

IMAGING-BASED LENSLESS POLARISATION-SENSITIVE FLUID STREAM ANALYSER FOR AUTOMATED, LABEL-FREE, AND COST-EFFECTIVE MICROPLASTIC CLASSIFICATION



Author:

Fraser Derrick Charles Montandon

Supervisor:

Associate Professor Fred Nicolls

January 2024

Submitted to the Department of Electrical Engineering at the University of Cape Town in fulfilment of the academic requirements for the degree of Master of Science in Electrical Engineering.

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

Abstract

The presence of microplastics in the environment is of concern with the actual distribution of this pollution remaining relatively unknown. The ocean is of particular interest as the monitoring of microplastics in this area presents a challenge in that in situ fluid stream solutions are not readily available and traditional sampling methods are labour-intensive and costly. Additionally, the lack of consensus on sampling techniques makes comparing results difficult. Our proposed device demonstrates an imaging-based lensless polarisation-sensitive fluid stream analyser (FSA) for automated, label-free, and cost-effective microplastic classification. The FSA performs analysis at high flow rates with a custom-designed illumination circuit that reduces motion blur and provides quantitative sample information using a polarisation-sensitive image sensor. Digital in-line holography (DIH) and birefringence numerical computation are utilised in the processing workflow. The device can be used for either quantitative polarisation-sensitive imaging and analysis or for further machine-learning-based activities, including the classification of samples. Both abilities are demonstrated in this study. Our analyser computes the two-dimensional birefringent characteristics of samples and we investigate the detection of synthetic polymer birefringent textures due to the optical anisotropy of these materials. We perform a comparative machine learning study with both learned and filter bank feature generation being assessed to aid the microplastic classification process. The FSA and classifier components are used to develop an end-to-end workflow that samples a fluid stream and determines the composition of marine and microplastic particles. We use two phytoplankton cultures to create a simplified marine environment for testing purposes. To demonstrate the performance of our classification methods we tested our device and workflow in a two-class configuration for marine microorganisms and plastics, as well as a five-class configuration for marine microorganisms and four individual plastic types (polyethylene (PE), polyethylene terephthalate (PET), polypropylene (PP), and polystyrene (PS)). Our analysis shows that high accuracy is achieved from the classifier implementation, with the simulated marine environment experiments further supporting the ability of the proposed implementation.

Acknowledgments

I would like to acknowledge the South African International Maritime Institute and the University of Cape Town for the funding contributions towards this thesis.

To Prof Fred Nicolls, I am grateful for your supervision throughout this process. You always made time to discuss the project, engineering, and life. The mentorship, insights, and advice you provided have been invaluable to my growth. Your gentle guidance and patience allowed me the space to forge my own path and brought me back on track whenever I ventured off-piste.

To Dr Clive Garcin at Disa Scientific, your support in the chemistry and biological domains is appreciated.

Thank you to the Marine Robotics Unit and the Digital Image Processing Laboratory; learning from you all was enriching.

Finally, to Ang and my family, which now includes little Raphaella, thank you for your love and support and for encouraging me to get across the finish line.

Plagiarism Declaration

1. I know that plagiarism is wrong. Plagiarism is to use another's work and pretend that it is one's own.
2. I have used the IEEE convention for citation and referencing. Each contribution to, and quotation in, this report from the work(s) of other people has been attributed, and has been cited and referenced.
3. This report is my own work.
4. I have not allowed, and will not allow, anyone to copy my work with the intention of passing it off as their own work or part thereof.
5. I know the meaning of plagiarism and declare that all the work in the document, save for that which is properly acknowledged, is my own. This thesis/dissertation has been submitted to the Turnitin module (or equivalent similarity and originality checking software) and I confirm that my supervisor has seen my report and any concerns revealed by such have been resolved with my supervisor.

FDCM

.....

Fraser Derrick Charles Montandon

January 22, 2024

Acronyms

BCE Binary cross-entropy. 55–57

CNN Convolutional neural network. 13, 22, 23

DIH Digital in-line holography. i, vii, ix, 3, 5, 12, 15–21, 24, 25, 30, 31, 39, 40, 44, 47, 49, 51, 57, 58, 120, 122

FOV Field of view. 9, 20, 24, 30, 31, 42, 62

FPR False positive rate. 6, 22, 73, 90, 91, 94, 98, 102, 111, 112, 115

FSA Fluid stream analyser. i, vii, ix, xi, 3–6, 18–21, 23, 25, 26, 31, 34, 38–44, 46–53, 55–58, 60, 62, 68, 69, 72–79, 81–83, 85, 95–100, 108, 109, 117, 119–121

IOU Intersection over union. 5, 56, 57

KNN K-nearest neighbour. xi, xii, 86, 90–93

LED Light emitting diode. 12, 13, 15, 44, 46

LPE Light pulse engine. ix, x, 40, 44–46, 66, 67, 71, 72, 77, 120

PCA Principal component analysis. 12, 85

PDF Probability density function. x, xi, 69–71, 113

PE Polyethylene. x, 3, 4, 7, 41, 49, 53, 68, 75, 76, 81, 82, 107–110, 116

PET Polyethylene terephthalate. x, 3, 4, 7, 53, 62, 75, 81, 82

PFA Polarisation filter array. ix, 2, 14, 16, 20, 24, 34, 41, 42, 51, 120, 122

PMMA Polymethyl methacrylate. 41, 76, 107–110, 116

PP Polypropylene. x, 3, 4, 7, 53, 75, 81, 82, 92

PS Polystyrene. [x](#), [3](#), [4](#), [7](#), [53](#), [75](#), [81](#), [82](#), [92](#)

RGB Red, green, and blue. [12](#), [14](#), [15](#), [20](#), [22](#), [23](#), [42](#), [46](#), [61](#)

SNR Signal-to-noise ratio. [2](#), [20](#), [23](#)

SVM Support vector machine. [xi](#), [11](#), [14](#), [22](#), [86](#), [88](#), [91–93](#)

UMAP Unified manifold and projection. [85](#), [101](#), [109](#)

Contents

Abstract	i
Acknowledgments	ii
Plagiarism Declaration	iii
Acronyms	iv
Table of Contents	vi
List of Figures	ix
List of Tables	xii
Chapter 1: Introduction	1
1.1 Introduction	1
1.2 Motivation	2
1.3 Problem and thesis statement	3
1.4 Objectives	3
1.5 Scope and limitations	3
1.6 Outline	4
Chapter 2: Literature review	8
2.1 General literature	8
2.1.1 Digital holography	9
2.1.2 Polarisation microscopy	9
2.1.3 Polymer birefringence and texture	10
2.2 Consulted scientific publications	11
2.2.1 Static microplastic classification	11
2.2.2 Imaging flow cytometry	12
2.2.3 Polarimetry	14
2.3 Significance of proposed work	15
2.4 Summary	17

Chapter 3:	Approach and theoretical framework	18
3.1	Overview of approach	18
3.1.1	Proposal for imaging-based FSA	20
3.1.2	Proposal for computer vision	21
3.1.3	Proposal for pipeline	23
3.2	Hardware information	23
3.2.1	Image sensor information	23
3.2.2	Flow cell information	24
3.3	Theoretical framework	25
3.3.1	Encoder-decoder networks for segmentation and classification	26
3.3.2	Filter banks	27
3.3.3	Lensless DIH	30
3.3.4	Polarisation and birefringence	33
3.3.5	Motion blur	36
3.4	Software development and workstation specifications	37
3.5	Conclusion	38
Chapter 4:	Imaging fluid stream analysis	39
4.1	Preliminaries	39
4.1.1	High-level overview of propagation	41
4.2	Materials and methods	42
4.2.1	FSA	42
4.2.2	Software workflow development	47
4.2.3	DIH	49
4.2.4	Polarisation sensitive birefringence imaging	50
4.2.5	Segmentation	51
4.3	Results	56
4.3.1	Trainable segmentation results	57
4.3.2	Experimental demonstration	57
4.3.3	Calibration	68
4.4	Discussion	71
4.5	Conclusions	72
Chapter 5:	Computer vision	73
5.1	Introduction	73
5.2	Materials and methods	75
5.2.1	Dataset	75
5.2.2	Learned feature classifier	78

