting for randomness of the observations, indicating that the assumptions behind the technique, and operation of the software are correct.

For each observation, the distribution curve is first generated, based on the standard deviation of the uncertainty in position ( $\sigma$ ), and the confidence of the measurement. The curve is of size  $3\sigma$  (99.7% of values lie within  $\pm 3\sigma$  of the mean in a normal distribution), and sub-sampled into a square grid with cell size twice the resolution of the cell size of the main mapping grid (per Nyquist). The sum of the sub-sampled points is calculated (effectively the estimate of area under the curve), and this total is used as a divisor to normalise the sub-sampled array such

that the sum of the sub-sampled points is unity. This sub-sampled curve now represents a single sample as a series of sub-sampled probability points which can be distributed over the abundance layers. A final modification to the curve is to apply the confidence, which is a simple multiplication of the height of all sub-sampled points, attenuating the contribution of an observation to the abundance layers based on the confidence level of the observation.

In the case of image reconstruction, the three abundance layers are the red, green, and blue planes of the image pixels. After all observation have been distributed over the individual layers, they are converted to a colour-image by representing the proportions of red, green, blue pixels as the abundance of each layer at the pixel position. These are combined to form the colour pixel for display in the final map.

### 5.5.2 Sum-to-one Constraint

This section validates the sum-to-one constraint through the use of randomly generated observations and examination of the reconstructed layers. Figure 5.6 show the results of a test using random observations consisting of four end-member types. The graphs display a 1000 m<sup>2</sup> area on the  $(x, y)$  plane, with the  $z$  axis showing the proportion of each end-member.

The random observations contained four end-members. Each observation had the sum-to-one constraint enforced. The requirement is that the Gaussian distribution of observations onto the map produces variable abundances of end-members for any map-cell, but the sum of abundance of end-members in any cell is 1.

The lowest graph in the figure shows a combined map where the proportions of each end-member at each  $(x, y)$  position is summed, representing the total allocation of end-members. The expected result is that a 100% allocation applies at all  $(x, y)$  locations, and this is confirmed by the plane at  $z = 1$ . The colour variation on the graph has been emphasised to highlight the extremely small variation in combined values. These variations are caused by accuracy constraints in the representation of small values in software.

### 5.5.3 Proportion Preservation

A test was performed using a single randomised observation mix for all samples (i.e., the same observation was distributed over the entire area). In this situation, the expected result is that each normalised layer should demonstrate the original

proportion for the associated end-member in the observation (despite random confidence, position, and standard deviation having been applied). This was observed to be the case (see Figure 5.7), where each normalised layer is a flat plane in  $z$ , giving the original mix proportion of that layer's end-member type in the original observation.

This test confirmed the operation of the algorithm in correctly distributing an observation over the Gaussian area with varying  $\sigma$  and confidence. The minor variations are artefacts of numerical inaccuracy inherent in floating point calculations.

## 5.6 Mapping Lena

---

This section examines the performance of the software through the use of a proxy for bottom habitat - a standard test image for computer graphics.

The effects of varying uncertainty in position ( $\sigma$ ) on the ability of the mapping algorithm was investigated by using an image as a 'ground truth', and simulating measurements from the image and reconstructions of the image using only those measurements passed through the mapping algorithm. The reconstructed images shown here generally overflow the borders of the source image, as the software displays the entire of the classified area. The original image extents are marked with a transparent box. It's interesting how in all cases the algorithm generates a very plausible extension of the original image.

The well-known 'Lena' image (see Figure 5.8) has been used as a standard in computer graphics for over 30 years. In the tests presented here, the image represents the seafloor, and the software performs a simulation of an AUV making transects over the image area and making observations. Dynamics of vehicle movement and uncertainty in position are simulated, and an attempt is made to reconstruct the image through the mapping process that would be performed by the AUV

The earlier discussion has shown that the mapping algorithm preserves proportional constraints, however this depends on the input data having a constant sum. The following tests use the red-green-blue colour triplet (RGB) pixel values of the Lena image as proxy for the proportion data for three end-member types. As these RGB triplets do not sum to constant value, the global proportional constraint is not valid. However, earlier tests showed that the constraint is preserved. These additional tests are intended to explore the effect of uncertainty and confidence on the mapping accuracy.



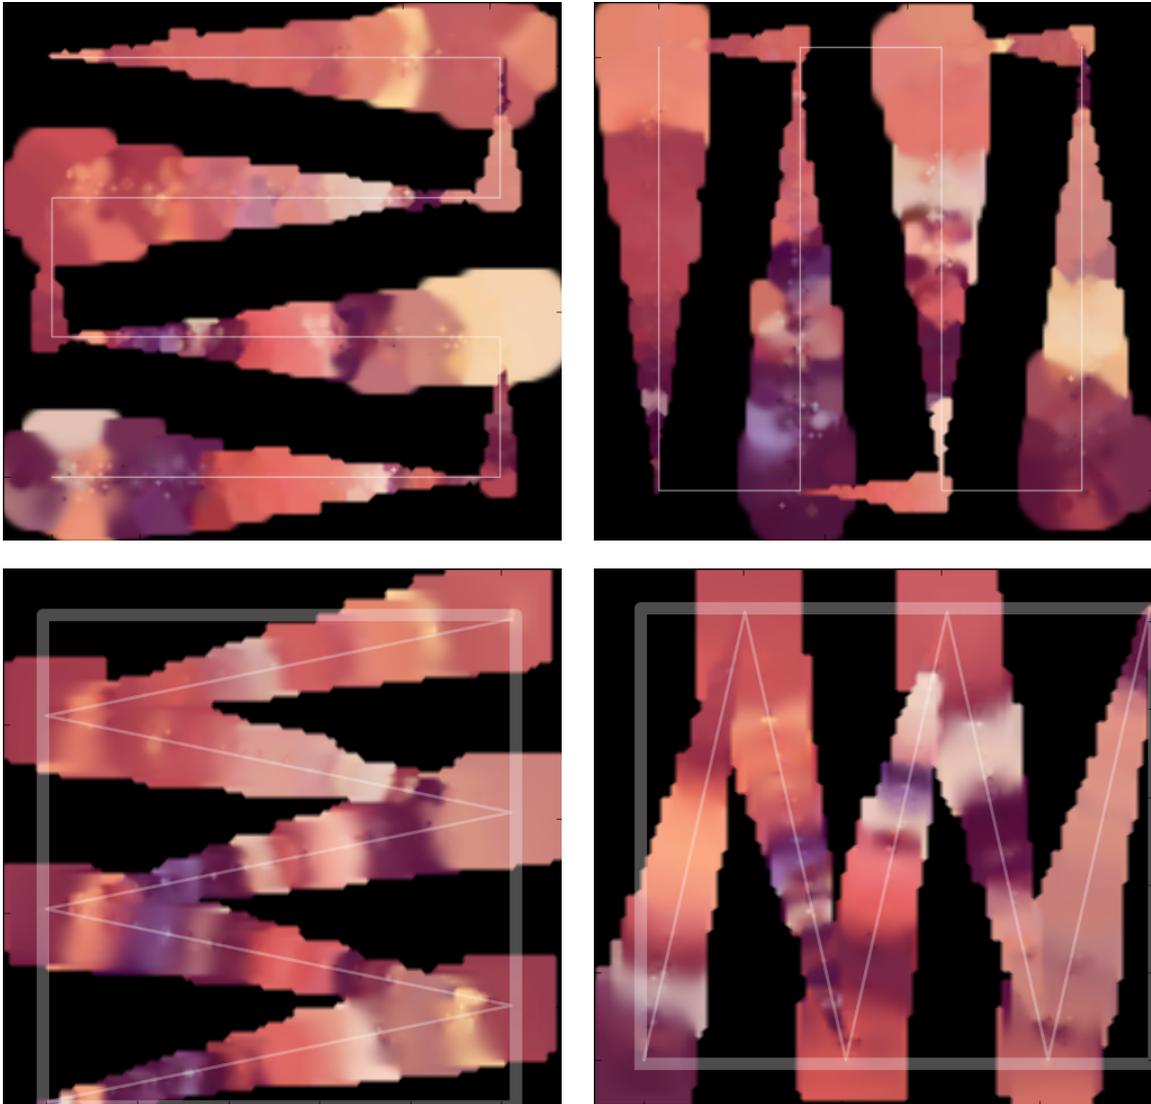
**Figure 5.8:** Ubiquitous computer-graphics test image used to represent the ‘ground truth’ for algorithm testing. This is Lena Söderberg. ©1972 *Playboy Enterprises, Inc.*

The Lena image is  $512 \times 512$  pixels, and a scale of 1 m/pixel is used for the simulations. The vehicle speed can vary, and the map’s grid resolution is 5 m/cell (approximately twice the GPS accuracy). The sparse mapping (Figure 5.13) used a sampling of 50 m/sample, whereas all other test simulations of vehicle movement use sampling of 1 m/sample.

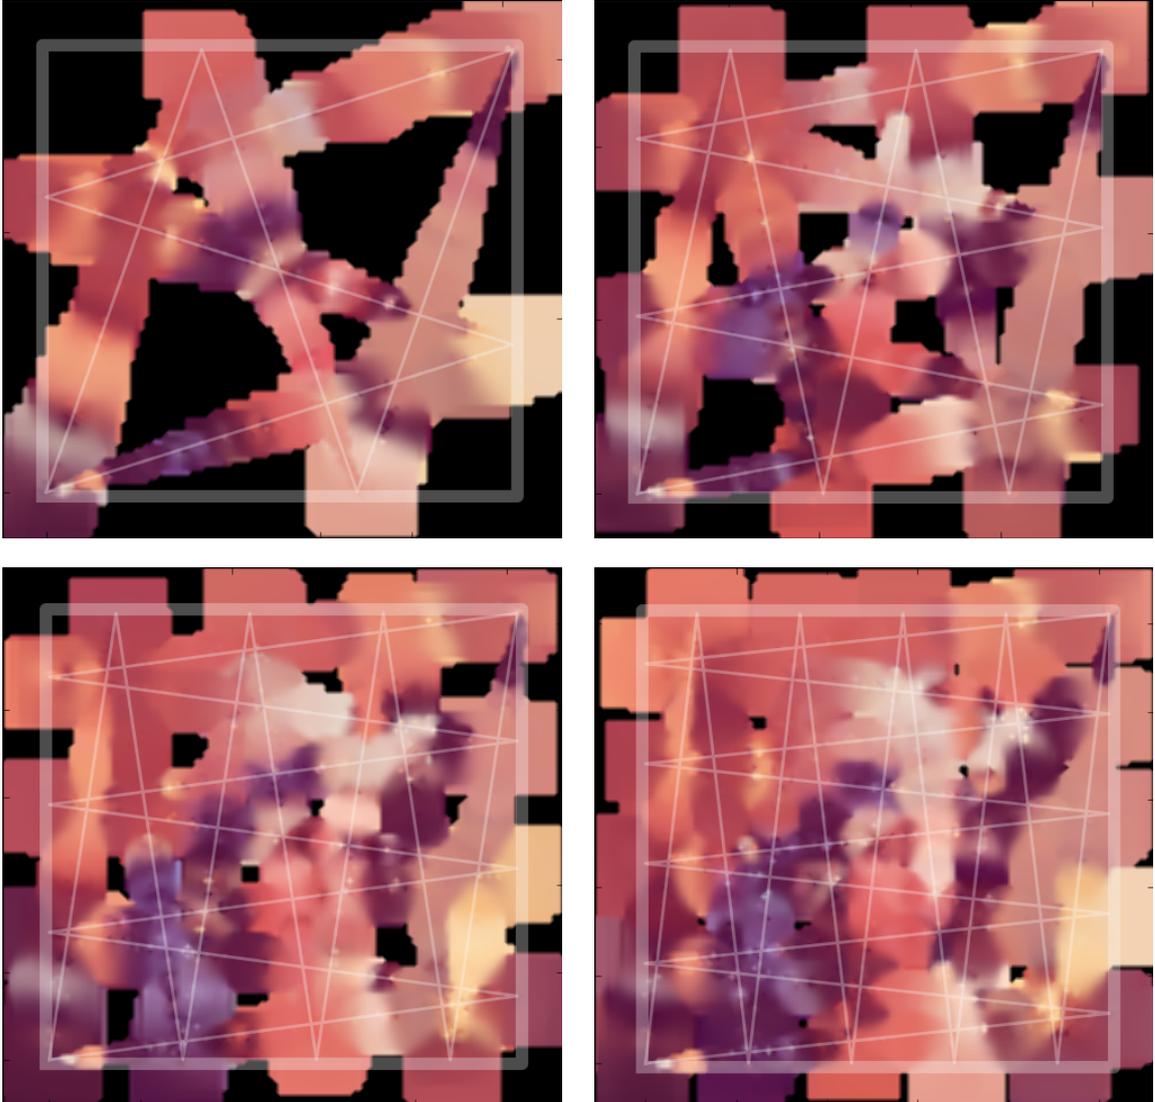
### 5.6.1 Transects

The simulations presented here transect a path along the image, recording observations at locations along the path. The vehicle is assumed to be underwater for each leg of the path, and consequently the vehicle’s position becomes increasingly inaccurate. This inaccuracy is modelled and controls the  $\sigma$  value passed to the Gaussian function used to distribute the observation on to the map.

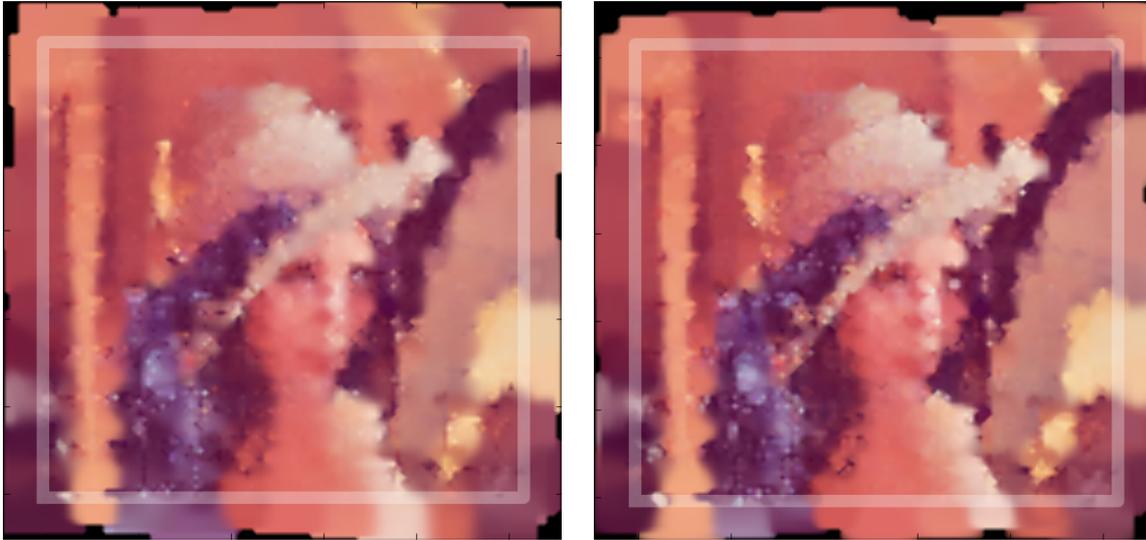
Traditional transects divide an area to be mapped into a rectangular grid and scan in two orthogonal passes: horizontal and vertical. Such an arrangement is shown in the white path lines in the top row of Figure 5.9. This scanning method does have



**Figure 5.9:** Transects (horizontal at left, and vertical at right) define the path of the AUV. Both paths are performed during the mapping, but are shown separately here to emphasise the process. As uncertainty ( $\sigma$ ) increases along the transect segments, the distribution of an observation is spread over increasing area. The top row shows traditional transects, and the lower row shows alternate transects used in this work. The alternate transect avoids overlapping line segments and offers longer continuous underwater legs.



**Figure 5.10:** Increase in transect ordinality improves image reconstruction coverage and quality. These examples show transects of order 2 to 5, referring to the number of points along each edge. The image of Lena is clearly recognised in the last example.



**Figure 5.11:** Biasing the relative contribution of observations with lower  $\sigma$  improves resolution. These images are created with the same  $\sigma$  and using order-10 transects. The rightmost image has more detail.

some disadvantages; the frequent changes of direction and short segment-lengths at the extremes of the area being mapped, and the overlapping transects at the edges.

This work uses an alternative transect method, as shown in the lower row of the same figure. In this method, the same rectangular grid is used to develop a zigzag transect pattern which avoids overlapping on the edges, and provides longer underwater transects. This pattern is described as ‘order- $n$  transect’, where  $n$  is the number of turn-points on the edge of the mapping area.

Both transect patterns offer a return-to-origin capability, an important consideration when operating the vehicle in inaccessible areas with only one or two people.

In the reconstructions shown in Figure 5.9, which are order-3 transects, the vehicle transects along the path shown by the white lines. Vehicle speed is 1 m/step, and the vehicle traverses each transect underwater. The  $\sigma$  increases linearly, although the actual value used for randomising location offsets is randomised within the  $\sigma$  maximum. The increase in  $\sigma$  is apparent when viewing the increased area over which measurements are distributed as the transect length increases.

The left images show the map reconstruction after the horizontal scanning transects have completed. At right, the vertical transects are shown. A full

mapping requires completion of both the horizontal and vertical transects.

The wide transparent border shows the dimensions of the original image. It is apparent that the mapping algorithm creates a mapping for areas outside this border, because of the  $\sigma$  uncertainty distributing observations over a larger area. The interesting offshoot of this is that the algorithm does an excellent job of extending images to larger sizes with plausible content.

### Halfway

Although the AUV positional uncertainty increases in a linear fashion as the vehicle transects underwater, on surfacing and re-establishment of a GPS fix, the positional uncertainty reverts to the accuracy of the GPS. It is reasonable to assume that just before that point, the uncertainty was closer to the GPS uncertainty than the accumulated uncertainty when the vehicle was underwater.

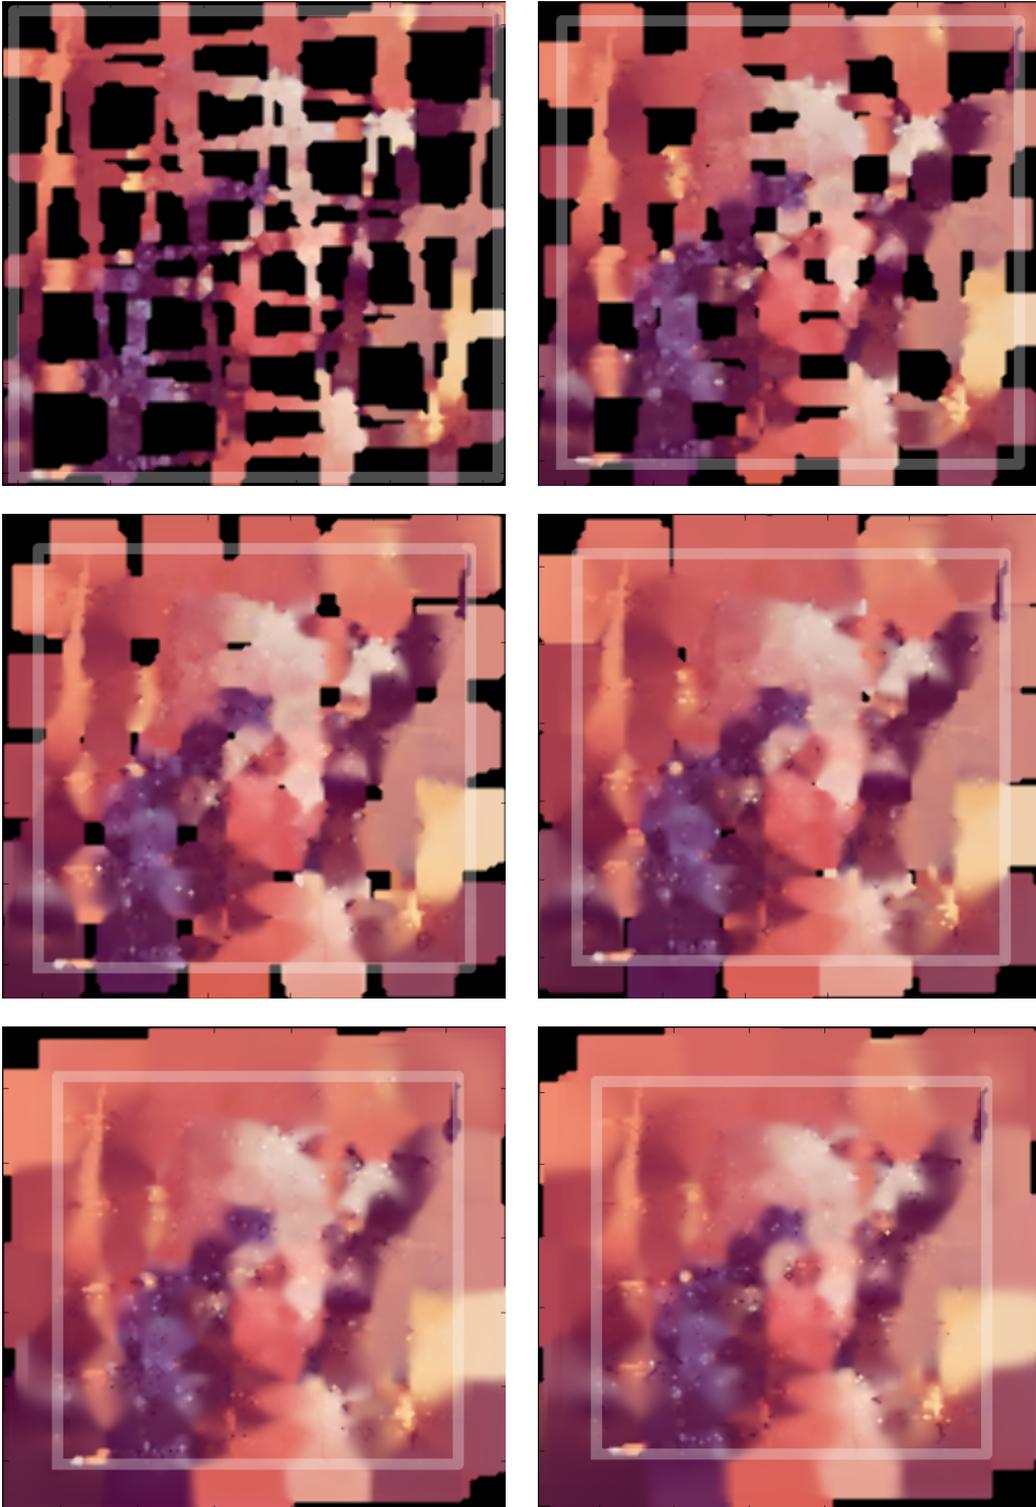
An assumption could be made that the vehicle has made a direct track from the last known GPS position to its current GPS position, in which case the uncertainty in position remains at the base level of the GPS accuracy. However, the case may be that the vehicle movements underwater were not in a direct line, in which case the maximum value of uncertainty can be assumed to be the halfway point between the last GPS fix and the current position, and the uncertainty increases from both GPS fix positions linearly to that point.