5.2.3	Filter bank classifier	81
5.3	Results	90
5.4	Discussion and analysis	94
5.5	Conclusions	95
Chapter 6:	Pipeline	96
6.1	Preliminaries	96
6.2	Workflow	98
6.2.1	Algorithm implementation	98
6.2.2	Recreated marine environment and protocol	101
6.3	Materials and methods	104
6.3.1	Experiment 6A: Materials and methods	105
6.3.2	Experiment 6B: Materials and methods	106
6.3.3	Experiment 6C: Materials and methods	107
6.4	Results	110
6.4.1	Experiment 6A: Combined microplastic analysis in a marine environment	111
6.4.2	Experiment 6B: Multi-class microplastic analysis in a marine environment	114
6.4.3	Experiment 6C: PE PMMA sphere analysis	116
6.5	Discussion and analysis	118
6.5.1	General workflow	118
6.5.2	Experiments 6A and 6B: Mixed media	118
6.5.3	Experiment 6C: PE and PMMA spheres	118
6.6	Conclusions	119
Chapter 7:	Conclusion	120
7.1	Summary of main topics	120
7.2	General conclusions	121
7.3	Contributions	122
Bibliography		124
Appendix A: Supporting data		130
A.1	Confusion matrices for filter bank feature extraction	130
A.2	Dry screening procedure	130
A.3	Costing of the FSA	132
Appendix B:		133

List of Figures

3.1	Illustration of an imaging-based FSA.	19
3.2	Project development.	19
3.3	Sony polarisation image sensor illustration [30].	24
3.4	Illustrations of the prototype flow cell. (a) Prototype flow cell. 1. Fluid entry point. 2. Cavity housing for an illumination source. 3. Fluid exit point. 4. Cavity housing for an image sensor. (b) Prototype flow cell cutaway showing round glass observation window.	25
3.5	Standard U-Net structure [20].	27
3.6	Basic on-chip DIH setup.	31
4.1	Polarisation propagation. Light propagates from the excitation source as randomly polarised light (RPL). It then passes through the circular polariser to produce circularly polarised light (CPL). Next, it illuminates an optically anisotropic sample (in this example) to become elliptically polarised light (EPL), and finally this light is captured by the PFA. The PFA shows the four different linear polarisation orientations per a super-pixel.	41
4.2	The realised FSA from two different views indicating important components and areas.	43
4.3	LPE circuit illustration.	45
4.4	Oscilloscope trace from LPE analysis.	46
4.5	Custom designed components.	47
4.6	Processing workflow of the FSA. The input frame is processed and expanded to seven output channels. The output channels can either be used for qualitative observation or further downstream quantitative processes.	48
4.7	Image correction routine. Scale bar in all figures indicates 50 μm	49
4.8	Resolution test target — full frame (propagated).	58
4.9	Resolution test target analysis — cropped orange region from Figure 4.8.	59
4.10	Intensity profile resolution test target analysis. (a) Successfully resolved elements. (b) Point of failure.	59

4.11	Preprocessed propagated mid-depth intensity of microplastic samples (low-intensity regions).	60
4.12	Depolariser analysis at 650 nm.	61
4.13	Pseudo-coloured retardance and slow axis orientation (cropped).	61
4.14	Birefringence processing of PET. (a) Preprocessed hologram. (b) Propagated intensity. (c) Pseudo-coloured birefringence retardance and slow axis orientation. (d) Pseudo-coloured birefringence retardance and slow axis overlay/fused with propagated intensity. Scale bar in all figures indicates 50 μm . Colour scale legend available in Figure 4.13.	62
4.15	Comparison of the raw hologram at beginning of pipeline and the resulting segmentation output (cropped).	63
4.16	Image segmentation when applying workflow for network-based segmentation (cropped).	63
4.17	Propagated intensity and predicted output patches (128 \times 128) pixels. Scale bar in all figures indicates 50 μm	64
4.18	Motion blur test (40 mL/minute).	66
4.19	Propagated intensity images of PE spheres using continuous excitation. (a) 20 mL/minute. (b) 40 mL/minute. (c) 60 mL/minute. Scale bar in all figures indicates 50 μm	67
4.20	Propagated intensity images of PE spheres using the LPE. (a) 20 mL/minute. (b) 40 mL/minute. (c) 60 mL/minute. Scale bar in all figures indicates 50 μm	67
4.21	Normalised histogram (PDF approximation) PE. (a) PE initial run. (b) PE duplicate run.	70
4.22	Normalised histogram (PDF approximation) PMMA. (a) PMMA initial run. (b) PMMA duplicate run.	71
5.1	Layout of adapted ResUnet classifier.	79
5.2	Intensity mid-depth propagation of synthetic polymer samples, where objects are the low-intensity regions in the frame. (a) PS. (b) PE. (c) PET. (d) PP. The small particle fines mentioned in the sample refining process are visible in (b). Scale bar in all figures indicates 50 μm	81
5.3	Retardance channel of the same four samples from Figure 5.2. (a) PS. (b) PE. (c) PET. (d) PP.	82
5.4	Dimensionality reduction four-class (PE, PP, PS, and PET), showing three of five dimensions.	87
5.5	Dimensionality reduction for two classes (marine and polymer).	88

5.6	Classifier confusion matrices showing the TPR on the diagonal. (a) Random forest. (b) SVM. (c) Discriminant. (d) Naive Bayes. (e) Tree. (f) KNN.	93
6.1	Pipeline overview. The input frame acquired by the FSA undergoes the preprocessing and base feature extraction routine. The expanded seven-channel stack is passed to the neural network for segmentation and classification. The particles identified as polymer type are then sent for further feature extraction with the filter bank. This high-dimensional feature vector undergoes dimensionality reduction. Finally, the random forest classifier will perform a pixel-wise prediction on whether the polymer flagged sample is PE, PET, PS, or PP.	99
6.2	Raw cultures. Left: <i>Chaetoceros calcitrans</i> . Right: <i>Isochrysis galbana</i> . .	103
6.3	PMMA and PE sphere propagated intensity and birefringence retardance channels. (a) PMMA propagated intensity. (b) PMMA birefringence retardance. (c) PE propagated intensity. (d) PE birefringence retardance. Scale bar in Figures 6.3(a) and 6.3(b) indicates 50 μm	108
6.4	PE and PMMA mapping.	110
6.5	Synthetic polymer concentration adjustment. The error bars indicate the standard deviation and the sample points indicate the mean value of the repeated runs of the experiments.	112
6.6	Synthetic polymer concentration adjustment — area analysis.	112
6.7	Normalised histogram (PDF approximation). Synthetic polymer and marine.	113
6.8	Mixed media experiments.	115
6.9	Stepped concentration analysis for two synthetic polymers. (a) PMMA concentration is changed whilst PE is held constant. (b) PE concentration is changed whilst PMMA is held constant.	117
B.1	LPE circuit.	133

List of Tables

4.1	Binary segmentation results.	57
4.2	PE sphere measurement run.	69
4.3	PMMA sphere measurement run.	70
5.1	Adapted ResUnet structure.	80
5.2	Feature list	84
5.3	Learned feature two-class classifier results (mean).	90
5.4	ResUnet classifier results — individual classes.	90
5.5	U-Net classifier results — individual classes.	90
5.6	Feature extraction — two-class classifier results (mean %).	91
5.7	Confusion matrices two-class with TPR on diagonal: random forest (left), KNN (right).	91
5.8	Feature extraction four-class classifier results (mean).	93
5.9	Random forest classifier results — individual classes.	94
6.1	Cell count of phytoplankton cultures.	102
6.2	PE and PMMA microsphere comparison.	107
6.3	PE and PMMA confusion matrix — model test set.	116
6.4	PE PMMA confusion matrix live data capture.	116
6.5	PE PMMA live data capture counts.	117
A.1	Confusion matrices two-class with TPR on diagonal: random forest (left), KNN (right).	130
A.2	Confusion matrices two-class with TPR on diagonal: tree (left), SVM (right).	130
A.3	Confusion matrices two-class with TPR on diagonal: naive Bayes (left), discriminant (right).	130
A.4	Costing of the FSA.	132

Chapter 1

Introduction

Microplastic detection is an active research area with many competing analysis techniques [1–5]. However, low-cost fluid stream analysis of microplastics presents an unmet need that this work attempts to address. This chapter provides a summary of the document, offering details on the structure and implementation of the project.

The chapter starts by introducing the problem and providing background context. The motivation that led to the project investigation is then discussed, followed by the problem and thesis statement. The objectives of the project are detailed and this is followed by a summary of the contributions. Next, the scope and limitations are discussed, and finally the chapter concludes with the presentation of the thesis outline.

1.1 Introduction

Defined as synthetic polymers between 1 μm and 5 mm in diameter [6], the global distribution of microplastics in ocean waters remains relatively unknown due to lack of detailed measurement data [7]. Traditional sampling techniques are labour intensive, costly, and require analytical equipment in laboratory settings [1–3]. Laboratory techniques are sensitive and accurate but involve retrieval and thorough preparation of raw samples before analysis can occur [5]. Therefore these existing techniques are low-throughput and ex situ solutions. Due to resolution limits, current detection techniques do not adequately cover the full range of microplastic sizes [4]. Although Fourier-transform infrared spectroscopy and other spectral techniques employed in a laboratory setting are well suited to microplastic detection, they are not readily adaptable to in situ analysis; this is due to the strong water absorption bands that infrared sensing has to contend with. Furthermore, the equipment is expensive, slow, and not readily suited to multi-object heterogeneous solutions. Additionally, portable hyperspectral cameras are

useful but prohibitively expensive, limiting their practicality as field-detection devices.

In situ analysis of microplastics remains a nascent area. Recent digital holographic studies in static microplastic analysis move toward an imaging-based solution [8–10], though these used less portable lensed off-axis holographic configurations [11] and were limited to machine learning classifier implementations. Synthetic polymers are optically anisotropic (birefringent) and polarisation microscopy provides an alternative analysis tool to gain insight into the molecular arrangement of these materials [12]. Applying lensless in-line holographic polarisation-sensitive sensing to in-stream analysis of microplastics may provide useful features for classification, as has been done in a lensed off-axis study [13]. Formulating birefringence information into quantitative measurements for use with temporal polarisers and traditional image sensors has been effective [14]. A similar technique used in a static biomedical study employed a polarisation filter array (PFA) sensor instead of multiple temporally separated images to quantitatively resolve birefringence information in a single shot [15]. Imaging flow cytometry combined with digital in-line holography (DIH) techniques is a lightweight and cost-effective approach for marine fluid stream analysis, and this is further supported by recent advancements in computer processing performance and smaller image sensor pixel sizes [16, 17].

This project aims to work towards automating sampling and enhancing in-stream microplastic detection methods using lensless holographic polarisation-sensitive imaging. However, in-stream detection presents a challenge in the form of a low signal-to-noise ratio (SNR) due to turbidity and the presence of other particulates.

Although the main focus of this study is microplastic detection, the workflow presented in this research can hopefully serve as a framework for future studies in imaging-based fluid stream analysis, with the presented methodology and techniques being adapted and extended to explore other applications.

1.2 Motivation

Environmental monitoring of microplastics, particularly in marine habitats, has emerged as a significant concern [18]. The detection of microplastics in such settings calls for the development of portable and low-cost detection systems. The University of Cape Town is researching the development of sensing tools for microplastic detection, and this project contributes towards that objective. Technological advances in optical sensors and more affordable circuitry have created an opportunity to investigate the development of low-cost optical detectors for microplastic classification, and these factors have motivated this project's direction.

1.3 Problem and thesis statement

Detecting microplastics in marine environments is resource-intensive and time-consuming. The sampling and analysis of microplastics in ocean water have been conducted sparsely, resulting in a limited understanding of their distribution. There has been insufficient exploration into developing affordable, deployable, or portable techniques and devices for this purpose. Consequently, only limited progress has been made in this field, with devices having low throughput or lacking portability. The absence of portable and rapid testing detectors has presented a significant obstacle to researchers who require in situ testing capabilities and thus has limited the mapping of the distribution of plastic pollution in ocean water.

This thesis investigates optical methods for low-cost, label-free, automated fluid stream detection of microplastic particles by developing and evaluating a fluid stream analyser (FSA) that integrates DIH, and spatial polarisation sensing in combination with a classification back-end, resulting in an end-to-end pipeline or workflow.

1.4 Objectives

This study aims to develop a low-cost FSA capable of classifying select synthetic polymers (polyethylene (PE), polyethylene terephthalate (PET), polypropylene (PP), and polystyrene (PS)) of size range 50 μm to 1000 μm present in mixed media solutions. The mixed media solutions are created to simulate ocean water. We define three objectives to support the development of our detector.

- Fluid stream analysis: Develop a low-cost high flow rate lensless polarisation-sensitive imaging-based FSA.
- Microplastic classifier: Develop a microplastic classifier that uses the FSA.
- Pipeline: Develop an automated software workflow that combines the FSA and classifier objectives to provide end-to-end microplastic detection.

1.5 Scope and limitations

The analyser developed in this study serves as a complementary tool to existing diagnostic techniques. The instrument is intended to enhance sensing capabilities, particularly in resource-limited settings, by providing the capability to detect microplastics in situ. It can also function as an initial detection device, providing a preliminary assessment that may guide the need for more comprehensive detection using traditional techniques

where required.

Importantly, this study focuses exclusively on the detection of specific plastic types, namely PE, PET, PP, and PS, and does not consider other types of plastics. The size range of interest is 50 μm to 1000 μm and samples were used in their virgin state.

The recreated marine environment was designed to be easily reproducible for future research. While it does not capture the more complex biological interactions present in actual marine samples, it allows for generating a reliable ground truth that would be challenging to determine in a real-world deployment. However, the study does not address the potential effects of biofouling and other environmental factors, which could impact the analyser's performance in a natural marine environment.

The protocols, workflow, and detector developed in this work are intended to serve as a proof of concept that supports the advancement towards in situ analysis of microplastic particles. The feasibility of utilising polarisation-sensitive optical sensing is a theme in this work and the study provides a foundational framework for its use. However, it also acknowledges the need for further refinement and validation in real-world scenarios. This project aims to present a conceptual framework that can be taken forward to be optimised and integrated into embedded systems for future studies that may require field testing and deployment.

1.6 Thesis outline

The thesis is organised as follows:

Chapter 2, Literature review: The literature review presents key works consulted in this study. First, we briefly discuss the tertiary literature that forms part of the general consulted works. This literature includes background on Fourier optics holography, polarisation microscopy, and synthetic polymers. Next, we cover selected scientific publications and primary literature on static imaging-based microplastic classification and detection methods. This is followed by literature on imaging flow cytometry and polarimetry applications. The chapter concludes with a summary of the current state-of-the-art and a discussion of the significance of the proposed work.

Chapter 3, Approach and theoretical framework: This chapter outlines the approach underlying the development of an imaging-based microplastic detector. We propose the three subsystems or contributions for the project, namely: FSA, computer vision classifier, and pipeline. We then detail the essential components used, such as the polarisation image sensor and the flow cell.

The chapter also addresses the theoretical framework, including encoder-decoder networks (for segmentation and some classification models), filter banks for feature extraction, lensless DIH techniques, polarisation, birefringence, and motion blur.