### Order- $n$ Transects

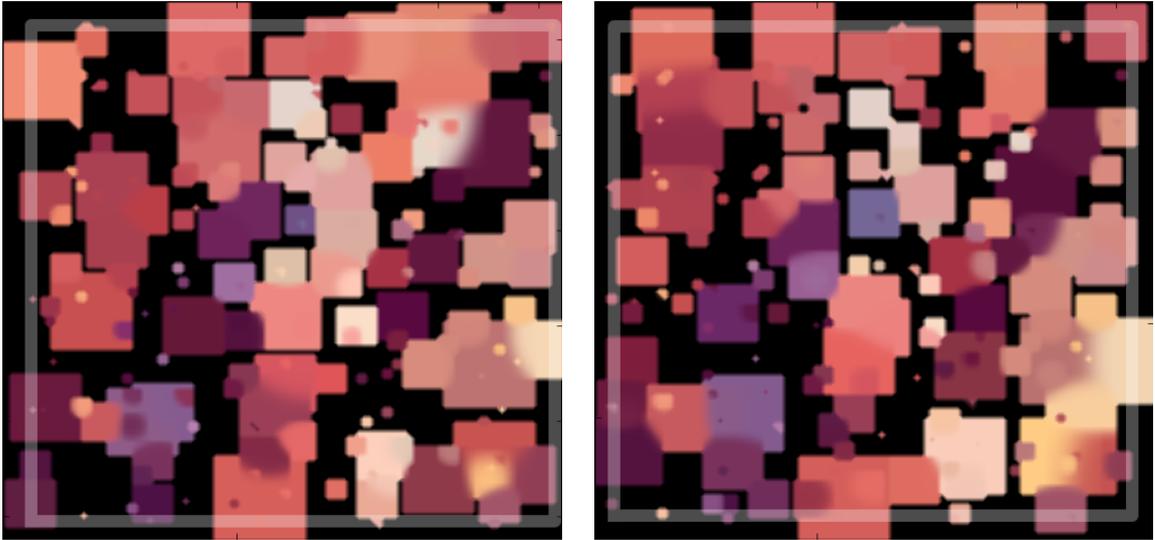
Tests shown in Figure 5.10 show the effect of increasing the ordinality of transects. Again, the white lines show the transect paths, and  $\sigma$  increases as the vehicle transects underwater. The coverage of the mapping increases with ordinality and, with the  $\sigma$  used here, full coverage is achieved by approximately order-5 transects (corresponding to a maximum separation of 100 m). This is a misleading measure, as it is dependant upon  $\sigma$  and speed.

### 5.6.2 Squaring the Circle

In Figure 5.13, the sampling frequency has been reduced to one sample per 50 m to highlight the unusual nature of the distribution algorithm. Although it appears to be the case that squares are used for distributing observations onto the map, this is not the case. Note that intersections of square areas are graduated and curved,



**Figure 5.12:** Variation in  $\sigma$  affects coverage but not significantly the sharpness of the reconstructed images. Examples are all using order-5 transects.



**Figure 5.13:** Successive runs of the mapping algorithm using low-density (50 m) sampling on an order-5 transect. The dual rectangular and circular nature of the distribution is apparent, particularly where there are intersections. The runs produce different observations because of slight variations in the vehicle position and the sparse sampling.

and that there is fine structure within these intersecting areas. The mapping is not simply a Gaussian blur of a sub-sampled image.

The squares are present as a side effect of the classification of layer cells. After observations have been distributed to cells, as described in Section 5.4, there may be some cells which have only one end-member detected. Furthermore, those cells may have only been allocated a very small classification count, for example, those on the extreme edge of the Gaussian. Even though we have a map cell containing only one end-member type, and it has only a minuscule amount of that type, the classification is derived from the sum of proportions divided by the count of distributions in the layer, which will equal 100%. So, the cell is classified as containing only that end-member. We see a square on the map. The coverage layer offers information about the distribution counts, and can be used to qualify the classifications of any cell.

The complex boundary between some of the squares' intersections is caused by the inherent Gaussian in the internal representation of the square, and the interaction between the sum of proportions and the count of distributions when there are two or more end-members present. The more observations and end-members detected, the less square-like the classified areas become.

### 5.6.3 The Effect of sigma

As noted,  $\sigma$  increases based on the duration of the underwater transect, as *Searise* is calculating its position through dead-reckoning. The increase in  $\sigma$  has the effect of spreading observations over the area of uncertainty of the position at which each observation is made. Figure 5.12 explores the effect of increasing  $\sigma$  on the reconstruction, using order-5 transects.

It is apparent in these examples that  $\sigma$  increase can vary significantly without impacting on the sharpness reconstructed image. The main effect of increasing  $\sigma$  is to improve the coverage.

### 5.6.4 Confidence

The reconstructed images in Figure 5.11 demonstrate the effect of biasing confidence in observations based on the magnitude of the uncertainty in position. The confidence attenuates the normalised Gaussian, reducing the contribution of an observation relative to other observations. Those observations with a higher degree of certainty (lower  $\sigma$ ) in position contribute more strongly to the classification map.

The image on the left uses a 1.0 confidence for all observations; no biasing is in effect. The image on the right uses  $1/\sigma^2$  as the confidence for observations, and this has brought out detail in the mapping not previously apparent.

## 5.7 Application

---

After a transect, when the AUV re-establishes a GPS fix, the position is known to within the accuracy of the GPS. In principle, a linear interpolation of position to the previous GPS fix could be used to determine location of observations made when the vehicle was underwater, and positional uncertainty could be constant along the path. An alternative method is to treat uncertainty as increasing from both GPS fix positions, reaching a maxima at the midpoint. This approach does not presuppose that offsets in position are linear for the whole path.

### 5.8 Summary

---

This chapter described the production of a multi-layer map showing the distribution of end-members as proportional probability layers. The production of a coverage map and its role in the representation of quality of results was also described.

A comprehensive discussion of the factors affecting reconstruction of mapped areas was presented, and an alternate transect pattern described. The algorithm was validated using a computer-graphics image as a representation of the seafloor.

*“However beautiful the strategy, you should occasionally look at the results.”*

Winston Churchill (1874-1965)

*“The best way to predict the future is to create it.”*

Peter Drucker (1909-2005)



# Field Testing

## 6.1 Introduction

---

This chapter gives details of a mission performed by the autonomous underwater vehicle (*Searise*) to validate the unmixing and mapping algorithms described in Chapter 4 and Chapter 5. The maps presented in this chapter, and the detected end-members, are generated automatically from the data collected during the mission without any human intervention.

## 6.2 Cloudy Bay Lagoon

---

Cloudy Bay Lagoon, Bruny Island, Tasmania (CBL) is a fairly remote shallow lagoon (approximately 2 hours drive southeast from Hobart, including a ferry trip) situated on the southern tip of Bruny Island (see Figure 6.3, Figure 6.2 and Table 6.1). It is a designated marine conservation area, in fairly pristine condition. A 2005 survey (Mount et al., 2005) of CBL identified diverse habitat in this area (see Figure 6.5). There is an oyster farm in the southern bay of the lagoon, operating outside the marine conservation area.

CBL is an excellent site for testing *Searise* operation and the algorithms described in this work, as:

- The lagoon has previously been mapped.
- The seafloor contains a diversity of sediment types.

## CHAPTER 6. FIELD TESTING

---



**Figure 6.1:** Photo of Cloudy Bay Lagoon, Bruny Island, Tasmania. The sample mission was performed approximately center of picture on the far shore.

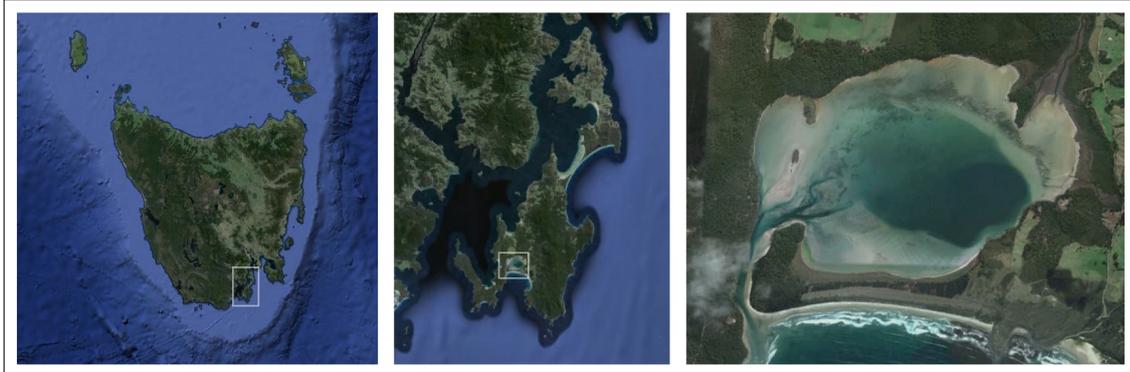


**Figure 6.2:** Satellite image of Cloudy Bay Lagoon. Figure 6.1 was at location A, and the mapping mission was performed at location B, 3.1 km away.

## 6.3. THE MISSION AT CLOUDY BAY LAGOON

---

©2013 DigitalGlobe



**Figure 6.3:** Location of Cloudy Bay Lagoon, Bruny Island, southeast Tasmania. The lagoon is 3.4 km from shore to shore at its widest point. Tasmania is approximately 260 km across.

- A variety of seagrasses and marine algae are known to be present.
- The lagoon is isolated from marine traffic.
- There is easy access to suitable sites for deployment of *Searise*.

A site a few hundred metres site to the north of the inlet, on the western shore of the lagoon, was chosen for this study (site B on Figure 6.2). This site is ideal because of its accessibility and the presence of a variety of underwater benthos. Access to the site is from the road along the western shoreline of the lagoon.

## 6.3 The Mission at Cloudy Bay Lagoon

---

The goal of the mission was to identify and autonomously characterise the seafloor in the study site, and validate the operation of the unmixing and mapping algorithms through the production of habitat maps as described in Chapter 5. The unmixing algorithm's halting condition (see Section 4.9.1) was set to 8 end-members, although Figure 5.2 shows the same site mapped using only 4 end-members.

A series of transects using Universal Transverse Mercator (UTM) coordinates from Google Earth (see Table 6.2) were specified for the mission. *Searise* performed these transects in order to provide a comprehensive coverage of the area and collect a suitable data set for the mapping algorithms.

The mission was performed on 23<sup>rd</sup> August 2013. Execution of the mission plan

## CHAPTER 6. FIELD TESTING

---

**Table 6.1:** Mission details from Cloudy Bay Lagoon.

|                       |   |
|-----------------------|---|
| Location <sup>a</sup> | UTM Zone 55G, Easting 516292 Northing 5192185<br>43°25'22" S, 147°1206" E |
| Mission Start         | 2013.08.15 12:42:22 AEST <sup>b</sup><br>1376534542 epoch s <sup>c</sup>  |
| Mission End           | 2013.08.15 14:53:11 AEST<br>1376542391 epoch s                            |
| Duration              | 7849 s  |

---

<sup>a</sup>Start/end location of transects (see Table 6.2).

<sup>b</sup>Format: year.month.day 24-hour:minute:second as Australian Eastern Standard Time.

<sup>c</sup>Epoch time counts seconds since 00:00:00 Coordinated Universal Time (UTC), Thursday, 1 January 1970, not counting leap seconds. This representation is convenient for time calculations.

was nominal, and no problems were encountered. Details of the mission are given in Table 6.1, and a summary of the data collected is given in Table 6.3.

### Processing of the Collected Data

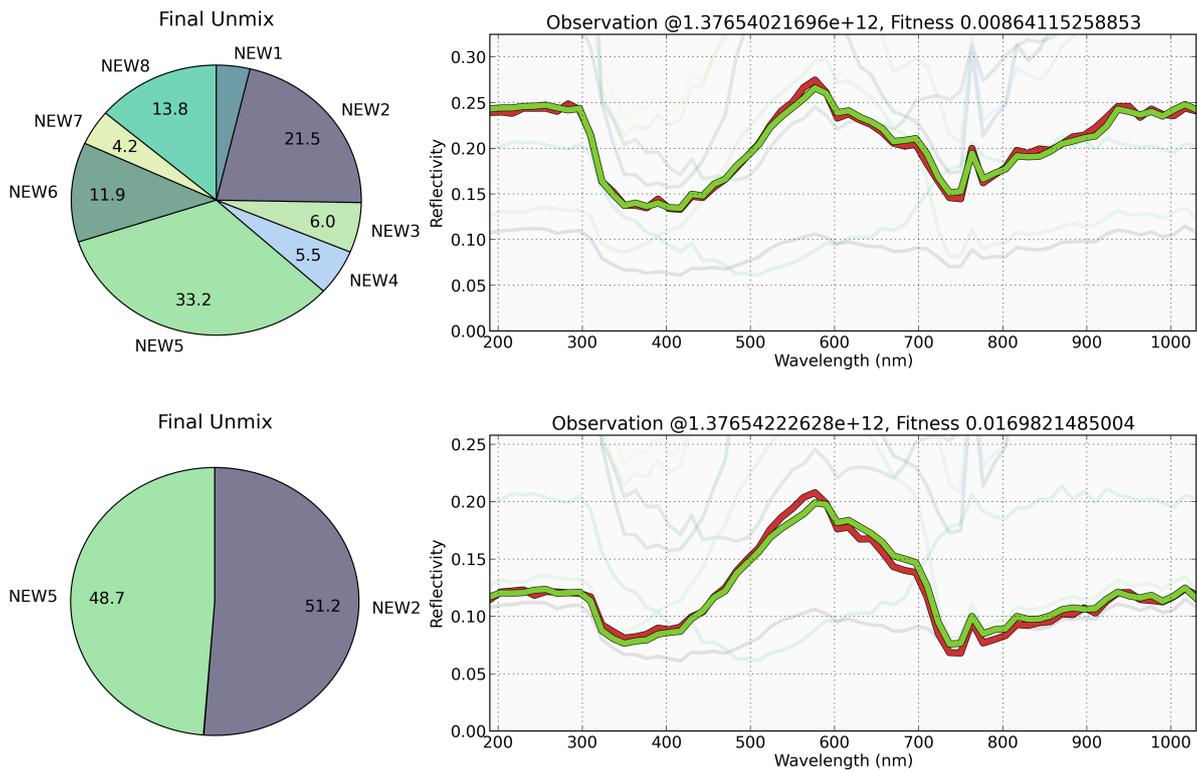
The unmixing algorithm processed the spectrometer data (see Figure 6.4 for two unmixed observations), producing a set of eight automatically detected end-members and a mix estimate using these for each observation. These end-members and the results of unmixing the observations were then processed by the mapping algorithm, which produced the maps shown below.

The author developed the unmixing and mapping software in Python. The graphs shown here were created by Python code also written by the author, using the `matplotlib` and `PIL` libraries.

### Identification of Images for End-members

Table 6.4 lists details of the end-members added to the library by the unmixing algorithm, in order of frequency of expression in the unmix results. The epoch times represent the time of the observation that was used as the holotype representation of each end-member, and these times are used to find the closest

### 6.3. THE MISSION AT CLOUDY BAY LAGOON



**Figure 6.4:** Samples of unmix results for two observations. The red line in the graphs is the observation made by the spectrometers. The green line is a spectrum reconstructed from the mix of end-members shown in the pie chart to the left of each graph. Some unmix estimates are complex (top), and some are very simple (bottom).

## CHAPTER 6. FIELD TESTING

---

**Table 6.2:** Coordinates of *Searise* transects for Cloudy Bay Lagoon mission. Coordinates are UTM Zone 55G Easting and Northing in metres. The transects were repeated four times to improve coverage.

| Leg | Easting | Northing |
|-----|---------|----------|
| 1   | 516292  | 5192185  |
| 2   | 516360  | 5192168  |
| 3   | 516313  | 5192195  |
| 4   | 516362  | 5192186  |
| 5   | 516319  | 5192222  |
| 6   | 516369  | 5192210  |
| 7   | 516331  | 5192241  |
| 8   | 516316  | 5192166  |
| 9   | 516349  | 5192242  |
| 10  | 516341  | 5192160  |
| 11  | 516365  | 5192226  |
| 12  | 516292  | 5192185  |

camera images. The easting and northing give the physical location of the holotype (UTM Zone 55G, in metres).

The camera has a 4-sample buffer, and was recording at 5 s intervals (0.2 Hz), so sample times are adjusted by +20 s when retrieving images. The holotype images are shown in Figure 6.14 and Figure 6.15.

### 6.3. THE MISSION AT CLOUDY BAY LAGOON

**Table 6.3:** Data collected during Cloudy Bay Lagoon mission.

| Sensor               | Type                  | Samples | Records              |
|----------------------|-----------------------|---------|----------------------|
| <i>Searise</i>       | state                 | 156,934 | records <sup>a</sup> |
| Depth                | combined <sup>b</sup> | 443,432 | samples              |
| Camera               | USB9000               | 1,556   | JPG images           |
| GPS                  | location              | 14,267  | fixes                |
|                      | uncertainty           | 15,006  | estimates            |
| Position             | calculated            | 313,518 | estimates            |
| IMU                  | 3DM-GX1               | 73,667  | samples              |
| Altimeter            | PA500                 | 7,017   | depth readings       |
|                      | CruzPro               | 16,502  | depth readings       |
| Optical Spectrometer | H04855                | 9,480   | spectra              |
|                      | H04856                | 3,424   | spectra              |

<sup>a</sup>*Searise* engineering data represents the internal state of the control system software, and includes positional information.

<sup>b</sup>Merged depth estimate from the two SSI Technologies depth-sensors installed in the tail unit. The P51-100 unit's reading is used above 40 m, and the P51-500 unit's reading is used below 50 m. A weighted, averaged reading is used between 40-50 m to provide smooth transition between the two units.