Chapter 4, Imaging fluid stream analysis: The FSA is introduced and developed in this chapter. The analyser is configured in a lensless layout, incorporating DIH and spatial polarisation sensing. The hardware elements are first detailed, including the flow cell, polarisation image sensor, fibre-coupled laser diode, 3D printed supporting components, and a custom-designed light pulse driver circuit.

An image preprocessing routine is established to reduce aberrations and artefacts in the acquired frame. Following this, we detail software elements enabling the input image frame from the FSA to expand into a seven-channel output. The output channels are developed by incorporating holography propagation and birefringence processing routines for lensless optical anisotropic sensing. We apply Jones calculus [14, 19] and the angular spectrum method [11] to match our optical setup with the computational backend. For pixel-wise particle segmentation, we train a customised supervised encoder-decoder network.

In reviewing the results from the segmentation implementation, it is noted that our proposed segmenter achieved a mean intersection over union (IOU) of 94.46%, which is an improvement over a standard U-Net [20] configuration trained alongside it.

Next, we present the experimental demonstrations. Testing of the DIH propagation was performed with a static resolution test target and an intensity profile analysis, yielding a lateral resolution of $6.6 \pm 0.8 \mu\text{m}$, and confirming the expected resolution of our system. Subsequently, we passed microplastic samples through the FSA in a fluid stream.

For the polarisation birefringence demonstration, we first tested the subsystem with an anisotropic liquid crystal depolariser, with results consistent with the manufacturer's specifications. Then, a sample of microplastic specimens was passed through the analyser, and we mapped the retardance and slow axis to a hue and saturation colour space, which is useful for visualisation purposes.

A segmentation demonstration was performed on various samples of both marine phytoplankton and microplastics, including challenging examples closer to the resolving limits of the system.

A motion blur study was then conducted to assess performance improvement when using the excitation driver for pulsed light driving. Finally, a calibration procedure was performed with two types of scientific calibration spheres (including one unseen

polymer type and size). The resulting calibration output agreed with the manufacturer's specifications of the spheres.

Chapter 5, Computer vision: This chapter develops and presents the microplastic classifiers. These classifiers used the outputs from the FSA developed in Chapter 4 for further feature extraction. Two approaches were explored: a neural network learned feature classifier and a filter bank feature classifier. This investigation aimed to select the most appropriate or suitable model for two different implementations, namely a binary marine and microplastic classifier and a multi-class polymer classifier.

A small dataset of marine phytoplankton and microplastic samples was created for classifier training. The learned feature classifier, employing a custom encoder-decoder structure, was trained for the binary marine and microplastic problem. Following this, the filter bank classifier was developed. A sequential routine of filter banks, feature selection, and supervised dimensionality reduction led to the feature set used to train multiple classifier implementations for analysis.

The results from the learned feature approach suggest that our proposed implementation was the preferred candidate for the two-class (marine and microplastic) task with a mean accuracy of 98.66%. Additionally, it achieved a false positive rate (FPR) of 1.25% when treating microplastics as the positive class, representing an improvement over a standard U-Net implementation used for comparative purposes. Regarding the multi-class polymer classification problem, the routine using filter bank feature extraction with a random forest classifier attained the highest mean accuracy at 95.03% for the four microplastic types.

Chapter 6, Pipeline: This chapter develops and presents the pipeline designed to amalgamate the FSA and classifier developed in the previous two chapters. Our method also establishes a testing and evaluation protocol. The pipeline is intended to provide a platform for testing and assessment beyond standard machine learning classifier results. The end-to-end workflow obtains an image frame from the FSA as an input and performs preprocessing, segmentation, classification, and logging of each particle that passes through the analyser within the specified size range. The protocol followed to establish a simulated marine environment, acting as a ground truth for experiments is detailed. Three experiments were developed to assess the pipeline.

Experiment 6A performed a combined microplastic analysis in a marine environment; the concentration of a combined microplastic blend was adjusted whilst the marine concentration was held constant. This experiment aimed to establish if the workflow can detect microplastic particles as a single class amongst marine matter. Experiment 6B

conducted a multi-class microplastic analysis in a marine environment; the concentration of a single type of microplastic (PE, PET, PS, or PP) was adjusted whilst a solution of marine and other plastic particles was held constant. This aimed to further classify observed microplastic samples based on their individual class types. Experiment 6C performed a PE and PMMA sphere comparison with the task of classifying similarly shaped objects based on their birefringent texture.

Experiment 6A achieved an FPR of $0.66\% \pm 0.1\%$ (for microplastics), an improvement over the 1.25% achieved by the stand-alone proposed encoder-decoder classifier. This improvement is thought to be due to the ensemble effect of the adopted pixel-wise classification scheme. Experiment 6B produced noticeable stepped concentration changes for a particular microplastic in response to when we adjusted its concentration, whilst at the same time marine and other plastic particles in the solution showed less change. Experiment 6C yielded promising results with 93.06% of PE and 79.41% of PMMA spheres being correctly classified. These results support the view that the workflow may offer further utility in fluid stream microplastic analysis.

Chapter 7, Conclusions: The final chapter of the thesis provides a summary of the work and conclusions drawn. Recommendations for future work and improvements are discussed. Finally the contributions are covered with additional comments.

Chapter 2

Literature review

This chapter presents a literature review providing details on resources consulted in this work. Aspects relating to digital holography for imaging analysis and select material properties of plastics are essential and are introduced in this section. Additionally, current state-of-the-art imaging-based microplastic classification, marine imaging flow cytometry, and general polarimetry methods are documented.

We start by referencing general literature that has provided foundational support for this study; these resources are used to develop the theoretical framework in Chapter 3. Static microplastic classification is then covered, followed by imaging flow cytometry, and polarimetry. The significance of the proposed work is then discussed, and we conclude by dealing with the gaps in the prior works.

2.1 General literature

The key areas presented in the general literature are digital holography, polarisation microscopy, and polymer birefringence. Digital holography is a technique that analyses the amplitude and phase information of a light wave. By recording the interference pattern created when a reference wave interacts with an object wave, detailed image reconstruction can occur [11]; this study focuses primarily on its use in microscopy applications. Polarisation microscopy uses polarised light to study the optical properties of materials, particularly those that exhibit birefringence or optical anisotropy [12]. By analysing how these materials interact with polarised light, polarisation microscopy can provide insight into their compositional characteristics. Polymer birefringence is the difference in refractive indices along different axes within a polymer material; this optical property results from molecular alignment or orientation during polymer manufacture and may be studied using polarisation microscopy [12].

2.1.1 Digital holography

Goodman's book [11] provides a comprehensive introduction to Fourier optics, offering a solid foundation. The literature covers classical linear techniques that are particularly relevant to holography, with the necessary support to delve deeper into this area. The seminal work by Gabor [21] is thoroughly covered. Although modern computing and sensing technologies are not discussed, the techniques presented are translatable to contemporary environments.

Holographic methods, including both in-line and off-axis configurations, are valuable for imaging applications. The off-axis method typically requires a more stable environment and is less portable due to the use of mirrors and lenses, but offers easier and more practical phase recovery methods [11]. The in-line alternative is a compact and lightweight setup that provides more stability but can produce images suffering from twin-image artefacts [22]. This occurs especially in the case of dense samples such as where the reference wave is limited and has to contend with strongly scattering object waves. An advantage of the in-line approach in a lensless configuration is that it can offer a wide field-of-view (FOV) equal to the image sensor area. This however depends on the relative sample, sensor, and light source distances [22]. Both holographic methods use numerical propagation to computationally focus on a sample object, enabling accurate imaging of the flowing particles.

In a high flow rate microfluidic environment, an "on-chip" lensless configuration may be preferred for the image sensor as it eliminates restrictions imposed by a shallow depth of field and limited focal area caused by using objective lenses in a lensed setup [16]. With this lensless configuration, the spatially resolvable resolution depends on the image sensor's pixel-pitch size. Thus, a smaller pixel size is helpful for improved resolution.