## CHAPTER 6. FIELD TESTING

---

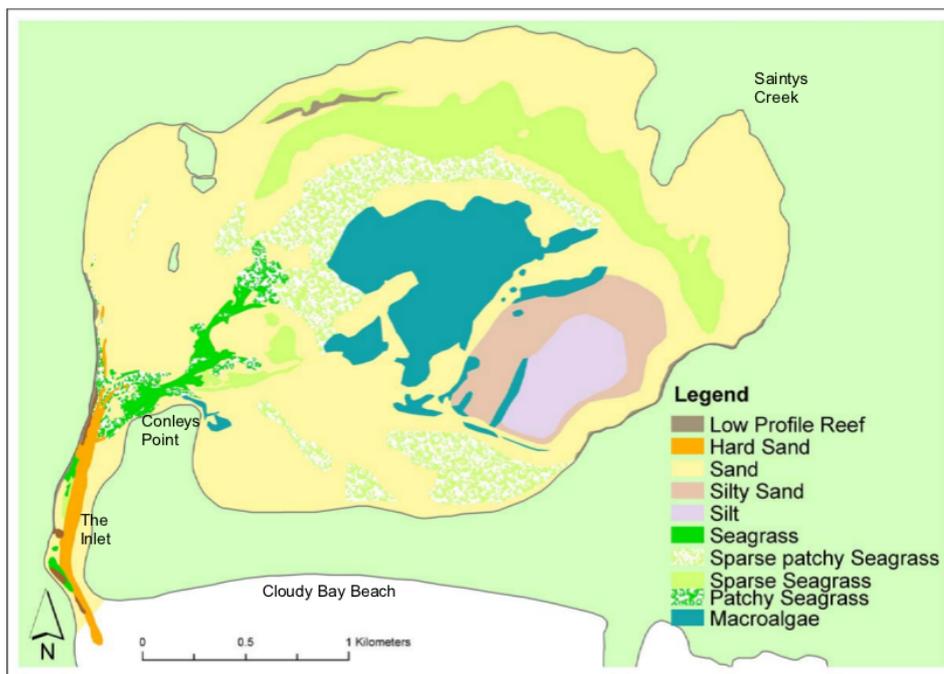
**Table 6.4:** End-members detected, listed in order of abundance. The times were generated by the unmixing software, and refer to the time of each observation chosen as the end-member holotype by the autonomous end-member detection algorithm. Using these times to reference the image data allows a visual of the end-member to be retrieved. Figure 6.14, Figure 6.15, Figure 6.16 and Figure 6.17 show the end-member spectra and associated camera images, retrieved using the times shown here.

| E.M. | Mix %  | Epoch s    | Easting | Northing |
|------|--------|------------|---------|----------|
| NEW2 | 38.88  | 1376538402 | 516296  | 5192210  |
| NEW5 | 23.28  | 1376541487 | 516338  | 5192221  |
| NEW8 | 9.35   | 1376538110 | 516346  | 5192209  |
| NEW6 | 8.48   | 1376537194 | 516352  | 5192216  |
| NEW3 | 8.36   | 1376535190 | 516338  | 5192191  |
| NEW4 | 5.16   | 1376540884 | 516341  | 5192162  |
| NEW7 | 4.68   | 1376535169 | 516356  | 5192190  |
| NEW1 | 1.81   | 1376541569 | 516347  | 5192219  |
|      | 100.00 |            |         |          |

**Table 6.5:** Interpretation of end-member mix.

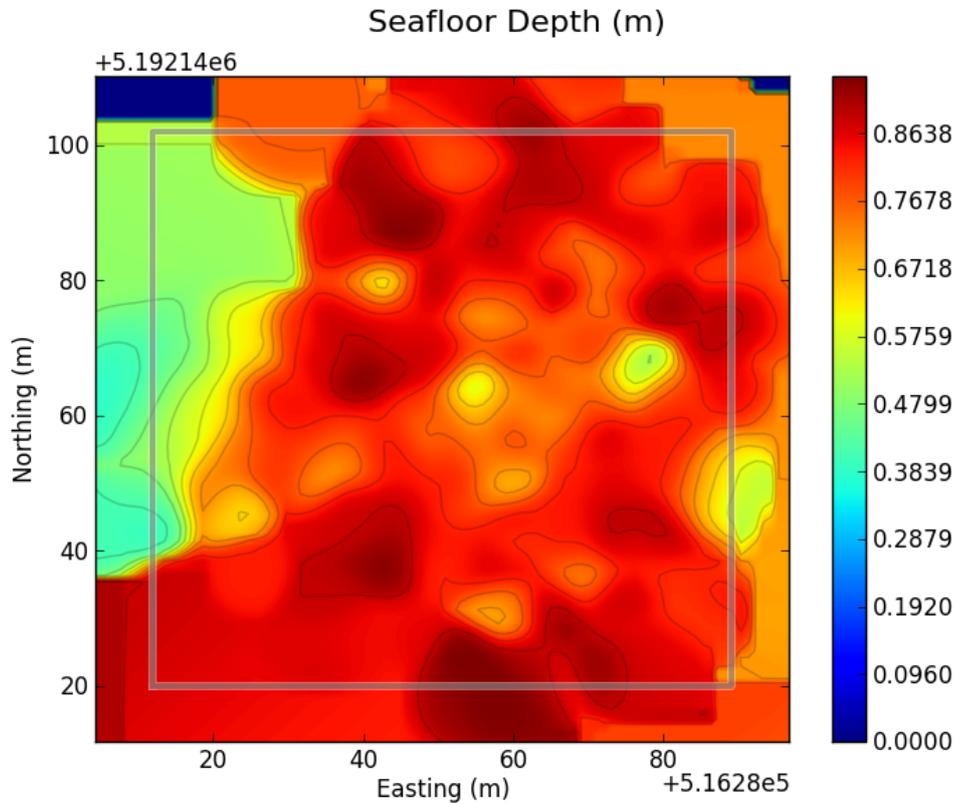
| E.M. | Interpretation  |
|------|---|
| NEW1 | Sand, with one image containing approximately 40% area of <i>Gracilaria</i> sp.   |
| NEW2 | Predominant yellow/brown sand, with patches of lighter sand and a darker material ringed around small burrows.  |
| NEW3 | Light-coloured Sand with islands of vegetation of type <i>Gracilaria</i> sp. covering approximately 40% of images. There may also be oysters present as a substrate for the vegetation. |
| NEW4 | Mixture of <i>Gracilaria</i> sp. (30%) and a green macro-algae (possibly <i>Ulva</i> sp.) on a sandy substrate.   |
| NEW5 | Darker sand/mud, with small raised mounds from burrowing activity. May also have very sparse seagrass <i>Zostera muelleri</i> .   |
| NEW6 | Lighter sand substrate (20%), with an equal mix of <i>Ulva</i> sp. (40%) and <i>Gracilaria</i> sp. (40%).   |
| NEW7 | Medium brightness sand with possible algae-covered small rocks very sparsely spaces.  |
| NEW8 | Light intensity sand with <i>Ulva</i> sp. (50%).  |

### 6.3. THE MISSION AT CLOUDY BAY LAGOON



**Figure 6.5:** Cloudy Bay Lagoon habitat map (from Mount et al. (2005), pp.67). This map provides broad-scale classifications of habitat.

### 6.3.1 Seafloor Depth Map

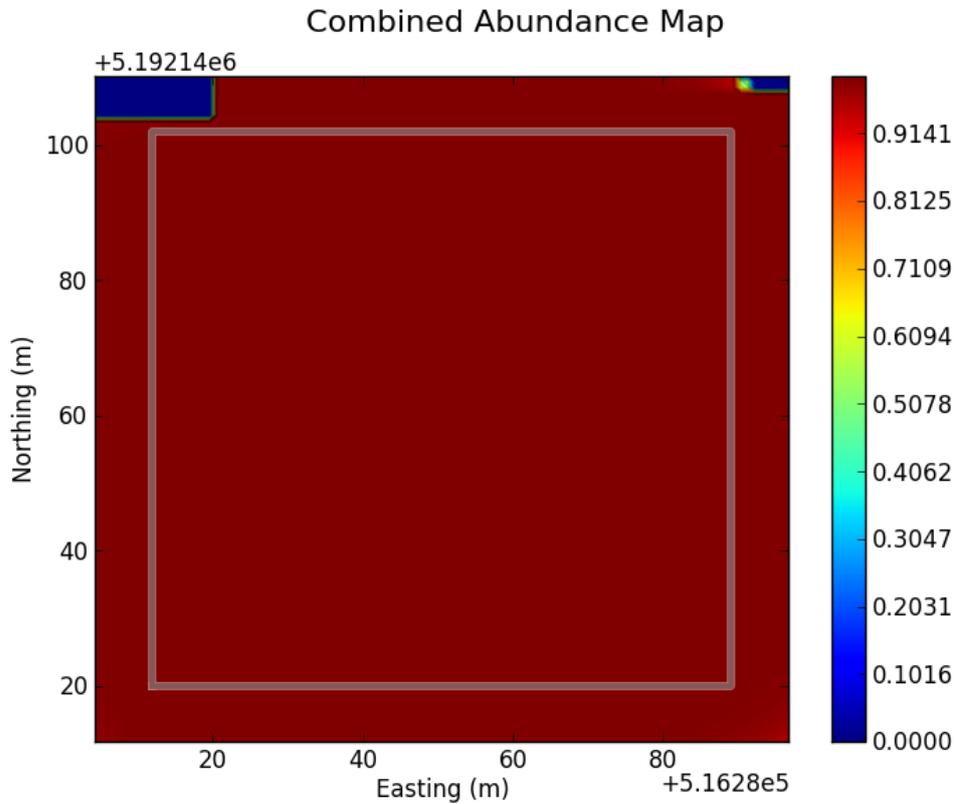


**Figure 6.6:** Cloudy Bay Lagoon seafloor depth map.

The map shown in Figure 6.6 displays seafloor depth from *Searise* data collected during the trial mission. The colour scale represents depth in metres. The map is reconstructed from the PA500 altimeter and P51-100 depth sensor data, which are summed to give seafloor depth. The observation distribution algorithm processed seafloor depth observations (see Section 5.4), compensating for uncertainty in the vehicle’s position.

In this and the following figures, the inset box shows the extremes of the area covered by the commanded path of the vehicle (see Table 6.2).

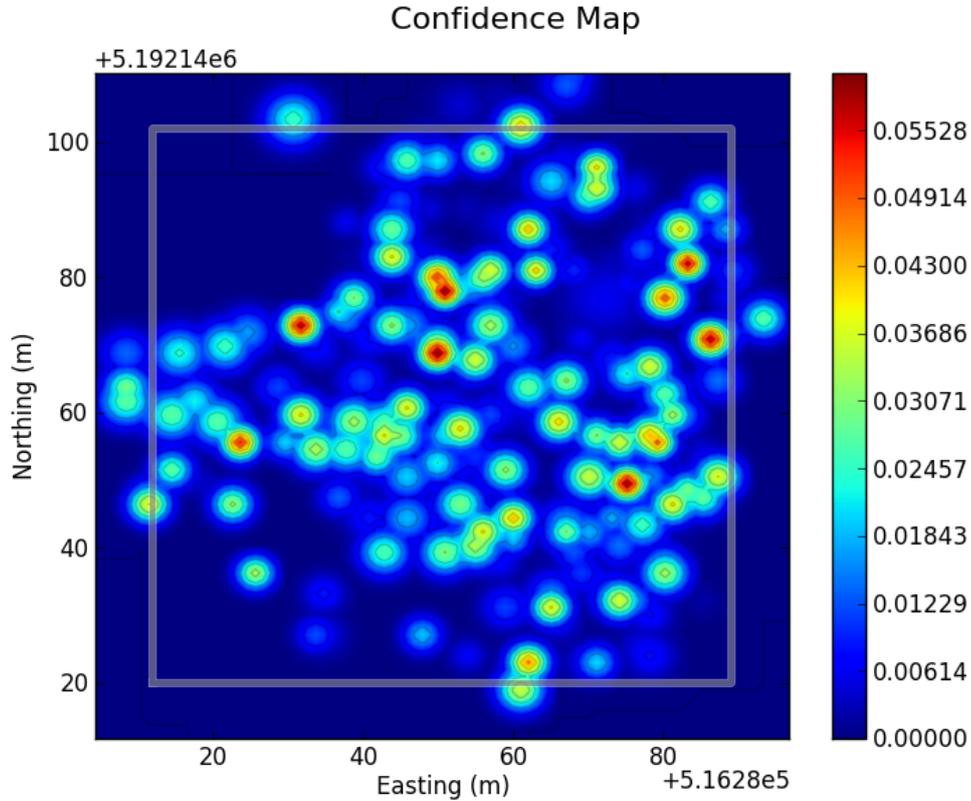
6.3.2 Combined Abundance Map



**Figure 6.7:** Cloudy Bay Lagoon combined abundance map.

The combined abundance map (see Figure 6.7) is used for validation that the observation distribution algorithm has operated correctly. This fairly bland, yet very important map confirms that the distribution of observations has resulted in a unity sum for all layers and map cells. The colour scale represents the sum of all layer cell values for any position in the map. The uniform sum to 1 is confirmation of the requirement demonstrated in Equation 3.3.3, Figure 5.6 and Figure 5.7. The blue areas at the top indicate map cells that received no distribution of observation data and are empty of data.

### 6.3.3 Confidence Map

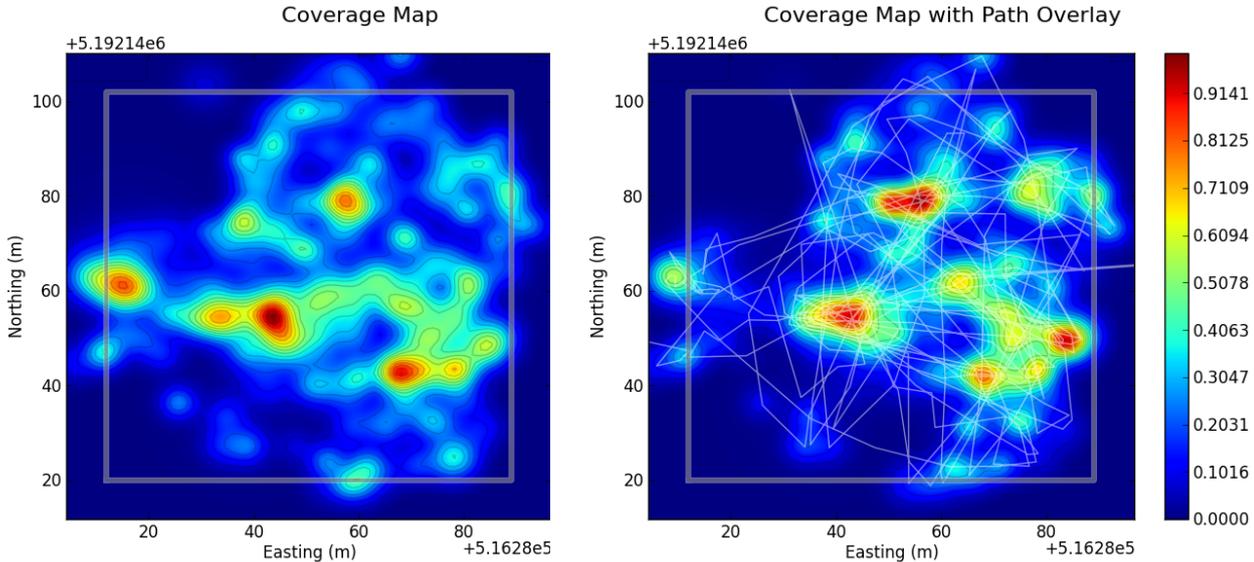


**Figure 6.8:** Cloudy Bay Lagoon confidence map.

Figure 6.8 is a confidence map for the Cloudy Bay Lagoon mission. This graph displays the maximum fitness value of unmix solutions at each location (i.e., the worst unmix result). The colour scale represents the fitness of observations, as explained in Section 4.5. Higher values are less fit unmix solutions. The apparent random distribution of poor-fitness areas, the low fitness values (indicating good unmix results) and lack of clustering suggests that the unmixing has not missed any major end-member groupings.

If clusters of poor fitness were present, this might suggest the presence of unrecognised end-members or greater variability within detected end-members, and these areas could be targeted for further investigation.

## 6.3.4 Coverage Map



**Figure 6.9:** Cloudy Bay Lagoon coverage map and path overlay.

The coverage map (see Section 5.3.1 and Figure 6.9) represents how well each map cell has been visited. It displays the count for each map cell, summed from the individual heights of the Gaussian distribution of observations to the cell (see Section 5.4). The colour scale indicates relative coverage proportions; the most-visited cells have a coverage value of one, and the least-visited cells have a coverage value of zero. This map can be used to plan future mapping transects to optimise science recovery by visiting less-well covered areas (see Section 7.4).

The map at right in Figure 6.9 has *Searise*'s path overlaid in white. Apparent jumps and discontinuities in the path are due to variation in reported global positioning system (GPS) position as *Searise* resurfaces, and sub-sampling of the path's data points.

### Caution About Interpretation

Caution should be exercised in placing too much credence in the classification of areas that have sparse coverage. The distribution of observations with high  $\sigma$  can lead to an observation being spread over a large physical area. At the extremes of this distribution, the Gaussian height can be very small, and the consequent fractional distribution of the observation to map cells at these extremes can also be

small. In cases where an observation is the only contribution to the classification of a map cell, even an extremely small Gaussian contribution is enough to cause that map cell to be classified as 100% of the given object type (e.g., see Figure 6.18 showing an isolated area of classification on the right side due to inaccurate GPS position data). This can be misleading and inaccurate. Mapping classification is likely to be inaccurate when the total Gaussian fractions distributed to the cell are small.

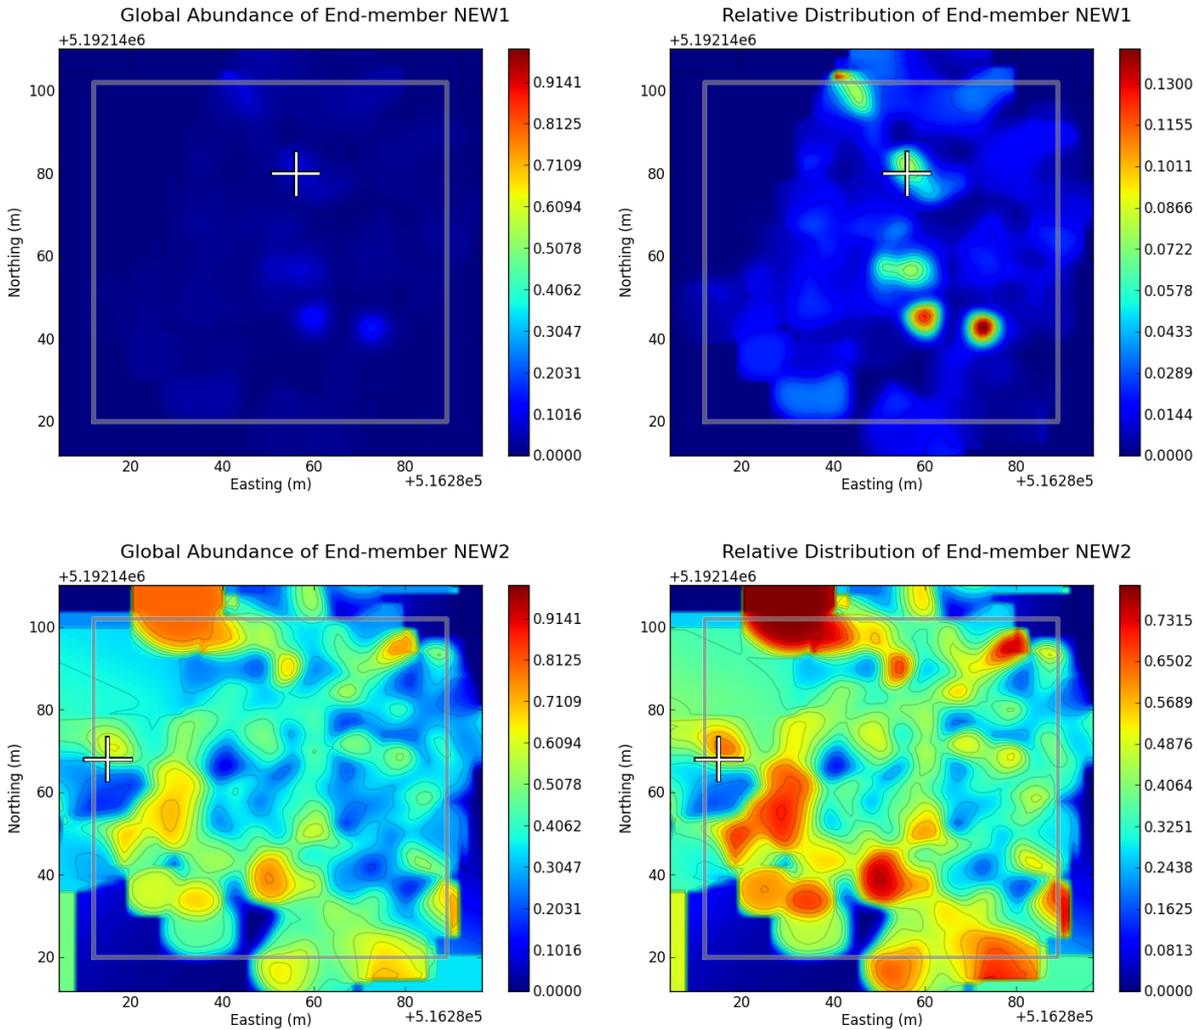
### 6.3.5 End-member Abundance Layers

The end-member abundance layers display the abundance of each of the end-members detected by the classification algorithm. The white cross on each plot is the location of the automatically detected end-member holotype. This location is where the worst match between observed spectra and remixed spectra occurred. It does not necessarily lie over the highest abundance of the end-member, as the abundance shown is affected by surrounding observations.

#### Global Abundance and Relative Distribution

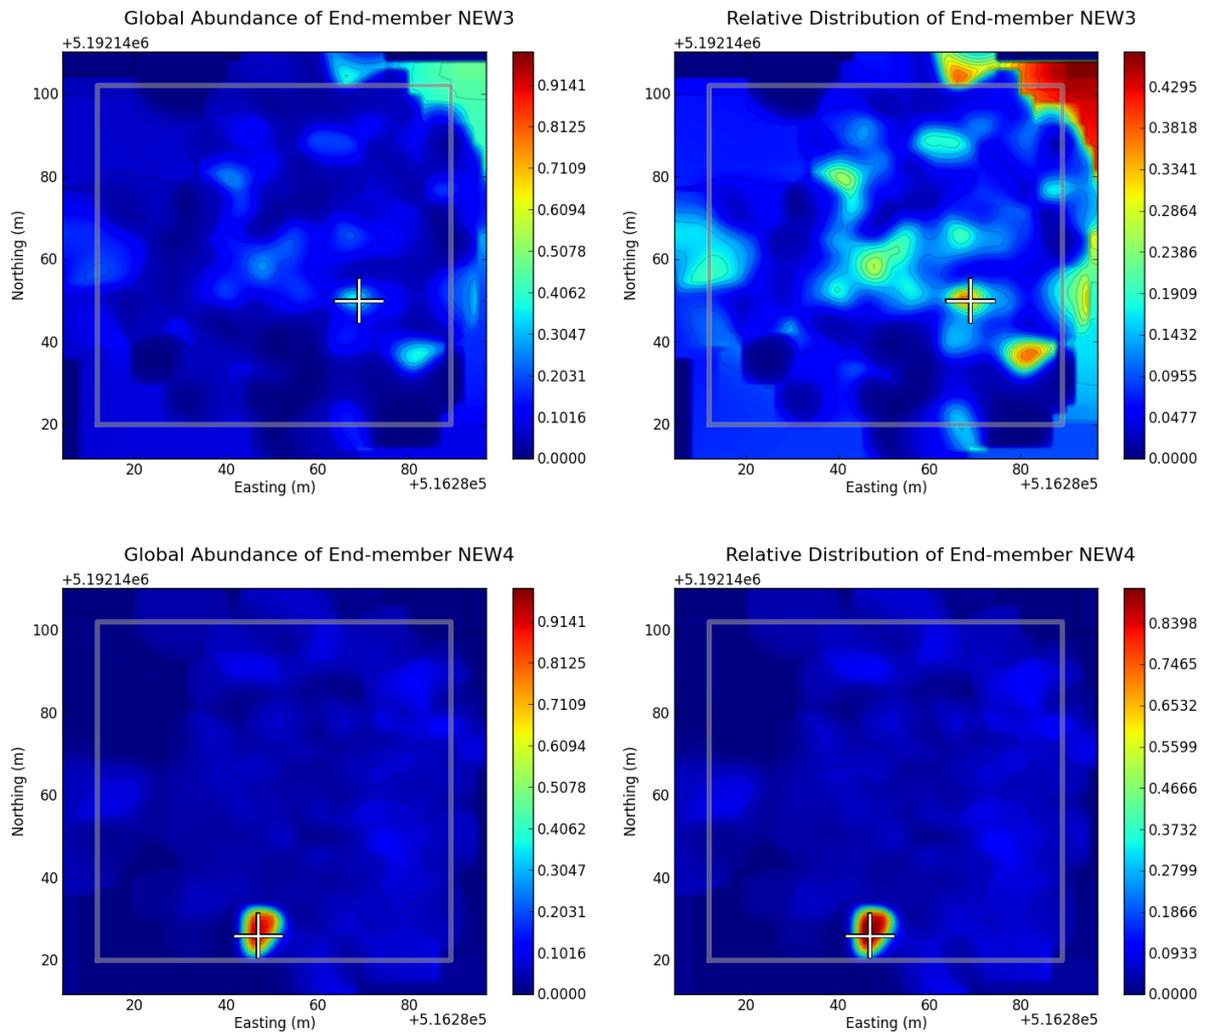
The global abundance and the relative distribution layers are two views on the same data set, with different scale ranges. The relative distribution is useful for determining where an end-member is most concentrated (the scale ranges from the minimum to the maximum expression of the end-member). The global abundance is useful for determining the concentration of an end-member relative to all other end-members (the scale ranges from 0 to 1). These two views on the same data are useful, as when an end-member has very low abundance it will display with very little colour variation in the global abundance layer and appear as mostly a single colour, or perhaps not be visible at all. The relative distribution layer optimises the scale and gives a better visual of the distribution of the end-member, regardless of its contribution to the abundance map.

### 6.3. THE MISSION AT CLOUDY BAY LAGOON



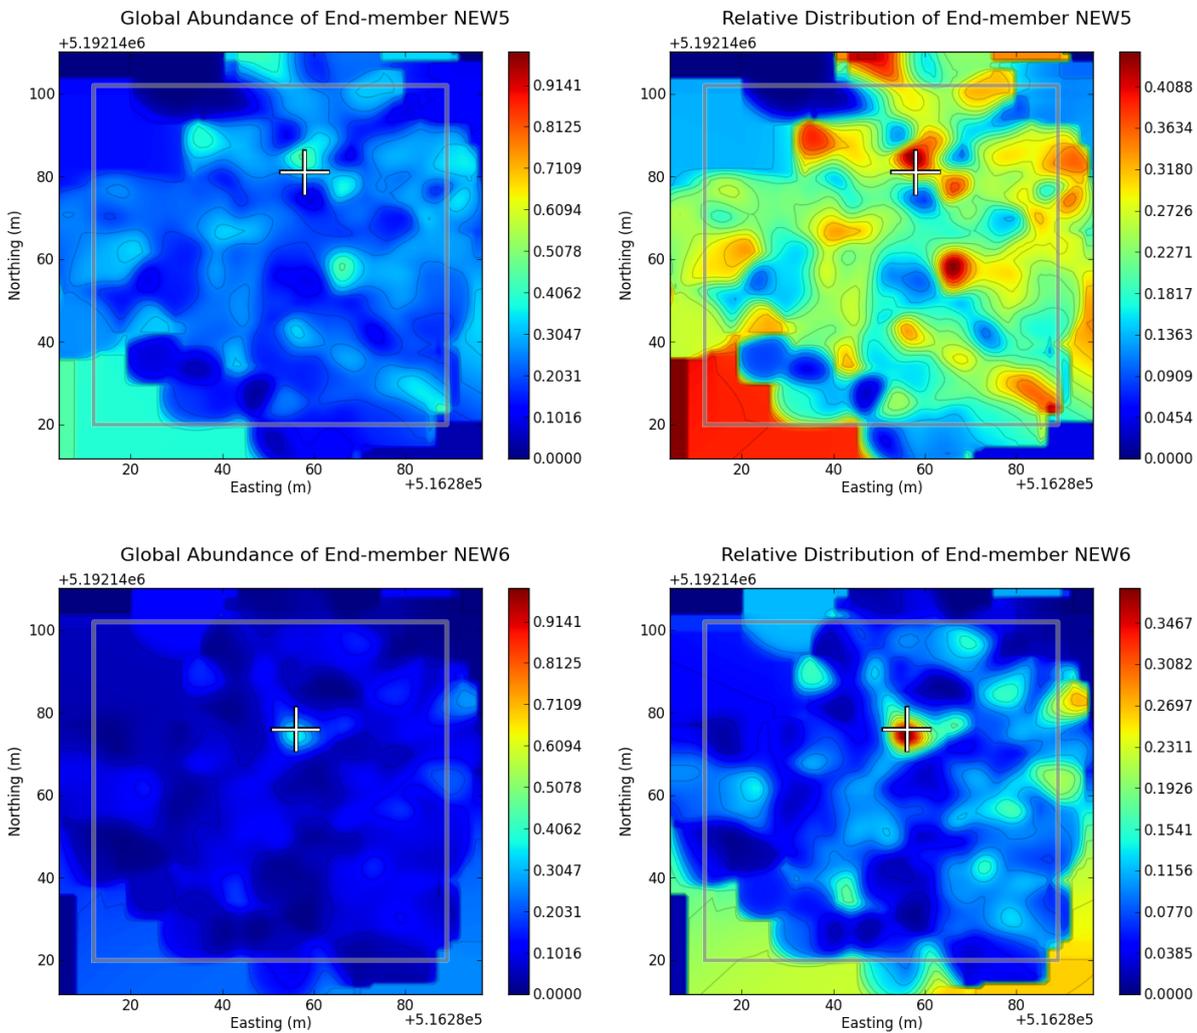
**Figure 6.10:** Abundance and distribution layers for end-members NEW1 and NEW2. Abundance estimates for each end-member are shown in Table 6.4. The white cross shows the automatically determined location of the end-member's holotype. These layers must be considered in context with the coverage map (see the caution in Section 6.3.4).

## CHAPTER 6. FIELD TESTING



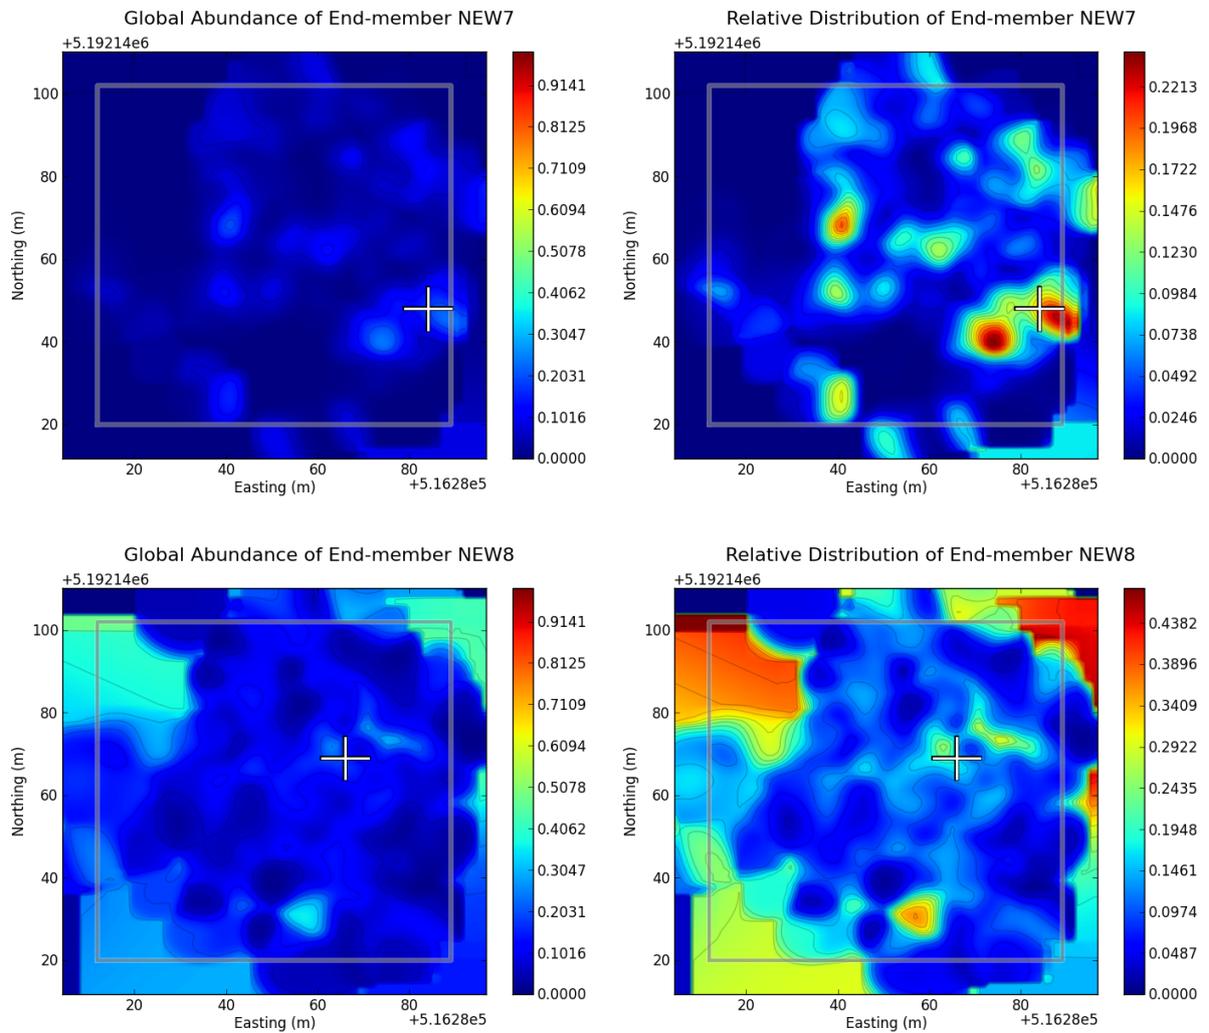
**Figure 6.11:** Abundance and distribution layers for end-members NEW3 and NEW4. Abundance estimates for each end-member are shown in Table 6.4. The white cross shows the automatically determined location of the end-member's holotype. These layers must be considered in context with the coverage map (see the caution in Section 6.3.4).

### 6.3. THE MISSION AT CLOUDY BAY LAGOON



**Figure 6.12:** Abundance and distribution layers for end-members NEW5 and NEW6. Abundance estimates for each end-member are shown in Table 6.4. The white cross shows the automatically determined location of the end-member's holotype. These layers must be considered in context with the coverage map (see the caution in Section 6.3.4).

## CHAPTER 6. FIELD TESTING



**Figure 6.13:** Abundance and distribution layers for end-members NEW7 and NEW8. Abundance estimates for each end-member are shown in Table 6.4. The white cross shows the automatically determined location of the end-member's holotype. These layers must be considered in context with the coverage map (see the caution in Section 6.3.4).

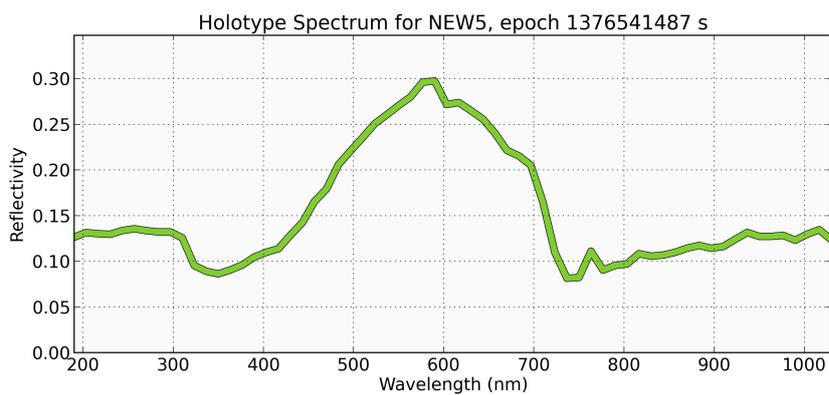
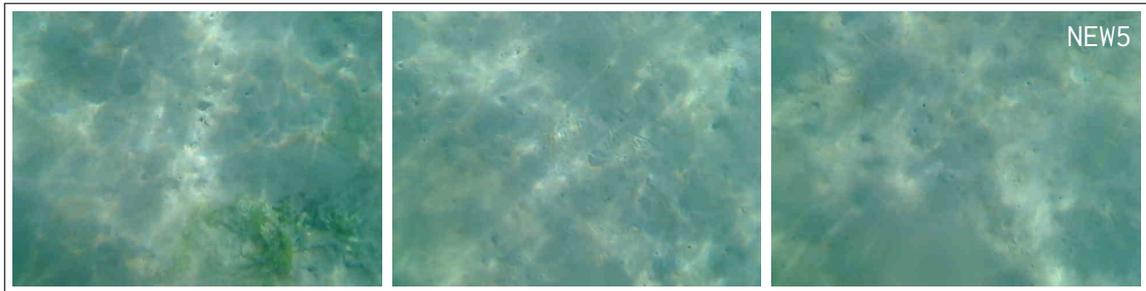
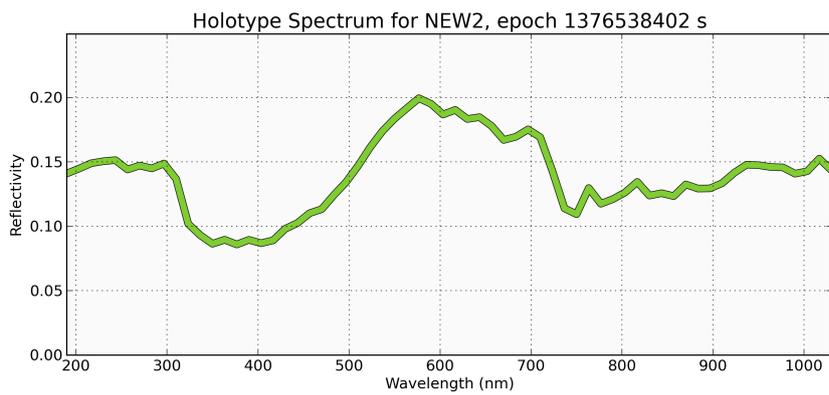
### 6.3.6 Holotype Images and Spectra

The unmixing algorithm detected eight composite end-members, represented by the holotype images and spectra shown in Figure 6.14 and Figure 6.15. These are displayed in order of abundance, as shown in Table 6.4. Every observation is reconstructed with optimal fitness using only these holotypes. The holotypes are representative of major groupings of habitat types (see Table 6.4 for some interpretations), and their distributions are shown in the earlier maps.

The location of each holotype is determined from the sample time of the observation that was detected as a worst fit during each iteration of the unmixing algorithm. The exact locations of holotypes are displayed on the abundance layers in Section 6.3.5.

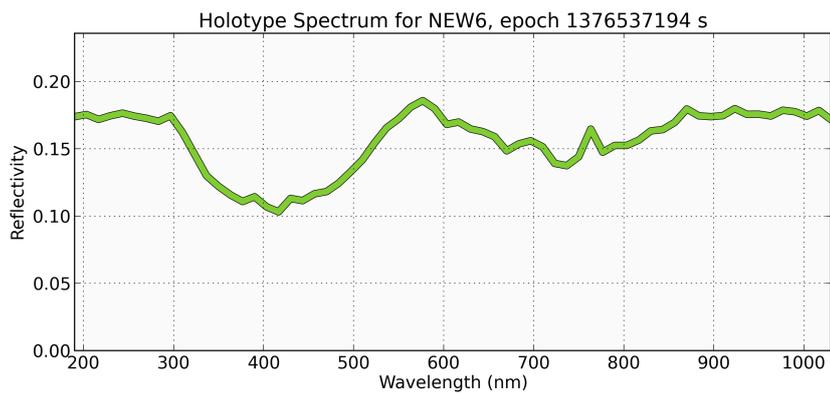
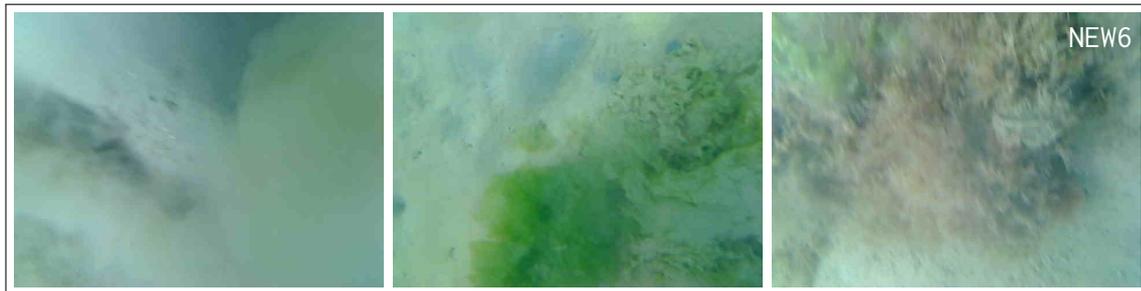
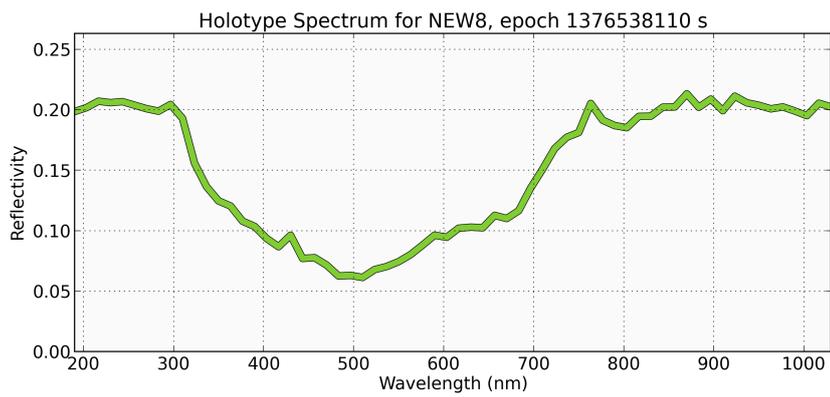
Each holotype's observation time, corrected for the camera's 20 s buffer, is used to find the most representative image of the holotype from the camera image data. There is unlikely to be a camera image at the exact time of a holotype's observation, as the camera was recording images at 0.2 Hz. The closest three images (in time) have been selected as the representative holotype images.