2.1.2 Polarisation microscopy

Sawyer and Grubb present microscopy methods for analysing polymers in their book [12]. The book covers a range of quantitative and qualitative techniques, providing detailed insights into the structural properties of synthetic polymers. Birefringence information is covered in detail and is particularly useful to this study. Chipman et al.'s book [23] provides the polarisation and birefringence support necessary to apply Jones calculus [19] to the tasks considered in this work.

Birefringence is a material property that exhibits different material refractive indices for light polarised in different directions; this property gives rise to the formation of double-refractive rays, which upon exiting the sample will cause interference [23]. Observing

these changes allows valuable information about the sample's optical anisotropy and molecular orientation to be obtained [12].

A benefit to performing birefringence measurements in the micron resolution domain is that they can serve as an analogue method to obtain quantitative information about the molecular orientation of an optically anisotropic sample [12]. Polarised light microscopy is commonly employed to gain insights into anisotropic structures that cannot be resolved using standard light microscopy techniques [14].

In polarising light microscopy, a transmission light microscope is used with additional components such as a rotatable stage, a polariser in the illumination system, and an analyser positioned between the objective lens and the eyepiece [12]. Typically, one or both polarisers can be rotatable, allowing the selection of light transmitted to be in a specific plane of polarisation. The polariser is dichroic, meaning it absorbs light of a particular polarisation state [23]. Often, a cross-polarisation or orthogonal configuration is used where no light is transmitted in the absence of an anisotropic object, while light is transmitted in the presence of an optically isotropic object [12].

2.1.3 Polymer birefringence and texture

Synthetic polymers exhibit different types of textural features. Two broad birefringent structural arrangements may be used to describe the synthetic polymers studied in this work and are often observed in combination. Sawyer and Grubb [12] describe these two structural arrangements as crystalline and amorphous. Crystalline structures are repeating or periodic patterns observed due to the arrangement of atoms or molecules in the crystal lattice. This leads to highly ordered chains and an optical arrangement that results in periodic or regular refractive index variations when exposed to polarised light. Amorphous structures are a less ordered arrangement in a molecular or structural sense. This leads to more random variations in the refractive indices of the material, resulting in general stochastic observations in polarisation states and refractive indices.

Crystalline structures are often used for their strength while amorphous structures can be useful for transparency and clarity applications. The manufacturing process affects these characteristics and their properties can be modified further with additives and adjustments to the production process. Therefore, it is expected that same-class polymer samples observed in the wild may vary under polarisation interrogation, not to mention the effects of degradation and biofouling that provide further noise to the observed measurement.

2.2 Consulted scientific publications

Some consulted studies are discussed in this section. Their key uses of technologies related to this project are presented and important aspects around their contributions are covered, along with any challenges encountered. This allows us to formulate the significance of our proposed work.

2.2.1 Static microplastic classification

Recent studies in developing static sample microplastic classification techniques using imaging-based methods are now covered. Three prominent studies are documented to provide background on what has been achieved and to indicate where further research can be performed. The type of imaging methods employed, classifier models, and target (or sample classes) analysed in their studies are of use in forming the foundation of our study.

Bianco et al. presented a system that used lensed holographic imaging and machine learning to classify static samples of microplastics as a single class against nine different diatom species [10]. Objects were manually segmented, and they trained a linear support vector machine (SVM) classifier on 28 global features. These features were generated from the holographic wrapped phase and included phase roughness and traditional regional metrics such as area, perimeter, eccentricity, length, and breadth as well as statistical features derived from the grey-level co-occurrence matrix. Object classification was thereafter performed on a 28-length feature vector per object. However, the holographic imaging system employed an off-axis interferometer, making it non-portable and susceptible to vibrational disturbances.

The decision was made to keep the phase wrapped due to sudden jumps and discontinuities observed when performing phase unwrapping on microplastic samples. The specimens used in their study were static samples imaged independently, allowing for accurate establishment and assignment of ground truth labels. No mixed-media tests were performed, and the classifier achieved a classification accuracy of $99.31\% \pm 0.19\%$. Nevertheless, the authors anticipated that the system would not perform similarly in the real world. They expressed concerns that microplastic specimens, such as microbeads (or microspheres), could be challenging to classify due to their morphological similarity to certain marine phytoplankton.

Bianco et al. introduced another system to improve their previous approach by investigating the fractal properties of microplastics and diatoms [9]. They used the same interferometer setup and performed the object segmentation process manually. A vector

of 13 features was extracted from the wrapped phase contrast of sample objects. These features included global measures of fractality, lacunarity, and regularity. To reduce the dimensionality of the feature vector they applied principal component analysis (PCA), resulting in a 2D mapping. Unlike their previous work, they did not rely on morphological measurements such as perimeter and area, however, they still computed a global feature vector for each object patch.

In this subsequent study, the number of unique diatom species increased to 55, while the number of microplastic types remained the same and were treated as a single class. A commercial test slide of diatoms was used and included 50 of the species studied. These diatoms were fixed between two glass slides and only one sample of each diatom class was present; the acquired data was therefore augmented to grow the dataset. The classification task was treated as a two-class approach: distinguishing diatoms from microplastics. Segmentation was performed manually, and similar to their prior work only static object classification was performed. No mixed microplastic and diatom samples were tested. The classifier achieved a classification accuracy of 98.49%.

Béhal et al. conducted a detailed investigation into utilising polarisation-sensitive features for polymer classification [24]. They explored the potential of polarisation-sensitive optical analysis of microfibers to generate rich classification features in a microscopy setting. Although they used limited raw data, they successfully classified four types of synthetic and two types of natural fibres. This study was conducted without training a classifier. Instead, it relied solely on observed values and focused specifically on derived polarisation characteristics such as birefringence.

2.2.2 Imaging flow cytometry

Studies in developing imaging flow cytometry for marine and/or microplastic applications are now covered. The fluid stream analysis component of these studies is necessary for the development of our project since we consider a marine application with biological and plastic samples.

Göröcs et al. developed a lensless imaging flow cytometer aided by deep learning for imaging phytoplankton samples in their study [16]. The device offered a flow rate of 100 mL/h, which is significantly higher than commercially available units. The system employed DIH with red, green, and blue (RGB) imaging and used a peristaltic pump to transfer the test fluid to a commercially available microfluidic channel. The fluid under examination was illuminated by single 120 μ s pulses from an RGB light emitting diode (LED) that was triggered by the image sensor. The image sensor captured three

frames per second. To aid the coherence of the LED, the spatial coherence was enhanced by employing a convex mirror; the temporal coherence was improved using bandpass spectral filters. Furthermore, the LED was overdriven due to the short excitation pulse width and the limited photon budget.

The device's primary purpose was qualitative imaging, with manual counting and logging of different types of observed phytoplankton being performed by an observer. Given the dynamic nature of the fluid sample, techniques like multi-height iterative phase retrieval were unsuitable for phase recovery [25]. Therefore, a convolutional neural network (CNN)-based approach was used for phase recovery, relying on the network training on similar waterborne objects. Field testing involved performing a net trawl to concentrate ocean water samples approximately 3000 times, which was subsequently diluted with filtered ocean water at a ratio of 1:50.

Işıl et al. further advanced the lensless imaging flow cytometer developed by Göröcs et al. to enable a phenotypic analysis of microalgae [17]. They introduced two methods to monitor the health of microalgae.

The first method involved morphological feature analysis, where statistical changes were measured by comparing histograms using Jefferies' divergence. This approach aimed to identify differences in the morphological characteristics of microalgae as an indicator of their health status.

The second method used three binary deep neural networks to form a classifier to identify three different microalgae species. The decision to employ three independent networks to determine whether an object belonged to a specific species class was made due to the difficulty in labelling unhealthy or non-class algae. Objects that did not belong to the specific class of microalgae were labelled as such, including dust, dirt, and other non-members of the target class. This approach was intended to improve performance by reducing misclassification.