## CHAPTER 6. FIELD TESTING



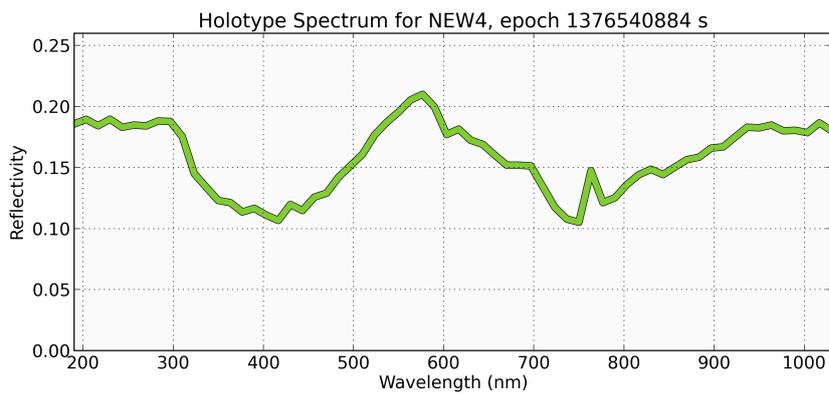
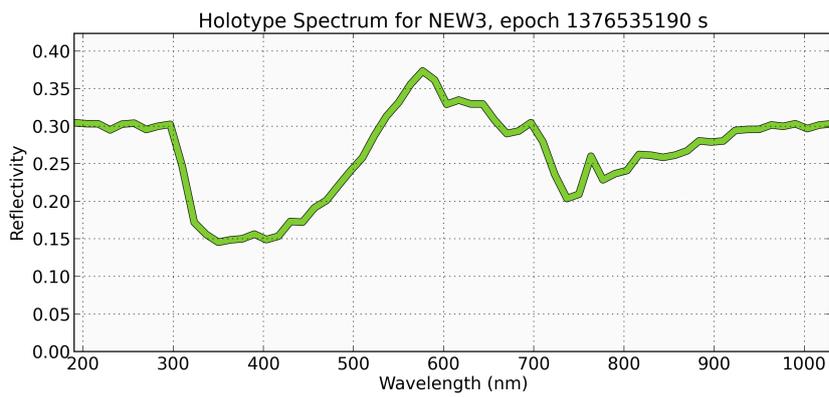
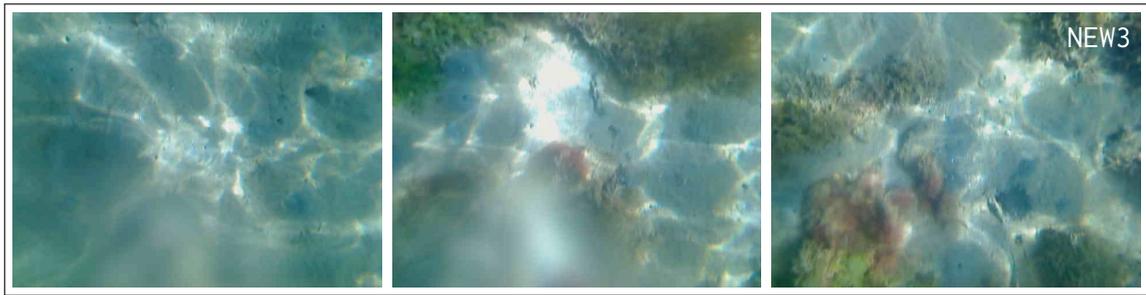
**Figure 6.14:** Images and holotype spectra for detected end-members NEW2 and NEW5. The unmixing algorithm determined the end-member spectra, and the associated camera images were retrieved using the holotypes' observation times. The time has been adjusted for the 20 s camera image buffer, and the three images taken closest to this time have been chosen.

### 6.3. THE MISSION AT CLOUDY BAY LAGOON



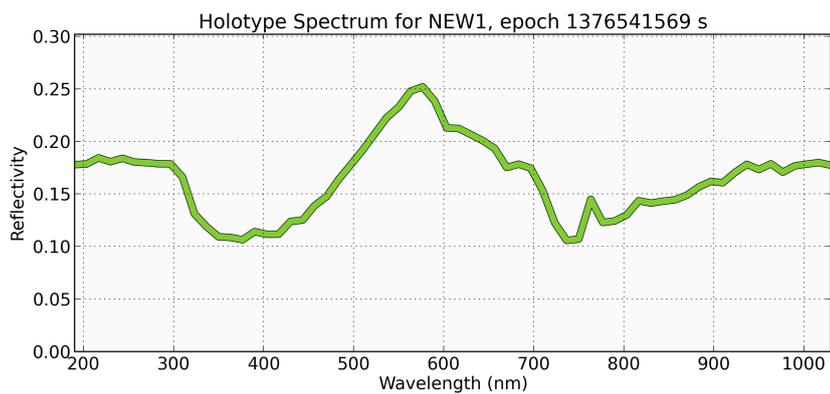
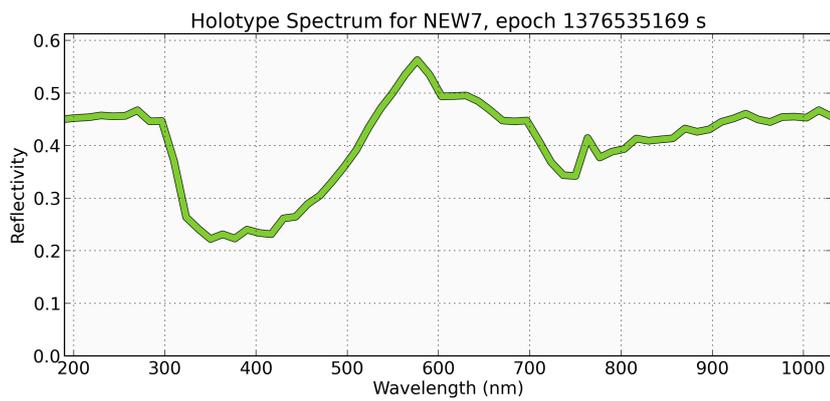
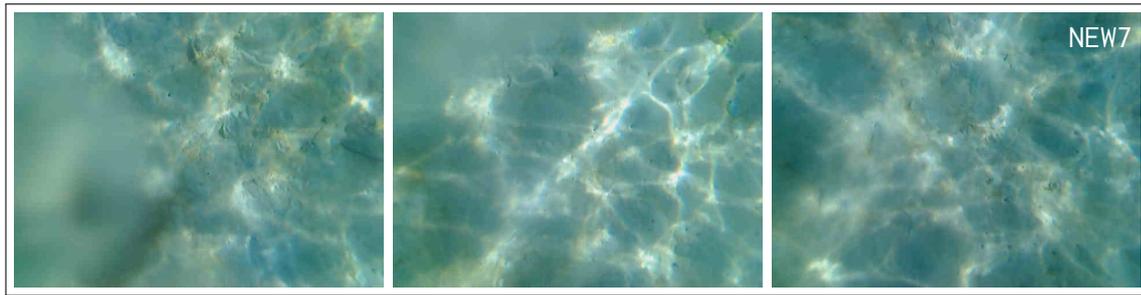
**Figure 6.15:** Images and holotype spectra for detected end-members NEW8 and NEW6. The unmixing algorithm determined the end-member spectra, and the associated camera images were retrieved using the holotypes' observation times. The time has been adjusted for the 20 s camera image buffer, and the three images taken closest to this time have been chosen.

## CHAPTER 6. FIELD TESTING



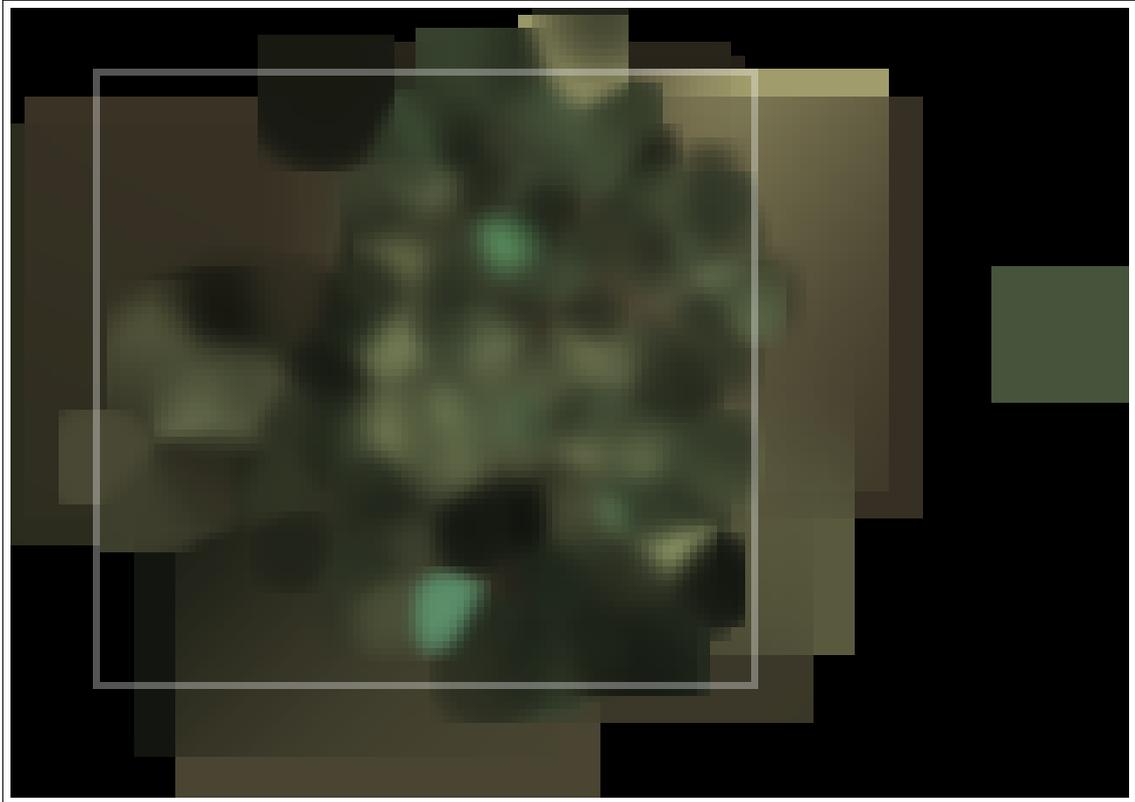
**Figure 6.16:** Images and holotype spectra for detected end-members NEW3 and NEW4. The unmixing algorithm determined the end-member spectra, and the associated camera images were retrieved using the holotypes' observation times. The time has been adjusted for the 20 s camera image buffer, and the three images taken closest to this time have been chosen.

### 6.3. THE MISSION AT CLOUDY BAY LAGOON



**Figure 6.17:** Images and holotype spectra for detected end-members NEW7 and NEW1. The unmixing algorithm determined the end-member spectra, and the associated camera images were retrieved using the holotypes' observation times. The time has been adjusted for the 20 s camera image buffer, and the three images taken closest to this time have been chosen.

### 6.3.7 Visual Reconstruction of the Seafloor



**Figure 6.18:** This image is a reconstruction of the seafloor, using the human-eye sensitive wavelengths red (650 nm) green (510 nm) and blue (475 nm) from the reflective end-member spectra, mixed in the same proportions as the end-members. This picture is not true-colour; the predominance of sand-type end-members is biased in intensity so that the vegetation is emphasised. The image extends past the mapped area due to the Gaussian distribution of observations (see Section 5.4). The large square area represents the same area as shown in the abundance layer maps presented earlier. The small square at right has occurred because of a very inaccurate GPS position (obtained just after surfacing); the square represents the area over which the observation at that point was distributed.

After the distribution of observations to the abundance layers, each layer cell (see Section 5.2.1) indicates the proportional contribution of the end-member it represents. The autonomous detection of end-members (see Section 4.9) has collected the spectra of the end-members. From these data it is possible to reconstruct an approximation of the visual appearance of the seafloor.

## 6.3. THE MISSION AT CLOUDY BAY LAGOON

---

### Determining True Colour

The human eye is sensitive to light in three wavelength bands, corresponding to red, green and blue. These wavelengths are at approximately 575 nm (red), 545 nm (green) and 475 nm (blue). Sensitivity of the eye to these colours peaks at the given wavelengths and reduces rapidly forming an approximate Gaussian for each colour.

The combination of these gives the spectrum of colours visible to humans.

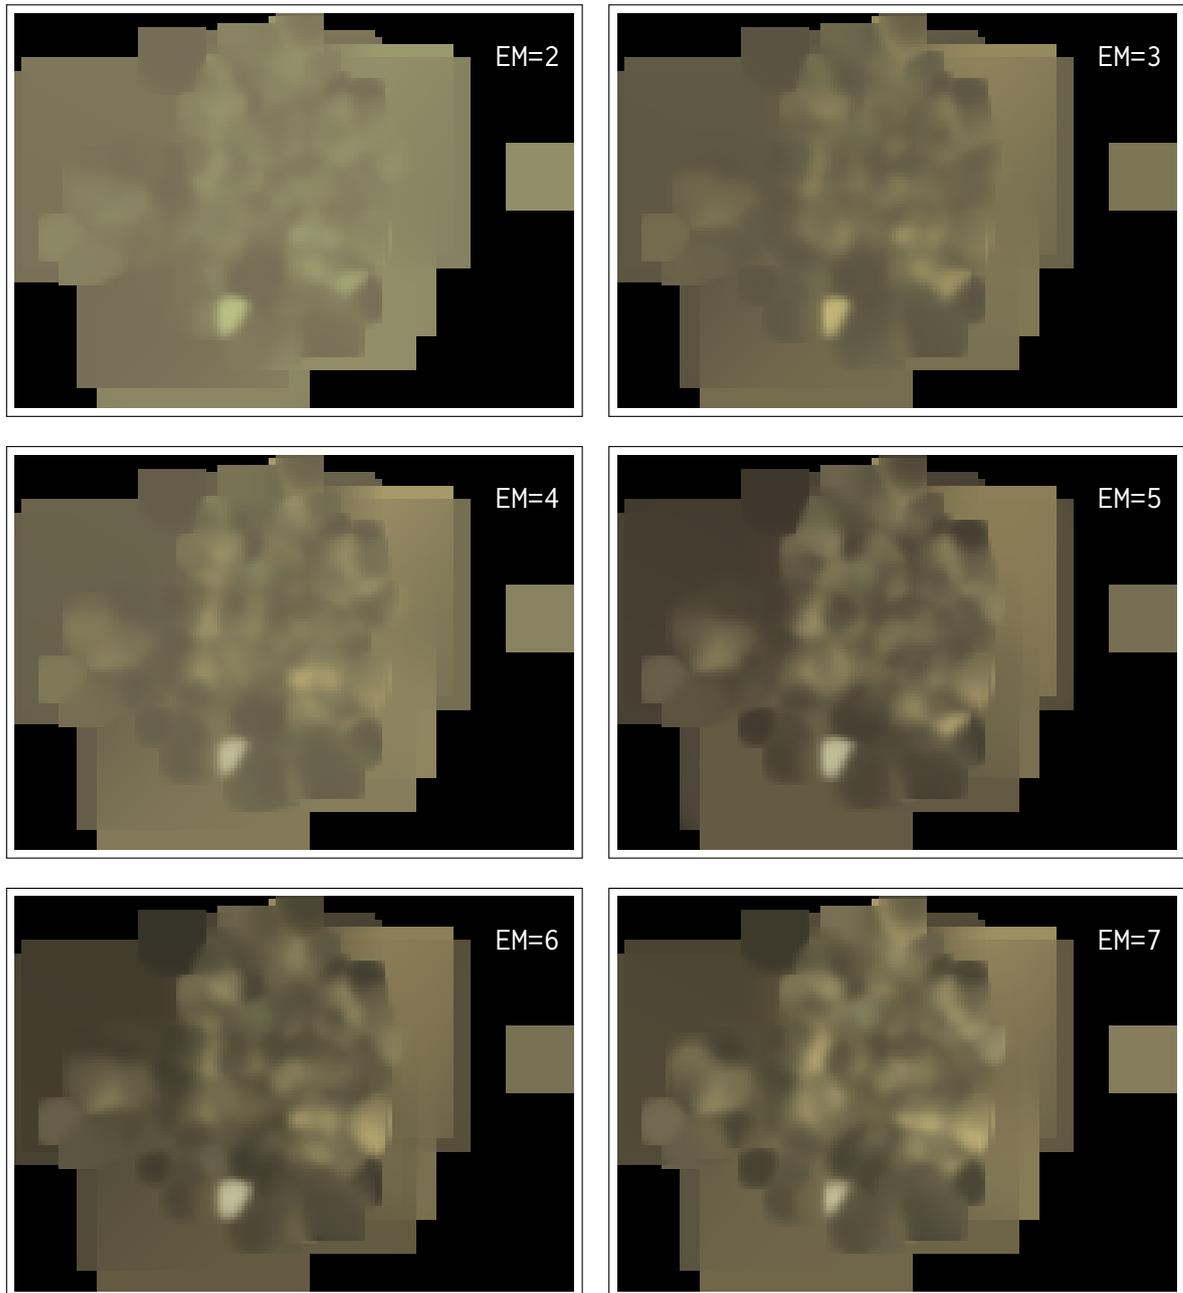
Each end-member spectrum is hyper-spectral, spanning a wide range of wavelengths (see Table 3.1 for details on the capability of *Searise*'s spectrometers). Selection of the red, green, and blue wavelengths from these end-member spectra and combining to form a RGB pixel in an image allows for a representation of the true-colour (as visible to a human) of the end-member. Combination of the true-colours of the end-members in the same proportion as the representation of the end-members in layer cells gives a true-colour representation of the area of the seafloor the layers' cells represent.

Figure 6.18 is a reconstruction of the seafloor from the Cloudy Bay Lagoon mission using the reflective spectra. The proportional contribution of end-members has been adjusted to reduce the prevalence of sandy substrates and emphasise vegetation. Patches of green evident in the image are directly a result of the end-member colours and distribution, but are only an indication of true-colour.

### 6.3.8 Comparative Abundance Map

The final set of maps presented here form a comparative abundance map (see Figure 6.20). Each of the first 8 images represents the abundance of a detected end-member type. A unique colour is used for each, and the scale is graduated so that the brightness indicates the isolated abundance of the end-member. The final combined map combines the 8 isolated abundance maps into a single image representing the locations of highest abundance of each end-member type.

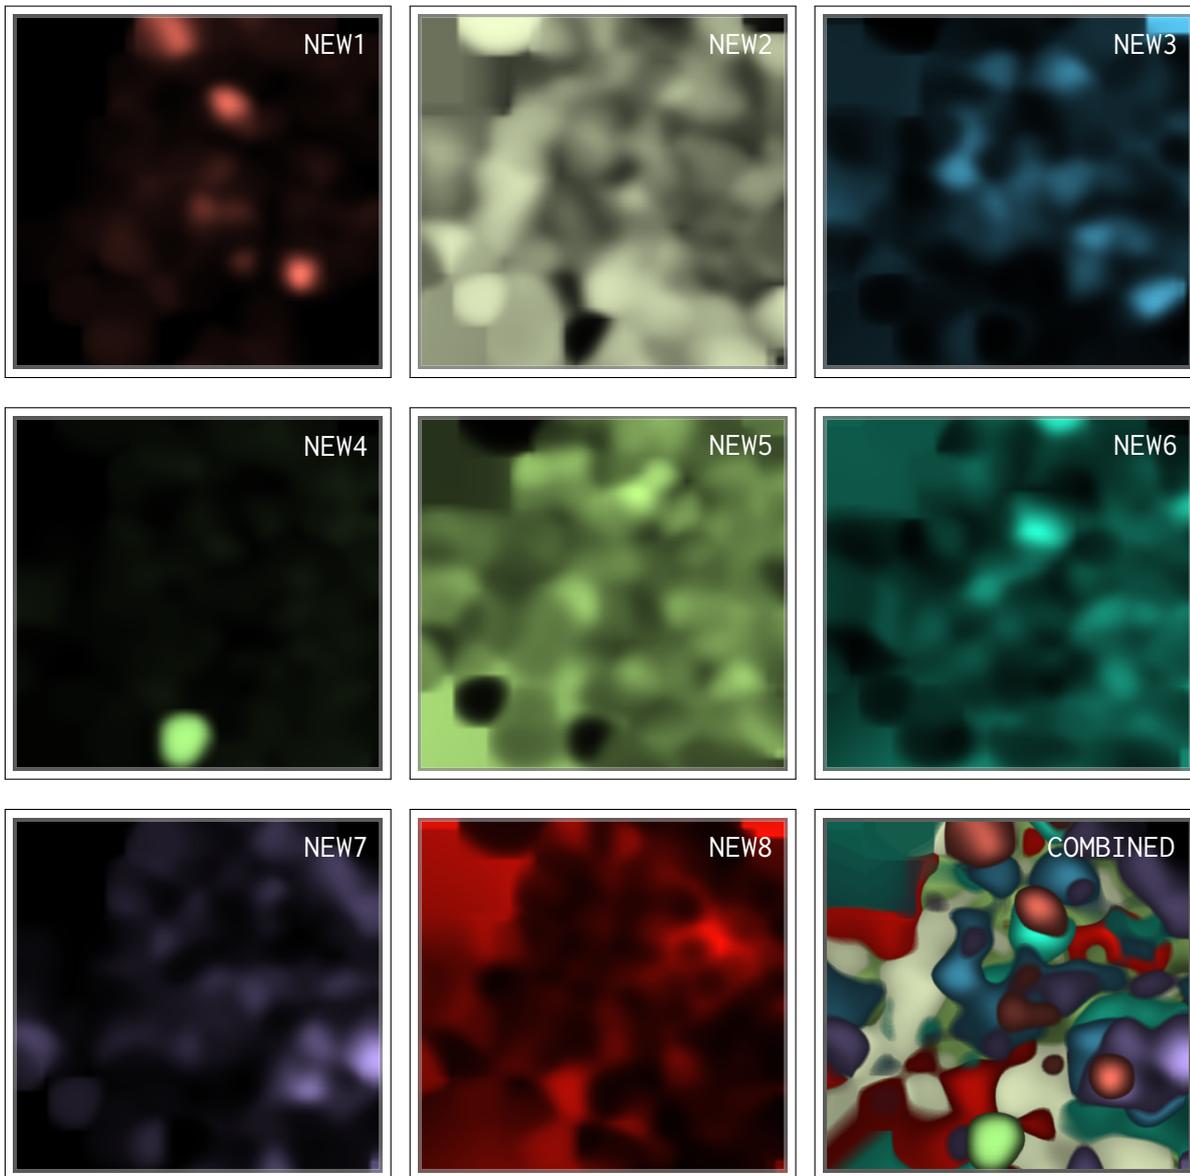
Caution should be exercised when interpreting this map. The abundance of some of the end-members is very small (see Table 6.4), but this is not reflected in the comparative map. Additionally, the map represents a single end-member type at any location, whereas typically there is a mix of end-members at any location. The map is best considered as an indication of where best to find any particular end-member (regardless of how prevalent that end-member is, compared to other end-members at the location).