The experiments conducted in the study were extensive and comprehensive. They monitored the algal population growth or decline in response to exposure to select heavy metals. Cultured microalgae samples were used for these experiments, providing controlled conditions for analysis.

Valentino et al. presented a system for lensed fluid stream detection of microfibers [13]. Building on the prior work of Bianco et al., this study employed polarisation-sensitive techniques in an imaging flow cytometry setup. A Mach-Zehnder interferometer was implemented using linear polarisers, beamsplitters, and lenses based on the theory of Colomb et al. [26] to obtain vertical and horizontal eigenpolarisation holograms. Jones

calculus [19] was applied to the propagated complex images to extract birefringent retardance and polarisation orientation features.

Sample objects for this study included virgin samples of PA6, PA6.6, PET, PP, cotton, and wool that were cut into uniform 200 μm lengths, although each class had a different diameter. An RGB rolling shutter image sensor, coupled with a 5X/0.12 objective lens, was attached to a commercially available microfluidic channel with a depth of 400 μm .

Global features were generated based on the morphological mask of the sample object. Amplitude and phase reconstructions were performed for both eigenpolarisations at the object's focus distance and the phase was kept wrapped. For classification purposes, training and testing objects were manually segmented. Shape, grey-level co-occurrence matrix, and grey-level run-length matrix features were extracted for each channel. A total of 228 features per object were obtained, which were reduced after conducting a t-distributed stochastic neighbour embedding analysis and a Pearson correlation study to eliminate highly correlated features.

A total of 432 object images were captured, with 72 images for each class. All samples were imaged separately. Standard traditional classifiers were tested. A cubic SVM achieved 87.2% accuracy as a binary classifier (synthetic versus natural fibres), while the multi-class classifier achieved 84.6% accuracy using 10% hold-out data for testing. Further testing with an ensemble of the best-performing classifiers resulted in 100% accuracy for the binary classifier and 95.3% accuracy for the multi-class classifier. The authors noted poor birefringent homogeneity in synthetic fibres and observed that natural fibres appeared more symmetric.

2.2.3 Polarimetry

Studies using PFA image sensors to extract birefringence information in a single shot are now covered. The use of these sensors in microplastic studies is not common and we therefore consult studies in the biomedical domain. Although these studies were not performed on marine or microplastic samples, they provide insight into using spatially-based polarisation sensing; this may have utility for fluid stream analysis where temporal filtering is unsuitable.

In a biomedical study, Bai et al. [15] incorporated Jones calculus [19] and applied the birefringence theory developed by Shribak and Oldenbourg [14] to display the birefringence retardance and slow axis orientation of monosodium urate (commonly known as

gout) and other crystals present in synovial fluid. While the work of Shribak and Oldenbourg relied on temporal-based polarisation filtering, Bai adopted a spatial-based polarisation filtering approach, taking advantage of the emergence of polarisation-sensitive image sensors.

In their study, Bai and the team used a PFA image sensor, an LED light source, a circular polariser, and a microscope. They developed an image processing and polarisation algorithm to support their hardware implementation. The system was designed for the biomedical domain and retrofitted to a light microscope with a $20\times/0.75\text{NA}$ objective lens. Their approach employed a single-shot method, simultaneously capturing all four linear polariser directions (0° , 45° , 90° , and 135°). This enabled quantitative analysis, providing advantages over traditional qualitative techniques in polarisation light microscopy. By utilising Jones calculus the captured intensity data could be computationally resolved, eliminating the need for a manual polariser analyser or compensator rotation to comprehensively inspect the sample's birefringence. However, this solution was only implemented for visual analysis to enhance the observer's ability to analyse the sample.

Liu et al. conducted additional research on the project to develop a deep learning-based holographic microscopy solution for gout diagnosis [27]. This work used a deep learning network to visualise the birefringent activity present in a sample. The network relied on training data comprised of specific examples of crystals.

2.3 Significance of proposed work

The following observations were made based on prior works.

Imaging flow cytometry combined with DIH techniques has been recognised as a lightweight and cost-effective fluid microscopy approach, and this is further supported by recent advancements in computer processing performance and smaller image sensor pixel sizes [16, 17]. These techniques used an RGB image sensor and have been applied as a microscopy viewer [16] and a classifier for marine phytoplankton [17] with notable performance.

Optical methods incorporating image sensors and generated features have been used in studies for microplastic classification [9, 10, 13, 24]. The optical phase was intentionally kept wrapped in all these studies due to variations in scattering. Additionally, global features were used for training classifiers. The systems involved large and less portable interferometers with lenses, which limited the potential processing throughput.

Prepared samples of synthetic microfibers were compared to natural fibres in two studies [8, 13], while another two studies examined microplastics and marine phytoplankton [9, 24]. Two of these studies analysed natural and synthetic fibres using samples of the same length but with different uniform diameters [13, 24].

Synthetic polymers exhibit birefringence [12], and techniques proposed by Colomb et al. [26] were employed in both in-stream [13] and slide-based [8, 24] settings to extract optical anisotropic features. While Valentino et al. [8, 13] trained a classifier for microplastic identification using a dual-arm interferometer, their systems lacked portability and an automated classification workflow.

Although the studies involving the classification of microplastics [8–10, 13] produced impressive classification results, none progressed toward an end-to-end workflow. Furthermore, mixed sample tests were not performed and instead these studies only reported classifier performance when tested on individual samples. Developing the experimental design may present an opportunity to deliver a more robust detector. Additionally, manual image segmentation was commonly used in these studies to obtain training and testing images for the classifier, preventing an automated environment. Investigating segmentation techniques would further assist in improving the detector workflow. It was suggested by Bianco et al. [10] that classifying morphologically similar objects such as microalgae and synthetic microbeads may pose challenges; thus, providing and assessing classification in this area is important.

The quantitative computation of birefringence using Jones calculus, as described by Shribak and Oldenbourg [14], has been used together with DIH for capturing static image slides of gout crystals [15, 27]. Specifically, Bai et al. [15] employed a PFA image sensor for their experiments. While this study did not involve microplastics, the techniques apply to studying other birefringent materials such as synthetic polymers.

The following proposed work deals with the gaps in prior works as follows:

- Developing a low-cost, high throughput method for microplastic detection that is better suited to resource-limited settings.
- Using a lensless DIH polarisation sensitive fluid stream imaging-based configuration to classify microplastics from background marine phytoplankton and debris.
- Implementing an end-to-end workflow (from image acquisition to sample measurement and logging) to classify microplastics from background marine phytoplankton and debris (ranging from 50 μm to 1000 μm in length).
- Developing a further assessment of analyser performance beyond machine learning

classification results by introducing additional experimental design.

- Providing successful classification of morphologically similar items, such as polymer microspheres and smooth marine phytoplankton by implementing a pixel-wise approach for texture-based feature vector generation.
- Implementing a short pulse length illumination circuit to reduce motion blur.

The proposed work is similar to Valentino et al. [13], which involved classifying microplastic particles in a lensed off-axis holographic imaging flow cytometry setting using polarisation-sensitive features. However, the following variations to improve outputs are proposed:

- Implementing a more stable, cost-effective imaging-based analyser based on polarisation-sensitive lensless DIH.
- Adopting a pixel-wise classification approach, rather than utilising global morphological features (to better handle morphologically similar objects).
- Establishing a system calibration protocol for accurate measurements.
- Investigating the effect of fluid flow rates and motion blur.
- Providing an end-to-end workflow that includes segmentation capable of handling changes in captured image sensor illumination intensity.
- Conducting mixed sample experiments to support the use of the analyser as a proof of concept.

2.4 Summary

This chapter reviewed literature applicable to detecting microplastics in an in-stream configuration. Gaps in the existing state of the art have been discussed with an opportunity for further development being identified. The significance of the proposed research has been presented, and the general approach as well as a theoretical framework are presented in the next chapter.

Chapter 3

Approach and theoretical framework

This chapter expands on the objectives of this study to provide a high-level overview of the development of the lensless polarisation-sensitive imaging-based FSA, microplastic classifier, and the amalgamation of these subsystems to form an automated pipeline or workflow. This is followed by information on the image sensor and flow cell. Finally, some essential theory necessary for the project application is presented, specifically in the machine learning, DIH, and polarisation domains.

3.1 Overview of approach

Developing an FSA for automated microplastic detection can be separated into key objectives that allow the detector to be realised. The following main objectives are defined:

- Development of a low-cost high flow rate lensless polarisation-sensitive imaging-based FSA subsystem.
- Development of a microplastic classifier subsystem that uses the FSA subsystem.
- Development of an automated software workflow or pipeline that combines the FSA and classifier subsystems to provide end-to-end microplastic detection.

The FSA is designed as a versatile device with potential for applications beyond its current scope. It is conceptualised as adaptable to further applications and could be used for tasks like biological detection and analysis in other environments. At its core, the hardware prerequisites for a lensless imaging-based FSA comprise three essential elements: an image sensor, a light source, and a microfluidic channel. Figure 3.1 illustrates this setup, wherein a fluid is introduced into the channel at one end and flows

through it. In the course of its journey, the fluid is both illuminated and imaged simultaneously. Eventually, the fluid exits the channel from the opposite end. This study applies polarisation-sensitive imaging and DIH to extract further object features. Both hardware and software considerations are required to support these techniques.




Figure 3.1: Illustration of an imaging-based FSA.

Microplastic classification is a specialised application that builds upon the outputs of the FSA. This process' automation hinges on computational algorithms, with machine learning being a necessary component. A comprehensive comparative analysis of suitable algorithms and tools is prudent to support the development of an automated software workflow for microplastic classification. As illustrated in Figure 3.2, sub-topics are

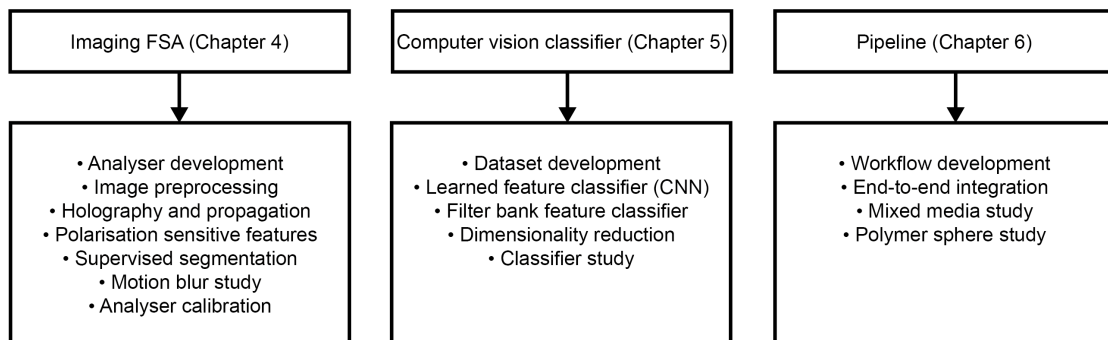


Figure 3.2: Project development.

identified to support the primary objectives. Chapter 4 focuses on implementing the FSA, providing a comprehensive overview of its development. Chapter 5 is dedicated to computer vision and the subset of machine learning techniques that support microplastic classification. While Chapter 4 develops the FSA, which may have multiple applications and can be considered a stand-alone device if required, the algorithms presented in Chapter 5 specifically address the microplastic classification problem and build on the base features produced by the FSA. The development of the classifier considers both filter

bank and learned feature extraction computer vision techniques. Chapter 6 extends the elements from the previous two chapters to create a pipeline and workflow for automated end-to-end classification of microplastic particles; this workflow is effectively able to take an input image stream from the FSA and output particle classification and measurements for logging purposes. A proposal is formed for these subtopics and developed in the following subsections.

3.1.1 Proposal for imaging-based FSA

The purpose of the FSA is to develop and implement the necessary hardware and software so that a fluid stream can be processed as follows: image the fluid stream, provide numerical propagation for the lensless holographic configuration, extract birefringence features, and perform segmentation in a high throughput manner. This subsystem outputs multi-channel spatially-aware polarisation-sensitive information that may be useful for downstream processes.

The FSA uses a prototype flow cell from Disa Scientific which provides the microfluidic channel for this system. We have chosen a monochrome PFA image sensor to capture polarisation-sensitive features. Our approach does not rely on RGB spectral data for feature generation, although the selected sensor can be substituted for a multispectral variant if required.

A light source excitation module is developed to mitigate the impact of motion blur and enable a high flow rate within the microfluidic channel. Additionally, a lensless on-chip DIH configuration is implemented that offers a stable large FOV, supported by numerical propagation techniques. To quantitatively measure birefringence, a polarisation feature extraction routine is established.

Segmenting individual foreground objects from the background when processing the image stream must be done before a detected object may be sent to a classifier for identification. Segmentation methods have been proposed using either classical computer vision [28] or learned feature neural network [20] approaches; it is vital to obtain reliable segmentation for accurate object measurement reporting. A small labelled dataset is created for supervised segmentation training. An encoder-decoder learned feature network performs segmentation. The segmentation must be reliable for varying particle concentrations, as well as changes in turbidity and mean image intensity.

The turbidity of the fluid under test may vary significantly depending on the sample's composition, particularly for diverse marine environmental samples. Therefore, it is necessary to investigate an image preprocessing algorithm to improve image SNR and

stability.

A series of experiments in the form of both performance metrics for quantitative analysis and experimental demonstrations are used to verify the proposed implementation. Analyser calibration and motion blur analysis are performed by observing known-sized calibration spheres. Tests with resolution targets are used to confirm the resolving abilities of the analyser with the implemented DIH routine. Demonstrations with birefringent samples and a depolariser are used to determine the quantitative birefringence-resolving abilities. Segmentation performance is assessed through both performance metrics and visual inspections.

3.1.2 Proposal for computer vision

The methodology for the computer vision task deals with processing an input (features developed by the FSA) and providing an output (classification result). As marine and microplastic datasets are not readily available, we must create a dataset to train the classifier effectively.

Since this project investigates microplastic detection in a marine environment, the fluid under observation contains two significant components within the size range from 50 μm to 1000 μm : microplastics and marine microorganisms such as phytoplankton. This study treats marine microorganisms and general non-microplastic debris as a single class of objects. However, microplastics are treated as either a single class or multiple classes, depending on the type of classification model that is considered. The ability to distinguish microplastics from marine microorganisms is required. Additionally, we may want to further classify detected microplastic samples according to the type of synthetic polymer.

Micronised synthetic polymers and marine phytoplankton cultures obtained for this study are used to create a simple dataset. Marine phytoplankton samples were supplied courtesy of the DFFE Marine Research Aquarium in Cape Town, South Africa and virgin micronised synthetic polymer powders were obtained from Shijiazhuang Tuya Technology, China. Additionally, microscopy calibration spheres were procured from Cospheric, USA.

Developing the microplastic classifier involves investigating two paradigms: a filter bank feature extraction and a learned feature neural network. Each approach has advantages and disadvantages, making it essential to conduct a comparative study to identify the most suitable candidate for the classification workflow. These classification methods use the multi-channel holographic polarisation-sensitive output from the FSA. A filter bank

based feature extraction is used to train classifiers such as SVMs and random forests in a pixel-wise feature vector approach. An analysis bank is created for feature extraction, with feature selection and dimensionality reduction being subsequently applied. A learned feature approach is used to implement and train an encoder-decoder CNN for pixel-wise inference.