**Figure 6.19:** Reconstruction with differing end-members. The reconstructions shown here are from unmix runs of 2 to 7 end-members. As with Figure 6.18, the seafloor is reconstructed using only the end-member spectra, and the information about the proportions of each of these end-members at each map location. The similarity between the reconstructions show that detected features are independent of the number of end-members, but the detail improves with end-member cardinality.

### 6.3. THE MISSION AT CLOUDY BAY LAGOON

---



**Figure 6.20:** Each of the eight end-members is represented here as a different coloured distribution map. The maps indicate the abundance of each end-member independent of other end-members. The final image combines the areas of highest abundance from each of the end-member distribution maps into a comparative abundance map. This shows the locations of highest abundance of each end-member type, but does not account for the relative abundance of each type.

### 6.4 Summary

---

This chapter presented the results of a *Searise* mapping mission to Cloudy Bay Lagoon. The data were collected and analysed using software developed by the author.

A series of maps were presented, showing the distribution of automatically detected end-members, bathymetry, confidence and coverage. Also shown were end-member spectra and geo-tagged images showing the location and visual form of detected end-members.

Finally, two overview maps were shown; a reconstructed map of the seafloor using true-colour from the detected end-member spectra, and a comparative abundance map showing the most likely locations to find any particular end-member.

The large variety of maps provides a comprehensive set of views of the mapped area. These demonstrate successful operation of the unmixing algorithm, the distributed observation mapping algorithm and the autonomous end-member detection and extraction algorithm.

The described systems perform automated classification and mapping of benthic habitat, achieving the primary research goal of this work.

*“Begin at the beginning,” the King said, very gravely, “and go on till you come to the end: then stop.”*

Lewis Carroll (1832-1898), *Alice in Wonderland*

*“By prevailing over all obstacles and distractions, one may unfailingly arrive at his chosen goal or destination.”*

Christopher Columbus (1451-1506)



## Review and Conclusion

### 7.1 Review of This Work

---

Monitoring of marine habitat is becoming increasingly important for the management of our marine environment. There is a need for more efficient and accurate systems to provide lower cost and improved accuracy of this monitoring. Traditional systems such as the interpretation of aerial and satellite photography have shortcomings in their inability to discriminate benthos on a fine scale.

This work has investigated a number of questions related to the automated classification and mapping of benthic habitat.

The autonomous underwater vehicle *Searise* was used for this research. Software developed for control of the vehicle and instruments provided capability for this machine to gather data for habitat mapping.

Unmixing spectral data to identify constituent end-members is a well-known technique. This method relies on a library of end-member spectra representing each object type to be identified. Light is attenuated by water, so it is difficult to use these unmixing techniques for underwater observation, as the observed spectra of end-members vary according to the attenuation of the light. In an attempt to correct this, various ‘de-attenuation’ algorithms are typically applied before unmixing.

Instead of attempting to correct for attenuation, this work uses a technique for the automatic detection of end-member groupings, allowing the un-mixing algorithm to optimise both the number of end-members detected and the accuracy of unmixing.

The sometimes-uncertain location of the autonomous underwater vehicle (AUV) when performing underwater measurement can also create some difficulty for the production of accurate habitat maps. The distribution of observations using a Gaussian spread function onto multi-layer maps is discussed as a technique for incorporating uncertainty into the mapping process.

The author developed software to perform the unmixing, automatic end-member detection, observation distribution onto multi-layered maps, and produce multiple map views of the mapped area. The system was successfully tested with a mission to Cloudy Bay Lagoon, Bruny Island by the autonomous underwater vehicle (*Searise*). The resulting maps show the algorithms and software are working successfully, producing habitat maps of the study area.

## 7.2 Revisiting the Research Questions

---

### 7.2.1 Unmixing Spectral Data

Spectral end-member libraries can be a suitable technique for accurate unmixing of hyper-spectral data, but the end-member spectra of these libraries must be in a normalised form independent of the light source (e.g., reflective spectra).

A number of techniques (e.g., linear algebra, non-negative least squares, quadratic programming, simulated annealing, genetic algorithms) have been shown to be effective for unmixing this spectral data. Linear algebra provides a mathematically perfect and quick method, but is not appropriate as it produces results not subject to the non-negativity and unity constraints. Iterative techniques such as quadratic programming, non-negative least squares, and genetic algorithms provide methods for enforcing these constraints and are suitable for unmixing hyper-spectral data.

The method implemented in this work is a memetic algorithm (a combination of genetic algorithms (GAs) and simulated annealing (SA) algorithms). This method is shown to be effective in unmixing hyper-spectral data and the implementation presented is adapted to autonomously detect end-members and create unsupervised classification of benthic habitat. It has proven to be successful at accurately unmixing underwater spectra to recover estimates of the proportions of contributing end-members.

The accuracy of unmixing is determined by reconstructing spectra using the unmix proportions and the end-member library, and comparing the similarity of the reconstructed spectra to the original observation.

## 7.2. REVISITING THE RESEARCH QUESTIONS

---

The techniques presented are shown to be robust and able to cope with both noise and unknown object types.

### 7.2.2 Autonomous End-member Detection

A pre-defined end-member spectral library is insufficient to achieve accurate unmixing. Such a library does not account for variations in the spectra of end-members, such as seasonal changes and differences between young and old plants. An algorithm has been presented here that creates an autonomously generated end-member library that allows multiple end-member spectra to represent single end-member types.

The use of a fitness measure for determining classification accuracy enables the detection and addition of new composite end-members to the end-member library, bypassing the problem of spectral variability within end-member types.

A high fitness value implies the presence of at least one unknown end-member, or significant spectral variance of an existing end-member. By allowing the system to append new end-members to its library, variations in spectra (both seasonal and inter-species) can be catered for, as can previously unknown end-member types. Post-classification analysis of the end-member library by a human provides a simple and expedient method for ensuring that variance is accurately accounted for in the mapping process.

### 7.2.3 Autonomous Mapping

This work has demonstrated that these processes produce useful multi-layer maps that are consistent with ground observations. In addition, the process produces a coverage map that represents the amount and accuracy of the mapping process for any given area in the map. This coverage map can be used to guide future mapping passes and improve the quality of the mapping data.

The distribution of observations to the maps using a Gaussian spread function is an effective way to account for errors in estimated vehicle position when producing maps.

It is not, however, feasible to create these maps in real-time on board the AUV. The developed system requires all observations to be available to determine the end-members in the library, and these are not available until a mission has completed. The processing requirements for the unmixing are also significant (a

single observation unmix can take a few seconds, but thousands of observations can take several minutes to process).

Some elements of the mapping process can be performed in real-time. Creation of the coverage map (useful for real-time navigation decisions) is one of those.

### 7.3 Contribution

---

- Unmixing Spectral Data
  - Attenuation/approach with up/down spectrometers and proportional representation of spectra
  - Application of unmixing to underwater data
- Autonomous End-Member Detection
  - Bypass identification of pure end-member spectra through detection of unknowns
  - Use of expert post-classification of end-members to provide pure end-member spectra ID
- Autonomous Mapping
  - Development of control systems for *Searise*
  - Multi-layer classification maps, including coverage map via Gaussian estimation

### 7.4 Further Work

---

#### Ground-Truth for Segmenting Classifiers

One of the outputs of the processes described in this work is a map showing the relative proportions of each end-member in the mapped area (see Section 5.3.3).

These proportion data may be useful as an input into classifiers that use texture and edge-based division of imagery to isolate and identify areas of similarity. The a priori proportions of objects can be used as a ground-truth against which the tuneable parameters used to control these segmenting algorithms could be adjusted so that their output segmentations are in agreement.

### Optimisation of Path using Coverage Map

The coverage map (see Section 5.3.1 and Section 6.3.4) can be used as an aid to determining optimal path to gather the most scientifically valuable data. The optimal path would balance visiting least-visited areas against cost (time/distance) to visit the areas and expected science return. This is essentially a travelling salesman problem.

A simpler method is to evaluate finite pre-calculated candidate mapping survey path alternatives for optimal science value based on an inferred Gaussian probability map of the survey field (Rigby et al. (2010)). The coverage map described here is an analogous representation, and can be used for the type of optimal path selection described.

### Cooperative AUV Mapping

The relatively slow movement of *Searise* and the need to provide fine-scale coverage of areas being mapped restricts the ability of the vehicle to map large areas quickly. The coverage map lends itself to use by multiple AUVs in cooperative mapping tasks, allowing the allocation of transects among multiple vehicles to efficiently map areas of interest.

### Mapping of Mid-Water Absorption

The use of upward-pointing and downward-pointing optical spectrometers allows for the calculation of water absorption in the mid-water column. Mark Baird (CSIRO Marine and Atmospheric Research) has used *Searise* spectrometer data to produce estimates of water absorption at varying wavelengths and depths. Similarly, the unmixing algorithms presented in this work could be used for autonomous detection of algae in the mid-water column.

## 7.5 Conclusions

---

The techniques presented in this work provide a practical method for using AUVs to map benthic habitat. The autonomous unmixing of spectra into constituent end-member proportions, the identification of end-members for addition to an end-member library, and the distribution of observations onto multi-layer maps based

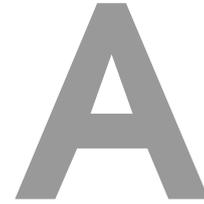
## CHAPTER 7. REVIEW AND CONCLUSION

---

on the uncertainty in vehicle position are all useful components of the classification and mapping system presented here.

This work has shown that autonomous unmixing of spectra representing complex mixes of benthos is practical and can be accurate. The examples presented in this work show an excellent correlation between the original spectral mix and the unmixed solutions calculated by a memetic algorithm. The unmixed spectra represent the proportion of end-members contributing to observations, and these were used to produce layered abundance maps, as demonstrated in Chapter 6. Additional layers representing confidence in observations, and coverage, can assist in planning future mapping of the observation area.

Combined together, the techniques presented in this work offer a practical method for autonomous classification of benthic habitat.



# GPS Algorithm

## A.1 GPS Accuracy

---

The recent availability of low-cost global positioning system (GPS) receivers has seen their increased use in a wide variety of vehicles, including submersibles like *Searise*. Shortcomings of these GPS devices include their ability to return position data only when the vehicle is surfaced, inconsistent accuracy (depending on the relative geometry of detected satellites and the time the vehicle has been submerged), and their inability to obtain an accurate ‘fix’ in challenging environments. When submerging and surfacing, the regular loss and re-acquisition of signal and the consequent variations in precision of position data make accurate estimation of position difficult.

The use of low-cost GPS receivers on board *Searise* imposes such limitations on the estimates of the AUVs position as discussed above. When underwater, *Searise* estimates its position using dead reckoning; the heading and thruster activity are used to estimate changes in position.

*Searise* has three inexpensive GPS receivers installed in its tail, which is approximately 30 cm above the water when the vehicle is surfaced<sup>1</sup>. Some GPS units have a relatively large degree of positional uncertainty when an initial fix is achieved. This may be because of the environmental conditions close to the water surface (e.g., multipath reflections), and slow satellite re-acquisition by the GPS

---

<sup>1</sup>When waves are of this magnitude, satellites low on the horizon are not detected by the GPS units. Signals from GPS satellites can also reflect off water and waves, causing problems with multipath interference.

## APPENDIX A. GPS ALGORITHM

---

units. Some GPS units are more accurate but do not rapidly achieve fixes in rough conditions, and some tend to ‘drift’ over time, based on satellite geometry. All units behave differently based on internal cold-start, warm-start, and hot-start states<sup>2</sup>. These differing shortcomings and behavioural characteristics make it difficult to select a single unit as the ‘best’ at any point in time, and can adversely affect a simple averaging algorithm.

Techniques for improving GPS accuracy have been proposed (e.g., Guo et al. (2001), Hasegawa and Yoshimura (2003)) but these generally require differential GPS, or additional sensors. This section outlines an approach that improves position estimates through the use of only the GPS receivers and a biased averaging algorithm. The algorithm uses the GPS units’ reported accuracy, and proximity to a pre-calculated position, to calculate a bias used when calculating an averaged position.

Characteristic of the GPS sensors used in this research is that when first acquiring a fix, a very large inaccuracy in position estimate (in the order of hundreds of metres) is typical. Only after a few seconds are the positional estimates more reliable. Operation of these sensors in an environment where GPS signal is regularly lost and re-acquired (when, for example, diving and surfacing an AUV or low in the water with multipath interference to satellite signals) means that the reported GPS position can vary wildly over a short time period.

Consequently, some sort of averaging or biasing of the signals towards those more accurate sensors is necessary to improve accuracy when measuring with multiple units. Operation of Kalman filters on these values does not offer significant improvement, because such filters place equal weighting on ‘good’ and ‘bad’ position estimates.

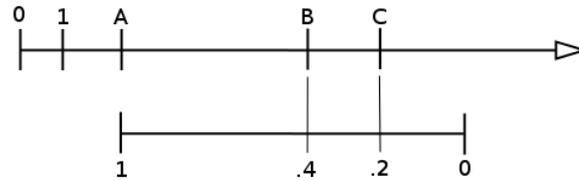
---

<sup>2</sup>GPS receivers take some time to obtain a lock on satellite signals, especially where the receiver is in a moving vehicle or satellite signals are obstructed. This time is dependent on the receiver’s state. There are three states: hot, warm and cold.

The hot state is when the GPS remembers its last calculated position, satellites in view, the information about the satellites in the constellation (the almanac), and the UTC time. In this state, the receiver makes an attempt to lock onto the same satellites and calculate position based on previous information. This is the fastest lock state but only works if the receiver is generally in the same location as when it was last active.

The warm state is when the GPS remembers its last position, almanac and UTC time, but not which satellites were in view. In this state a receiver attempts to obtain satellite signals and calculate a new position. It has a general idea of which satellites to look for because it knows its last position and the almanac identifies visible satellites in the sky. This process takes longer than a hot start.

The cold start is when the GPS loses all the information and attempts to locate satellites and calculate a GPS lock. This state takes the longest, as the receiver has to attempt to lock onto signals from any available satellites, which takes longer than using known satellites.



**Figure A.1:** Converting from HDOP or distance to fractional bias value. The accuracy of each GPS unit is used to bias the contribution of the unit to the averaged position. The top number line shows the HDOP values, and the bottom number line represents the conversion to a bias for each GPS receiver (A,B,C).

The algorithm presented is a computationally simple one that biases the contribution of those GPS sensors that are reporting more accurate positional data through their HDOP values. The algorithm performs a weighted average based on those HDOP values, as a proxy for accuracy.

Testing has indicated that the algorithm improves accuracy over that of a simple average, and is computationally simple enough to have no adverse effect on the limited computational resources of the vehicle.

### A.1.1 Dilution of Precision

Position data returned by GPS units are triangulated estimates of position calculated from time signals received from orbiting satellites. These position estimates vary in accuracy based on the number and relative position of the satellites within range of each GPS receiver. As an indication of the reliability and accuracy of positional data, GPS units report a dilution of precision (DOP) value for horizontal (HDOP) and vertical (VDOP) measurements. The DOP value is an analogue for positional precision, based on the geometry and availability of satellites. Lower DOP values indicate increased confidence in position, where 1 is the highest possible confidence level and higher values ( $> 20$ ) indicate that the GPS position may be very inaccurate.

### A.1.2 Algorithm

The algorithm performs a weighted average (each GPS position contributes disproportionately) based on two factors. Firstly, the HDOP value is used to bias in favour of GPS positions that have a higher confidence level, and secondly a pre-estimate of current position is used to bias in favour of GPS positions closer

to the pre-estimate. These biases are applied by using a multiplier for each GPS receiver's contribution to the weighted average.

### A.1.3 DOP Bias

The range of HDOP values can be considered on a number line, with good confidence at the left and low confidence values at the right. The goal is to convert a HDOP value into a useful fractional multiplier representing the contribution of a given GPS unit's HDOP to the averaged position, where 1 is the highest contribution and 0 is the lowest. Figure A.1 represents converting from a HDOP value (Equation A.1.1) to a fractional multiplier (Equation A.1.2). In the example figure, three GPS units are represented as A, B and C in increasing HDOP (decreasing confidence).

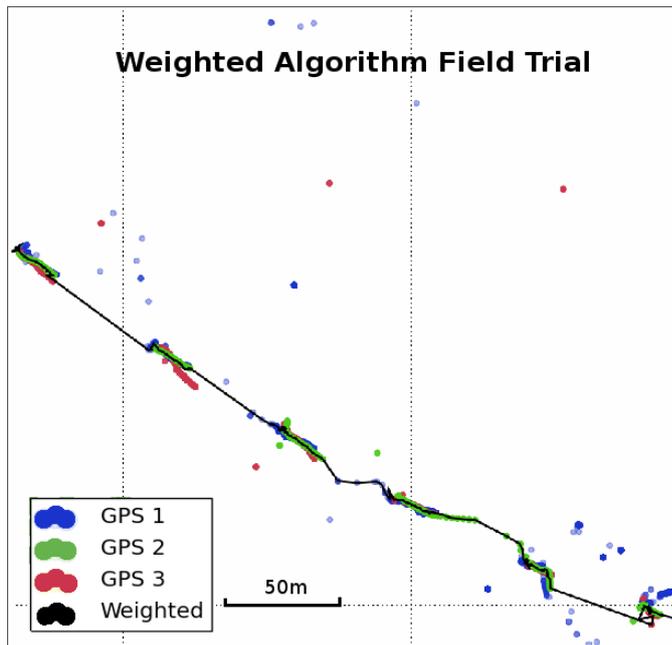
The range is calculated by determining the difference between best and worst HDOP values (in this case,  $C - A$ ). A multiplier is applied to the range to increase the contribution of the worst GPS HDOP, which would otherwise be zero in this system. This extended range is then used to produce a fractional position for each GPS on the number line, from 0 to 1 (calculated by  $(\text{HDOP}_i - \text{HDOP}_A)/\text{range}$ ). In the example in (Figure A.1), A is 0.0, B is 0.4 and C is 0.2. The number line is flipped by subtracting the fractional multiplier from one, as it is desirable to have higher contributions from GPS units with lower HDOP (i.e., with more confidence).

The optimal multiplier was determined empirically to be 1.44. This value minimised the variance of the calculated weighted average position for our GPS units. These calculations are formalised in the Equations A.1.1 and A.1.2, where  $\rho$  is HDOP,  $i$  is the GPS unit,  $m$  is the range multiplier and  $range_\rho$  is the amplitude over HDOP values for all  $i$ .

$$range_\rho = m * (\rho_{worst} - \rho_{best}) \tag{A.1.1}$$

$$bias_i(\rho) = 1 - (\rho_i - \rho_{best})/range_\rho \tag{A.1.2}$$

The *best* and *worst* refer to the magnitude of the HDOP values returned by the GPS units.



**Figure A.2:** GPS positions and calculated position from a deployment using the biased averaging algorithm. The coloured points indicate the position returned by the associated GPS unit. As the vehicle resurfaces, initial position estimates are inaccurate. The averaged position (black) provides a stable positional estimate.

### A.1.4 Proximity Bias

Because of the relatively high GPS sample rate compared to vehicle movement ( $< 2$  m/s), successive position estimates have relatively small changes. A simple predictive estimate of the position of the vehicle can be made using previously calculated positions and assuming no acceleration. This predictive estimate of position can be used as the ‘seed’ value for calculating distance-based weightings for each GPS receiver. The principle is that GPS units returning positions closest to where the AUV is most likely to be (calculated from previous estimates) are given more bias in contributing to the final estimate of position.

This calculation is similar to the bias for DOP, where  $\delta$  is the distance from the estimated position,  $i$  is the GPS unit and  $m$  is the range multiplier. The same range and multiplier factors are used for both HDOP and distance biases in this implementation.

$$range_{\delta} = m * (\delta_{worst} - \delta_{best}) \quad \text{(A.1.3)}$$

$$bias_i(\delta) = 1 - (\delta_i - \delta_{best}) / range_{\delta} \quad \text{(A.1.4)}$$

### A.1.5 Weighted Average

A final bias value for each GPS unit is calculated by multiplying the HDOP (Equation A.1.2) and distance (Equation A.1.4) bias values.

$$bias_i = bias_i(\rho) * bias_i(\delta), \quad \text{(A.1.5)}$$

where  $i$  is the GPS unit number. If a GPS unit does not have a fix then its bias is close, or equal, to zero and it does not significantly contribute to position estimates.

A weighted averaged position is then calculated from the coordinates and bias (Equation A.1.5) for each GPS unit: where *lat* and *long* are the coordinates of the position (easting/northing may be substituted), and  $n$  is the number of GPS receivers. This formula degenerates to a simple average when  $\forall i, bias_i = 1$ .

*Searise* uses GPS only when surfaced, and altitude is always assumed to be at present sea level, but for other uses the above calculation can be extended to the vertical component of position, using VDOP (vertical DOP) instead of HDOP.

## A.1. GPS ACCURACY

**Table A.1:** Statistics from the GPS averaging tests, showing the improvement in the averaged position after weighted-averaging of the three GPS units' data.

| Data Set         | $\sigma$ |
|------------------|----------|
| GPS 1            | 5.99     |
| GPS 2            | 3.55     |
| GPS 3            | 1.89     |
| Simple Average   | 2.35     |
| Weighted Average | 1.79     |

### A.1.6 Results

To test the response of the algorithm to stabilise GPS readings and to correct for situations where, for example, *Searise* has surfaced and a GPS unit initially returns a location with a very poor HDOP value, the behaviour of the algorithm on stationary outdoors and mobile in-the-field data sets was analysed. For the stationary trial, *Searise* recorded raw GPS values for 4 hours. Data from the stationary outdoor test show the limitations of the GPS units used, as all exhibited some degree of drift when satellites changed geometry relative to the vehicle.

An analysis of the variance and standard deviation of each of the GPS data sets, and the calculated simple average and weighted average is given in Table A.1. The simple average has a worse variance than the best of the GPS units (GPS #3). The weighted average has a better variance than all GPS units, and a 41.9% lower variance than that of the simple average. The weighted average has increased calculation of position to an accuracy of 1.79 m.

As demonstrated in Figure A.2, where there are GPS dropouts (only some or none of the GPS units are reporting position) and inaccuracies caused by drift and DOP (in the order of hundreds of metres) the algorithm successfully averages these GPS values with high HDOP and maintains an estimate to within metres of the true position. As the tests were performed on a stationary platform, its accurate coordinates could be determined. No GPS unit provides a consistent or accurate position, but in combination the calculated position is consistent and accurate. The black line includes position estimates when the vehicle is submerged and it can be seen that the endpoints of these dives are consistent with the re-established GPS fixes. These data were from an early mission using  $m = 1.25$ .

### A.1.7 Conclusion

Accurate position estimation is important for this work because the GPS position is used to place observations on the maps made by the vehicle. With more accurate position estimates, the quality of the maps is improved.

The algorithm demonstrates excellent averaging of inaccurate GPS fixes and reduced wander of position estimates compared to simple averaging. The computational cost is low and the technique is extendible to any number of GPS units. The system allows the use of inexpensive off-the-shelf GPS units in a challenging environment. Of particular importance is the ability of this system to use the best available positional data and its robustness when one or more of the GPS units becomes unavailable. Despite the varying quality and reliability of the GPS data in the shown examples, the algorithm has performed well at position estimation.

# B

## Instrumentation Software

The spectrometers communicate through a proprietary driver, packaged as a Java library. Manual communication with the spectrometers is through the *SpectraSuite* software provided by Ocean Optics. Data is gathered during missions by driver software, written by the author in Java. This software can store captured spectra either to text files, and/or to a PostgreSQL database running on board *Searise*.

The driver software dynamically controls the exposure (integration time) of the spectrometers independently. The dynamic control over integration time provides improved signal-to-noise ratio in the data, particularly in dark conditions. Care must be taken to ensure that the sensor does not over-saturate when using longer integration times. The spectrometers have 16-bit resolution (i.e.,  $2^{16}$  values) per sample band. Each captured spectrum is analysed to find the highest signal count in the 2048 channels. The integration time for the next measurement is set such that the peak signal count would have been at the 75% reading if the measurement was redone. The integration time for each reading is stored in the data along with the spectra itself, allowing for normalisation of the spectra to a standard unit of time when required.

The upper spectrometer collects a higher proportion of light. Interpolation is required to analyse a paired upper/lower signal at any particular point in time.

Each spectrometer records 2048 individual samples (bands), at equally spaced wavelengths (channels). Each spectrometer has a slightly different wavelength for each of the 2048 bands, so interpolation must be performed when combining and/or processing data from the two spectrometers.





## MOVETO Flags

## APPENDIX C. MOVETO FLAGS

---

Table C.1: MOVETO flag bit allocation.

|   |   |   |
|---|---|---|
| 0 | Linear Dive                             | The depth is set as a linear interpolation using the position as a guideline. <i>Mutually exclusive to bits 4,5,6.</i>  |
| 1 | Disable Waypoints                       | The steering always points towards the endpoint, not a calculated waypoint. Dynamic steering also invalidates waypoints.  |
| 2 | Dynamic Steering                        | The current actual position is compared to a waypoint position. The time fraction of the total ideal time is then used to give a total error “if things kept going this badly” and that error point is then mirrored around the actual desired line to give a pseudo target position. The sub is aimed for this pseudo target. This overrides (disables) waypoints.   |
| 3 | Alternate end-point radius              | When set, a true 3D distance from endpoint is calculated. This is used for the terminating distance check.  |
| 4 | Dynamic Diving, Distance-based          | The <code>segmentCount</code> is a count for the number of dives to perform for this MOVETO command. 0 is treated as 1. A segment length is calculated (total distance / this value) for each dive. The current ideal position (where the sub should be) is used to calculate a fractional (%) of the dive profile. This is used to set the desired depth). The commanded depth is clipped to a maximum of the value specified in <code>depth</code> . <i>Mutually exclusive to bits 0,5,6.</i> |
| 5 | Dynamic Diving, Time-based              | The <code>segmentCount</code> is a time for each dive (units: seconds). 0 is invalid. The modulus of the current task duration so far and this value is divided by this value to give a fractional (%) of the dive profile. This is used to set the desired depth. The commanded depth is clipped to a maximum of the value specified in <code>depth</code> . <i>Mutually exclusive to bits 0,4,6.</i>  |
| 6 | Dynamic Diving, Constant Segment Length | The <code>segmentCount</code> is a desired length for each dive segment (20 m minimum). The total distance is divided by this to give an integral number of dives possible. This integer is then divided into the total distance, giving an optimal dive segment length, which is used for each of the dives in the whole segment. The commanded depth is clipped to a maximum of the value specified in <code>depth</code> . <i>Mutually exclusive to bits 0,4,5.</i>                          |
| 7 | Depth by Pitch                          | Controls how diving is done. When 0, all flat thrusters push down and the pitch is maintained at 0. When 1, depth by pitch is used, and the rear motors push up.  |



## Optical Calibration

The optical transmission efficiency of the fibre-optic cables, and the spectrometers themselves, govern the proportion of available light at each wavelength that is detected by each spectrometer. Optical losses occur through the fibre-optic cable (roughly 10%) and in the spectrometer (up to 90%). The product of these two at each wavelength gives the proportion of the actual available light that a reading represents.

The following calibration data shows the optical efficiency for each of the two on-board spectrometers, and the associated fibre-optic cable connecting those to the external fibre-optic housing.



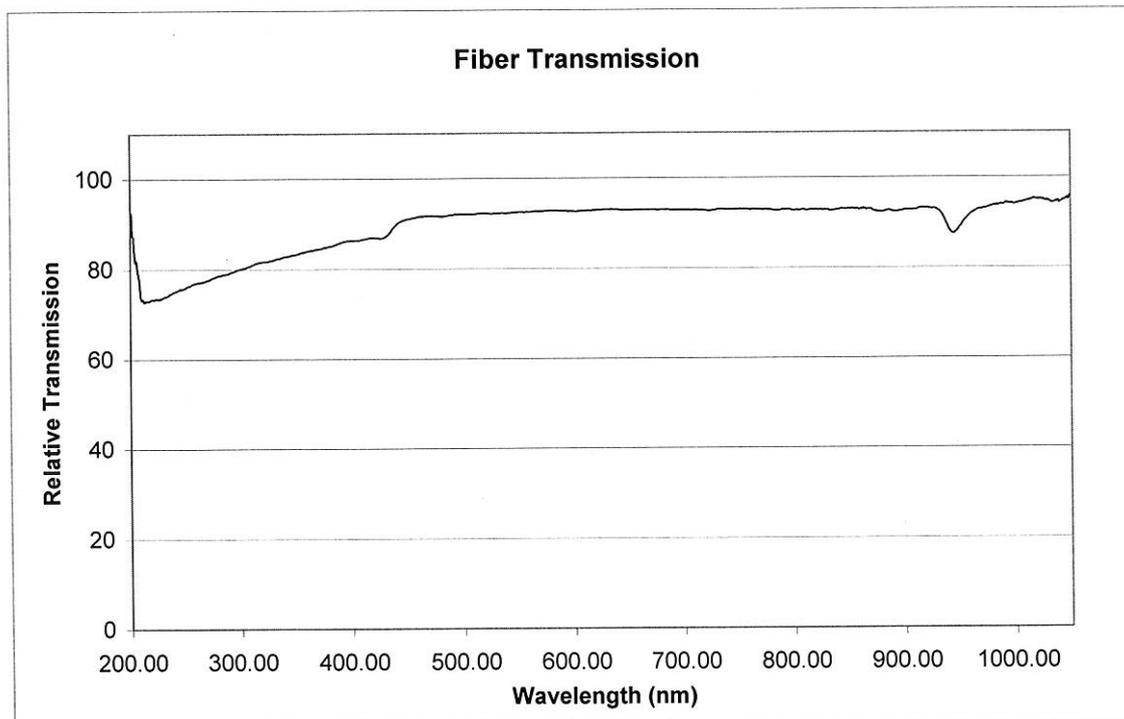
Part #: P400-025-SR  
Date: April 2, 2011  
Assembly #: OOS-001273-03  
Connector 1 #: SMA-905  
Connector 2 #: SMA-905  
Sales Order #: STOCK

[www.OceanOptics.com](http://www.OceanOptics.com)  
Phone:727-733-2447  
Fax:727-733-3962  
[Info@OceanOptics.com](mailto:Info@OceanOptics.com)  
830 Douglas Ave  
Dunedin, FL 34698

**\*Ask about our custom line of Optical Probes and Assemblies.\***

Fiber Type: SR  
Fiber Core Diameter:(um) 400um

Jacketing: Zip Tube  
Length (meters): 0.25



**Inspection Checklist X**

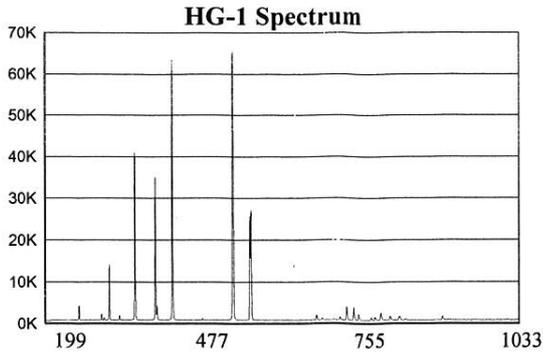
- Polish: X
- Concentricity: X
- Cap Placement: X
- Labeling: X
- Color Coding: X
- Ferrule length: X

Inspected by: Xiaohu Wang  
Xiaohu Wang

PRODUCT OF OCEAN OPTICS SHANGHAI PRC



Wavelength Calibration Data Sheet



Built for:  
 Order Number:  
 Model: **USB2000+XR1**  
 Grating: **GRATING\_#XR1**  
 Bandwidth: 199 - 1033 nm  
 Options: DET2B-200-1100 Detector, NoneLens, SLIT-25 Slit,  
 Serial Number: **USB2+H04855**

| $\lambda$ | Pixel # | Predicted $\lambda$ | $\Delta\lambda$ |
|-----------|---------|---------------------|-----------------|
| 253.652   | 144     | 253.635             | 0.017           |
| 296.728   | 239     | 296.665             | 0.063           |
| 313.155   | 275     | 312.854             | 0.301           |
| 334.148   | 322     | 334.216             | -0.068          |
| 340.365   | 336     | 340.387             | -0.022          |
| 346.620   | 350     | 346.640             | -0.020          |
| 365.015   | 391     | 365.207             | -0.192          |
| 404.656   | 480     | 404.700             | -0.044          |
| 435.835   | 551     | 435.824             | 0.011           |
| 467.816   | 624     | 467.822             | -0.006          |
| 479.992   | 652     | 480.000             | -0.008          |
| 508.582   | 718     | 508.574             | 0.008           |
| 546.074   | 805     | 546.068             | 0.006           |
| 557.029   | 831     | 557.054             | -0.025          |
| 585.249   | 897     | 585.231             | 0.018           |
| 594.483   | 919     | 594.475             | 0.008           |
| 602.999   | 939     | 602.997             | 0.002           |
| 626.649   | 996     | 626.630             | 0.019           |
| 696.543   | 1165    | 696.506             | 0.037           |
| 706.722   | 1190    | 706.729             | -0.007          |
| 717.394   | 1216    | 717.380             | 0.014           |
| 727.294   | 1240    | 727.277             | 0.017           |
| 738.398   | 1268    | 738.273             | 0.125           |
| 750.387   | 1298    | 750.417             | -0.030          |
| 763.510   | 1330    | 763.460             | 0.050           |
| 772.421   | 1353    | 772.364             | 0.057           |
| 794.817   | 1410    | 794.883             | -0.066          |
| 800.616   | 1424    | 800.723             | -0.107          |
| 805.950   | 1438    | 805.952             | -0.002          |
| 823.163   | 1482    | 823.177             | -0.014          |
| 826.452   | 1490    | 826.485             | -0.033          |
| 828.012   | 1494    | 827.993             | 0.019           |
| 829.811   | 1499    | 829.809             | 0.002           |
| 834.682   | 1511    | 834.669             | 0.013           |
| 842.465   | 1531    | 842.311             | 0.154           |
| 852.144   | 1557    | 852.204             | -0.060          |
| 881.941   | 1634    | 881.956             | -0.015          |
| 912.297   | 1715    | 912.326             | -0.029          |

This is a sample of calibration peaks used as there were more than can be shown on this page

**Calibration Coefficients**

First Coefficient: 0.4606344700  
 Second Coefficient: -1.6717e-005  
 Third Coefficient: -3.17392e-009  
 Intercept: 187.68707275  
 Regression Fit: 0.9999997616

**Stray Light Measurements (AU)**

Holmium Oxide (444nm): 2.47  
 Yellow Dye: 2.45  
 Blue Dye: 2.89  
 Molybdate: 2.47  
 OG550 Filter: 2.48  
 RG850 Filter: 3.49  
 FG3 Filter: 1.34

*Nidia Ramirez*

Calibrated By: Nydia.Ortiz  
 Calibrated: 28-July-2011

# Linearity Test

Serial Number **USB2+H04855**

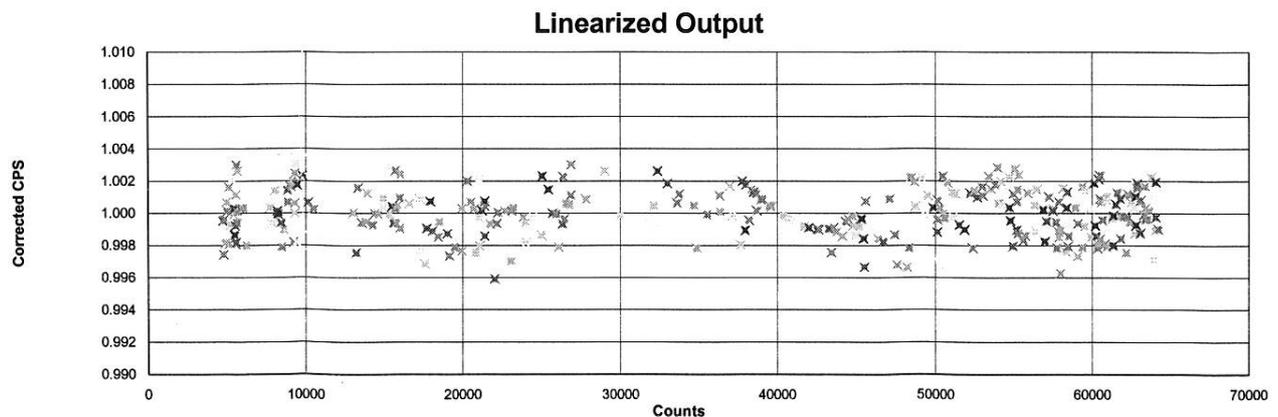
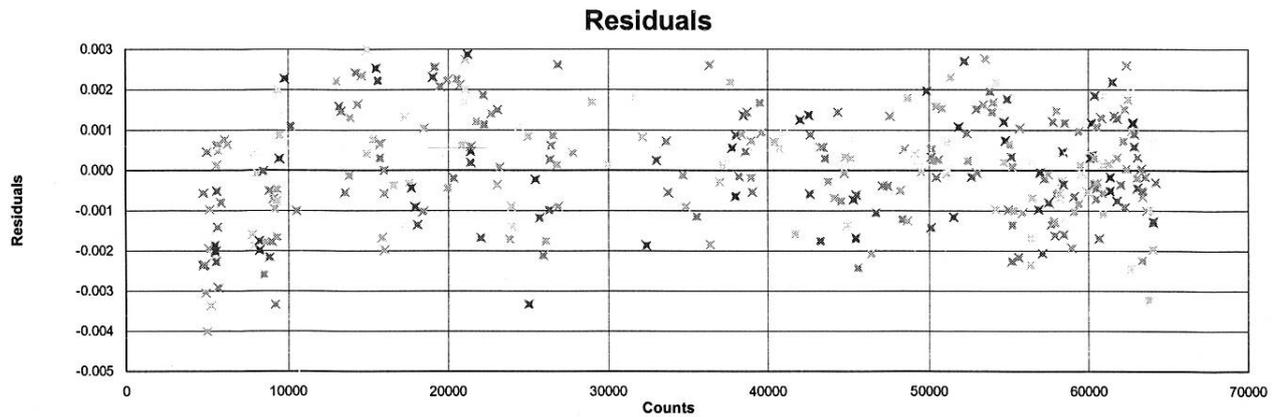
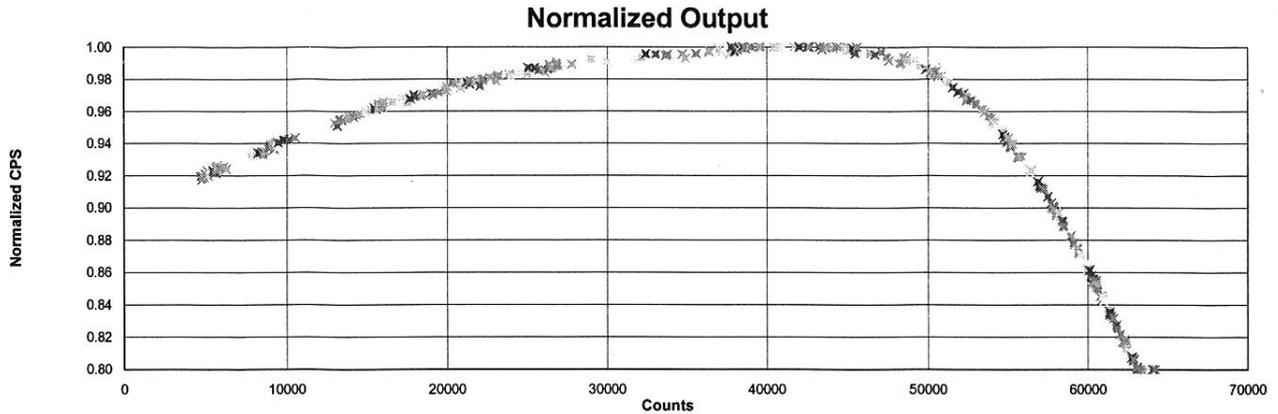
Tech: Nydia.Ortiz

Linearity: **99.85453**

Tested: 7/28/11

Test # 49,493.00

Intercept 0.893875  
Coefficient 1 7.44633e-006  
Coefficient 2 -6.50953e-010  
Coefficient 3 6.28417e-014  
Coefficient 4 -3.24904e-018  
Coefficient 5 8.49905e-023  
Coefficient 6 -1.0758e-027  
Coefficient 7 5.18069e-033

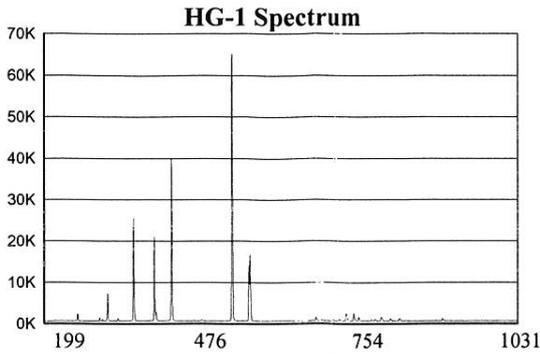


Max 1.00305915

Min 0.99591082



Wavelength Calibration Data Sheet



Built for:  
 Order Number:  
 Model: **USB2000+XR1**  
 Grating: **GRATING\_#XR1**  
 Bandwidth: **199 - 1031 nm**  
 Options: **DET2B-200-1100 Detector, NoneLens, SLIT-25 Slit,**  
 Serial Number: **USB2+H04856**

| $\lambda$ | Pixel # | Predicted $\lambda$ | $\Delta\lambda$ |
|-----------|---------|---------------------|-----------------|
| 253.652   | 144     | 253.632             | 0.020           |
| 296.728   | 239     | 296.698             | 0.030           |
| 302.150   | 252     | 302.149             | 0.001           |
| 313.155   | 276     | 312.920             | 0.235           |
| 334.148   | 323     | 334.134             | 0.014           |
| 334.557   | 324     | 334.560             | -0.003          |
| 365.015   | 393     | 365.137             | -0.122          |
| 404.656   | 482     | 404.719             | -0.063          |
| 435.835   | 553     | 435.925             | -0.090          |
| 546.074   | 809     | 546.111             | -0.037          |
| 557.029   | 834     | 557.050             | -0.021          |
| 585.249   | 901     | 585.262             | -0.013          |
| 594.483   | 923     | 594.442             | 0.041           |
| 602.999   | 943     | 602.979             | 0.020           |
| 696.543   | 1169    | 696.512             | 0.031           |
| 717.394   | 1221    | 717.363             | 0.031           |
| 724.516   | 1238    | 724.496             | 0.020           |
| 738.398   | 1273    | 738.442             | -0.044          |
| 743.890   | 1287    | 743.878             | 0.012           |
| 748.887   | 1299    | 748.868             | 0.019           |
| 750.387   | 1303    | 750.460             | -0.073          |
| 763.510   | 1336    | 763.473             | 0.037           |
| 769.454   | 1351    | 769.438             | 0.016           |
| 772.421   | 1358    | 772.443             | -0.022          |
| 794.318   | 1413    | 794.279             | 0.039           |
| 794.817   | 1415    | 794.814             | 0.003           |
| 800.616   | 1430    | 801.020             | -0.404          |
| 826.452   | 1496    | 826.465             | -0.013          |
| 828.012   | 1500    | 828.021             | -0.009          |
| 834.682   | 1517    | 834.662             | 0.020           |
| 837.760   | 1525    | 837.728             | 0.032           |
| 842.465   | 1537    | 842.465             | 0.000           |
| 849.535   | 1555    | 849.512             | 0.023           |
| 852.144   | 1562    | 852.222             | -0.078          |
| 877.675   | 1629    | 877.686             | -0.011          |
| 881.941   | 1640    | 881.915             | 0.026           |
| 912.297   | 1720    | 912.307             | -0.010          |
| 992.319   | 1937    | 992.282             | 0.036           |

This is a sample of calibration peaks used as there were more than can be shown on this page

**Calibration Coefficients**

First Coefficient: 0.4577443302  
 Second Coefficient: -1.60151e-005  
 Third Coefficient: -3.09239e-009  
 Intercept: 188.05905151  
 Regression Fit: 0.9999997616

**Stray Light Measurements (AU)**

Holmium Oxide (444nm): 1.95  
 Yellow Dye: 2.66  
 Blue Dye: 2.99  
 Molybdate: 2.58  
 OG550 Filter: 2.75  
 RG850 Filter: 3.54  
 FG3 Filter: 1.33

*Naiba Ramirez*

Calibrated By: Nydia.Ortiz  
 Calibrated: 28-July-2011

# Linearity Test

Serial Number **USB2+H04856**

Tech: Nydia.Ortiz

Intercept 0.903242

Linearity: **99.81149**

Coefficient 1 8.14149e-006

Coefficient 2 -7.1793e-010

Coefficient 3 6.26938e-014

Coefficient 4 -3.11155e-018

Coefficient 5 8.05218e-023

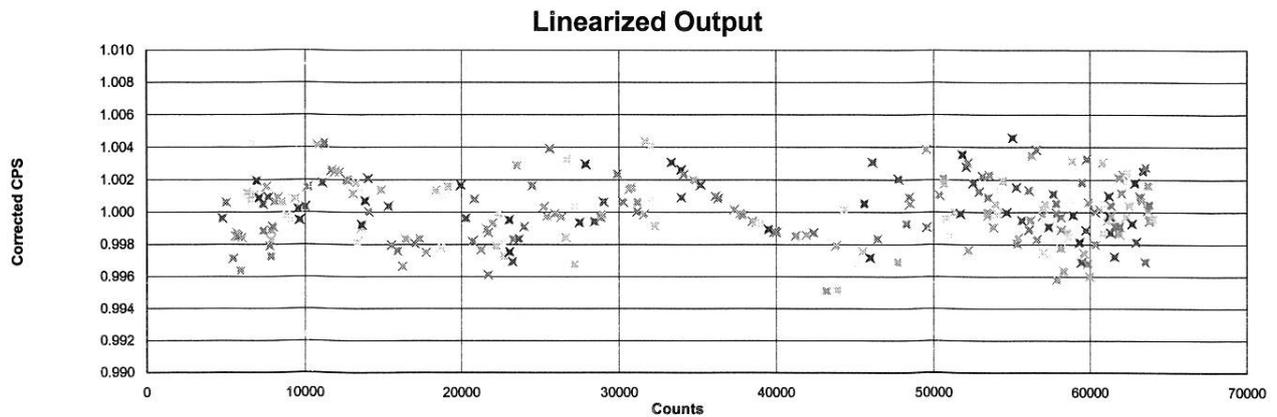
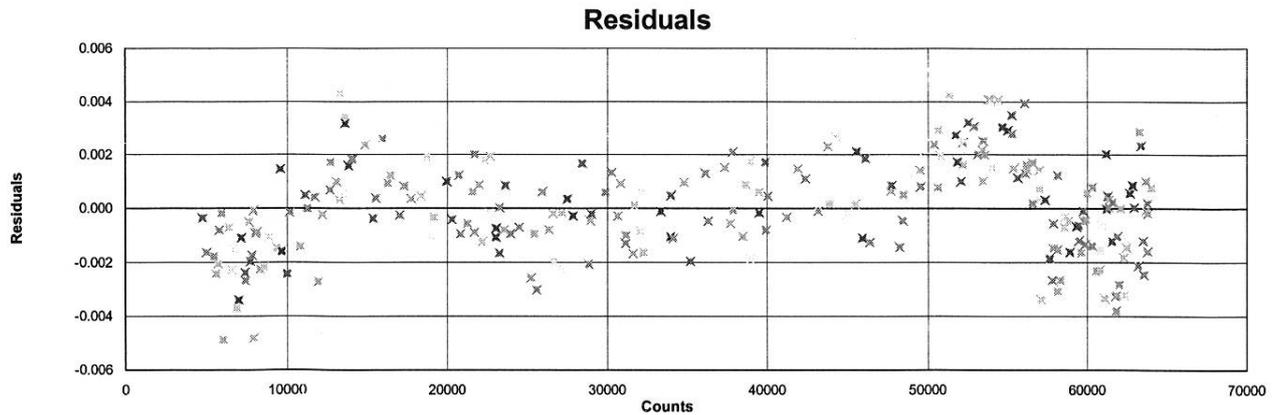
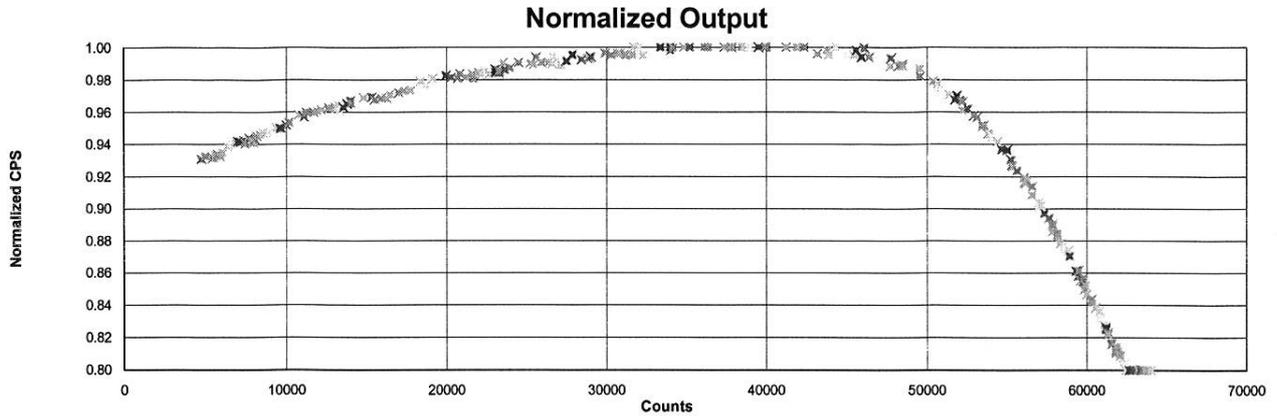
Coefficient 6 -1.02081e-027

Coefficient 7 4.94388e-033



Tested: 7/28/11

Test # 49,495.00



Max 1.00457156

Min 0.99512291

**Table D.1:** USB2000+XR1 Spectrometer Optical Efficiency.  
*(Data provided by Ocean Optics)*

| Wavelength (nm) | S    | P    | Average |
|-----------------|------|------|---------|
| 210             | 58.8 | 55.7 | 57.3    |
| 220             | 61.6 | 61.1 | 61.4    |
| 230             | 63.4 | 62.9 | 63.2    |
| 240             | 64.8 | 63.7 | 64.3    |
| 250             | 65.0 | 62.4 | 63.7    |
| 260             | 63.5 | 60.2 | 61.9    |
| 275             | 60.8 |      | 59.2    |
| 300             | 56.9 | 53.2 | 55.1    |
| 325             | 52.3 |      | 50.5    |
| 340             | 50.2 |      | 47.5    |
| 350             | 48.5 | 41.8 | 45.2    |
| 360             | 46.4 |      | 42.9    |
| 365             | 46.5 |      | 42.5    |
| 370             | 43.3 |      | 40.3    |
| 375             | 41.4 |      | 38.8    |
| 380             | 40.5 |      | 37.8    |
| 400             | 37.1 | 32.6 | 34.9    |
| 420             | 34.5 |      | 32.1    |
| 425             | 33.7 |      | 31.4    |
| 440             | 31.5 |      | 29.4    |
| 450             | 33.5 | 26.2 | 29.9    |
| 460             | 31.1 |      | 28.2    |
| 470             | 28.6 |      | 26.6    |
| 475             | 28.1 |      | 26.2    |
| 500             | 25.4 | 21.8 | 23.6    |
| 550             | 20.2 | 17.2 | 18.7    |
| 560             | 20.0 |      | 18.3    |
| 570             | 21.2 |      | 18.6    |
| 580             | 23.2 |      | 19.5    |
| 590             | 20.6 | 15.1 | 17.9    |
| 600             | 19.0 | 14.9 | 17.0    |
| 650             | 15.1 | 12.3 | 13.7    |
| 700             | 12.8 | 10.0 | 11.4    |
| 750             | 10.8 | 8.4  | 9.6     |
| 760             | 10.5 |      | 9.2     |
| 770             | 10.3 |      | 9.0     |
| 780             | 10.2 |      | 8.8     |
| 790             | 10.8 |      | 9.3     |
| 800             | 11.9 | 7.7  | 9.8     |
| 810             | 13.4 |      | 10.4    |
| 820             | 11.4 |      | 9.2     |
| 830             | 9.8  |      | 8.4     |
| 840             | 8.6  |      | 7.5     |
| 850             | 8.4  | 6.3  | 7.4     |
| 900             | 7.7  | 5.8  | 6.8     |
| 950             | 7.2  | 5.5  | 6.4     |
| 1000            | 7.4  | 5.0  | 6.2     |
| 1050            | 6.8  | 4.4  | 5.6     |
| 1100            | 6.3  | 4.2  | 5.3     |



# References

- Alves, M. A. and Carbonari, A. W. (2008). Fitting PAC spectra with a hybrid algorithm. *Hyperfine Interactions*, 181(1-3):127–130.
- Ambikapathi, A., Chan, T.-H., Ma, W.-K., and Chi, C.-Y. (2010). A robust minimum volume enclosing simplex algorithm for hyperspectral unmixing. In *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, pages 1202–1205.
- Aswani, S. and Lauer, M. (2006). Benthic mapping using local aerial photo interpretation and resident taxa inventories for designing marine protected areas. *Environmental Conservation*, 33(3):263–273.
- Bateson, C. A., Asner, G. P., and Wessman, C. A. (2000). Endmember bundles: a new approach to incorporating endmember variability into spectral mixture analysis. *IEEE Transactions on Geoscience and Remote Sensing*, 38(2):1083–1094.
- Bioucas-Dias, J. M., Plaza, A., Dobigeon, N., Parente, M., Du, Q., Gader, P., and Chanussot, J. (2012). Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches. *Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of*, 5(2):354–379.
- Catanzariti, L. (2013). BUDGET STRATEGY AND OUTLOOK. Budget Paper No. 1, 2013-14. *Commonwealth of Australia*, pages 1–418.
- Cho, H. J. (2007). Depthvariant spectral characteristics of submersed aquatic vegetation detected by Landsat 7 ETM+. *International Journal of Remote Sensing*, 28(7):1455–1467.
- Cho, H. J., Kirui, P., and Natarajan, H. (2008). Test of multi-spectral vegetation index for floating and canopy-forming submerged vegetation. *International Journal of Environmental Research and Public Health*, 5(5):477–483.
- Cho, H. J. and Lu, D. (2010). A water-depth correction algorithm for submerged vegetation spectra. *Remote Sensing Letters*, 1(1):29–35.
- Corke, P., Sikka, P., Roberts, J., and Duff, E. (2004). DDX: A distributed software architecture for robotic systems. In *Proceedings of the 2004 Australasian*

## REFERENCES

---

- Conference on Robotics & Automation*. Australian Robotics & Automation Association.
- de Souza, P. A., Davie, A., Timms, G., Butler, E. C., and Garg, V. K. (2010). Characterisation of suspended particles collected from an estuary in an urban and industrialised centre using magnets onboard an autonomous underwater vehicle. In *IEEE OCEANS, 2010*, pages 1–3. IEEE.
- Dekker, A. G., Byrne, G. T., Brando, V. E., and Anstee, J. M. (2003). Hyperspectral mapping of intertidal rock platform vegetation as a tool for adaptive management. *Technical Report, CSIRO Land and Water, Remote Sensing and Spatial Analysis*.
- Dierssen, H. M., Zimmerman, R. C., Leathers, R. A., Downes, T. V., and Davis, C. O. (2003). Ocean color remote sensing of seagrass and bathymetry in the Bahamas Banks by high-resolution airborne imagery. *Limnology and oceanography*, 1, part 2(48):444–455.
- Dunbabin, M., Roberts, J., Usher, K., and Corke, P. (2004). *A new robot for environmental monitoring on the Great Barrier Reef*. In *Proceedings of the 2004 Australasian Conference on Robotics & Automation*. Australian Robotics & Automation Association.
- Dunbabin, M., Roberts, J., Usher, K., Winstanley, G., and Corke, P. (2005). A hybrid AUV design for shallow water reef navigation. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*., pages 2105–2110. IEEE.
- Dunbabin, M., Usher, K., and Corke, P. (2006). Visual motion estimation for an autonomous underwater reef monitoring robot. In *Field and Service Robotics*, pages 31–42. Springer.
- Farzam, M., Beheshti, S., and Raahemifar, K. (2008). Calculation of abundance factors in hyperspectral imaging using genetic algorithm. *CCECE/CCGEI 2008*, pages 000837–000842.
- Fourqurean, J. W., Duarte, C. M., Kennedy, H., Marbà, N., Holmer, M., Mateo, M. A., Apostolaki, E. T., Kendrick, G. A., Krause-Jensen, D., McGlathery, K. J., and Serrano, O. (2012). Seagrass ecosystems as a globally significant carbon stock. *Nature Geoscience*, 5(7):505–509.
- Gelfand, A. E., Diggle, P., Guttorp, P., and Fuentes, M. (2010). *Handbook of Spatial Statistics*. CRC Press.

## REFERENCES

---

- Guo, F., Ji, Y., and Hu, G. (2001). Methods for Improving the Accuracy and Reliability of Vehicle-borne GPS Intelligence Navigation.
- Hartmann, K., de Souza, P., Timms, G., and Davie, A. (2009). Measuring light attenuation with a compact Optical Emission Spectrometer and CTD mounted on a low cost AUV. In *IEEE OCEANS, 2009*, pages 1–5. IEEE.
- Hasegawa, H. and Yoshimura, T. (2003). Application of dual-frequency GPS receivers for static surveying under tree canopies. *Journal of forest research*, 8(2):0103–0110.
- Hochberg, E. J. and Atkinson, M. J. (2000). Spectral discrimination of coral reef benthic communities. *Coral Reefs*, 19(2):164–171.
- Jackson, R. S., Jackson, R. S., Griffiths, P. R., and Griffiths, P. R. (1991). Comparison of Fourier self-deconvolution and maximum likelihood restoration for curve fitting. *Analytical Chemistry*, 63(22):2557–2563.
- Keshava, N. (2003). A survey of spectral unmixing algorithms. *Lincoln Laboratory Journal*, 14(1):55–78.
- Kirkman, H. (1997). Seagrasses of Australia. Technical report.
- Kirkpatrick, S. (1984). Optimization by simulated annealing: Quantitative studies - Springer. *Journal of statistical physics*.
- Kutser, T., Miller, I., and Jupp, D. L. (2006a). Mapping coral reef benthic substrates using hyperspectral space-borne images and spectral libraries. *Estuarine, Coastal and Shelf Science*, 70(3):449–460.
- Kutser, T., Vahtmae, E., and Metsamaa, L. (2006b). Spectral library of macroalgae and benthic substrates in Estonian coastal waters. *Proc. Estonian Acad. Sci. Biol. Ecol*, 55(4):329–340.
- Lesser, M. P. and Mobley, C. D. (2007). Bathymetry, water optical properties, and benthic classification of coral reefs using hyperspectral remote sensing imagery. *Coral Reefs*, 26(4):819–829.
- Lundblad, E. R., Wright, D. J., Miller, J., Larkin, E. M., Rinehart, R., Naar, D. F., Donahue, B. T., Anderson, S. M., and Battista, T. (2006). A Benthic Terrain Classification Scheme for American Samoa. *Marine Geodesy*, 29(2):89–111.
- Meng, L.-J. and Ramsden, D. (2000). An inter-comparison of three spectral-deconvolution algorithms for gamma-ray spectroscopy. *Nuclear Science, IEEE Transactions on*, 47(4):1329–1336.

## REFERENCES

---

- Mobley, C. D. (1994). *Light and water*. radiative transfer in natural waters. Academic Pr.
- Mount, R. E., Lucieer, V. L., Lawler, M. M., and Jordan, A. R. (2005). Mapping of estuarine and marine habitats in the southern NRM region.
- Orth, R. J., Carruthers, T. J., Dennison, W. C., Duarte, C. M., Fourqurean, J. W., Heck, Jr, K. L., Hughes, A. R., Kendrick, G. A., Kenworthy, W. J., and Olyarnik, S. (2006). A global crisis for seagrass ecosystems. *Bioscience*, 56(12):987–996.
- Pala, C. (2013). Giant Marine Reserves Pose Vast Challenges. *Science*, 339(6120):640–641.
- Pegau, W. S., Gray, D., and Zaneveld, J. R. (1997). Absorption and attenuation of visible and near-infrared light in water: dependence on temperature and salinity. *Applied optics*, 36(24):6035–6046.
- Rigby, P., Pizarro, O., and Williams, S. B. (2010). Toward adaptive benthic habitat mapping using gaussian process classification. *Journal of Field Robotics*, 27(6):741–758.
- Whitley, D. (1994). A genetic algorithm tutorial. *Statistics and computing*, 4(2):65–85.

# Acronyms

**XR-1** Ocean Optics USB2000+XR-1 miniature optical spectrometer. 31, 36, 41, 42, 44, 45, 47

***Searise*** the autonomous underwater vehicle. 6, 11, 13–19, 23, 25–28, 31, 36, 41, 45, 47–49, 83, 101, 105, 107, 114, 117, 129, 132, 134, 136, 137, 139, 144

**AUV** autonomous underwater vehicle. 2–5, 8, 24, 27, 34, 40, 44, 82, 83, 88, 95, 98, 100, 104, 134, 135, 137, 139, 142

**CBL** Cloudy Bay Lagoon, Bruny Island, Tasmania. 105

**CSIRO** Commonwealth Scientific and Industrial Research Organisation. 19

**CSIRO's ASLab** CSIRO's Autonomous Systems Laboratory, based at the Queensland Centre for Advanced Technologies. 11, 19, 23

**DDX** dynamic data exchange. 19

**FOC** fibre-optic cable. 33, 34, 36, 40–42, 44, 87

**FOV** field of view. 34, 35, 87

**GA** genetic algorithm. 6, 49, 53, 58, 61, 62, 86, 134

**GPS** global positioning system. 6, 8, 15, 17, 18, 20, 23, 24, 81, 83, 97, 100, 104, 117, 118, 128, 139–142, 144, 145

**IMU** inertial measurement unit. 27

**ISSL** CSIRO's Intelligent Sensing and Systems Laboratory. 11, 15, 19

**LA** linear algebra. 49, 58

**LMM** linear mixing model. 6, 33, 54, 58

**NNLS** Non-Negative Least Squares. 49

## Acronyms

---

**RGB** red-green-blue colour triplet. 95

**SA** simulated annealing. 6, 49, 53, 58, 61, 63, 134

**SAV** submerged aquatic vegetation. 4