Test datasets are used for algorithm assessment using standard machine learning performance metrics [29]. In the context of this study, there may be scenarios where the ratio of microplastic particles to marine organisms is low or the relative concentration of microplastic particles is low. In such cases it is necessary for the classifier to have a low false alarm rate or FPR, meaning it should avoid classifying real marine organisms as microplastic particles (thus leading to false positives). A low false alarm rate ensures that the classifier maintains a high level of specificity and must be appropriately examined when designing and refining the algorithm.

The sample particle sizes in this study vary from one object to the next, and in one captured frame multiple object instances and classes may be present. Therefore, a classifier must process objects of differing sizes, spanning more than one order of magnitude in length.

Filter bank feature classifier

Due to the birefringence of synthetic polymers, either stochastic or periodic textures may be available when using polarisation-sensitive analysis [12] and therefore may be used to generate features. Methods of generating features and determining appropriate features to train a classifier must be investigated. The curse of dimensionality becomes an issue when the feature vector dimension is large. A high dimensional feature space makes distance measurements difficult, and as a result meaningful metrics on which to train a classifier are hard to establish; dimensionality reduction assists with this issue. Plastics degrade to unique shapes, meaning that each microplastic particle can be considered morphologically unique and may vary significantly; global morphological and geometric features such as area and perimeter should not be solely relied upon to classify samples.

Learned feature classifier

The classifier needs to accommodate polarisation-sensitive feature channel inputs since the classifier is not being fed with a standard RGB intensity image. Multiple particles may be present in a single input image; individually classifying them accurately with morphological details is necessary. Semantic pixel-wise approaches to classification can be useful in this regard. The homogeneous appearance of smaller-sized diffraction rings

produced by particles does not naturally lend itself to CNN image-based classifiers in a multi-class exercise, as diffraction patterns may be challenging to interpret. Numerical propagation assists in bringing the object into focus, making it less morphologically homogeneous.

3.1.3 Proposal for pipeline

We develop an automated software workflow to classify microplastic particles in stream from data delivered by the FSA. The development of the automated pipeline ties into the imaging-based FSA and the selected classifiers to perform end-to-end detection, classification, and logging of microplastic particles. In other words, the end-to-end process obtains an image frame as an input and performs preprocessing, segmentation, classification, and logging of each particle that passes through the analyser within the specified size range. Moreover, this system must process a stream of input frames from an acquisition session or experiment. The workflow must be assessed with different microplastic concentrations (with low concentrations being of particular interest). The key attribute of this subsystem is the ability to provide automated processing.

3.2 Hardware information

We briefly detail the image sensor and flow cell used in this study, as the additional components and software were chosen and implemented to complement these items. The image sensor was specifically employed for single-shot polarisation-sensitive fluid stream analysis and is therefore a vital part of the FSA design. The flow cell and customisation thereof are also important for the FSA development and are introduced accordingly.

3.2.1 Image sensor information

Owing to the birefringence of synthetic polymers [12], measuring them effectively is requisite as relying solely on intensity-only measurements from an image sensor may not provide enough information on the sample to train a classifier effectively. We therefore employ a linear polarisation image sensor as a basic component to support birefringence sensing. This project moves away from using an RGB colour image sensor and instead adopts a monochrome sensor. This removes any possible reliance a classifier may have on RGB spectral-based measurements that can cue a classifier to target spectral components for classifying microplastic classes. As microplastics vary in colour, the likelihood of poor performance in a real-world setting would increase if the classifier relied on colour cues. Additionally, gains in SNR are achieved by removing spectral filters on the sensor.




Figure 3.3: Sony polarisation image sensor illustration [30].

Compact and affordable sensors have become available with the recent introduction of PFA image sensors in the machine vision market. Figure 3.3 illustrates the Sony IMX264MZR monochrome sensor [30], which offers four distinct nano-wire grating linear polariser orientations (0° , 45° , 90° , and 135°) per block, known as a superpixel. Each pixel on the sensor measures $3.45 \mu\text{m} \times 3.45 \mu\text{m}$, providing polarisation-sensitive intensity information through spatially arranged filters. This sensor offers an enhanced capability compared to a standard image sensor as it enables capturing polarisation information alongside traditional intensity data. Using this sensor means that linear polarisation orientations of incident light may be resolved in a single shot.

3.2.2 Flow cell information

To provide a high throughput environment for fluid samples, the chosen lensless DIH configuration is used with a prototype flow cell (Disa Scientific (Pty) Ltd, South Africa). The flow cell provides the microfluidic observation channel and is not supplied with any lighting or imaging components. Although commercial channels are readily available, this unit allows for customisation in that the channel depth and width may be adjusted. Furthermore, 3D-printed parts are designed and retrofitted to the unit to position the image sensor and light source. Standard commercial microfluidic channel offerings are typically limited to a few fixed channel widths and may not perfectly match the FOV dimensions of this image sensor [31]. This limits throughput if the channel width is narrower than the sensor width, with part of the image sensor not being used during data acquisition. On the other hand, if the channel is too wide, some particles may not pass over the sensor, leading to incomplete data capture [16]. Additionally, an unsuitable fixed channel depth may limit sample throughput.

Figure 3.4(a) shows the unit in a base configuration. The optical configuration may be likened to bright-field transmission mode microscopy. Figure 3.4(b) shows a cutaway of the unit where the observation window is present. The unibody base of the flow cell




Figure 3.4: Illustrations of the prototype flow cell. (a) Prototype flow cell. 1. Fluid entry point. 2. Cavity housing for an illumination source. 3. Fluid exit point. 4. Cavity housing for an image sensor. (b) Prototype flow cell cutaway showing round glass observation window.

provides a basic configuration for the FSA and can be customised by stacking spacer rings to achieve the desired cavity dimensions. The round glass window used in the flow cell is 2 mm thick borosilicate glass with a diameter of 45 mm and is removable if servicing is required.

3.3 Theoretical framework

This section presents theory drawn upon in this study and provides background on the concepts and methodologies employed. Selected machine learning, DIH, polarisation birefringence, and motion blur topics are presented.

We start by covering background information on encoder-decoder networks; this architecture is useful in segmenting particles and is also used for some classification methods. This is followed by filter banks, which are useful for feature extraction used for individual microplastic type classification. Next, lensless DIH is covered, which is necessary for the structural layout of our optics and provides support for both hardware and software development. Polymers and birefringence are then covered, which is again important for the structural layout of our optical and analyser development. Motion blur is finally covered to support the high throughput requirements of our device.

Additionally, and where relevant, further resources and recommended materials are provided for readers seeking a more in-depth understanding of the theory or further

exploration of the topics discussed.

3.3.1 Encoder-decoder networks for segmentation and classification

Encoder-decoder networks are useful for segmentation and classification tasks covered in the development of the FSA. Therefore, we will briefly introduce key concepts relevant to the subsequent chapters. This network structure was chosen because it is common for multiple objects to coexist in a single scene. Therefore, this fine-grained pixel-wise inference supports the detailed shape measurement analysis necessary for logging purposes.

The global interest in and adoption of neural network approaches for image segmentation provides an alternative to rule-based segmentation. A neural network can be supervised to learn how to segment an image or region. In the supervised learning paradigm, pairs of input examples and ground-truth labels form training data; this training data is processed to construct a function that maps new input data to a predicted or expected output label [32]. This means that instead of using rule or heuristic-based traditional methods, a neural network can be trained to learn and optimise the segmentation process by being provided with a raw example image (input) and a labelled segmentation mask (expected output). Similarly, for classification purposes, this structure may be trained in a multi-class format that provides pixel-wise inference for tasks beyond separating background from foreground.

Using neural networks to perform segmentation or classification relies on having sufficient training data: there is typically a correlation between the amount of well-curated training data and the network's performance [32, 33]. Although trainable approaches are time-consuming to train initially, they are considered a once-off process and actual inference speeds can be fast depending on the network structure. A widely recognised network model for pixel-wise semantic segmentation is the encoder-decoder network U-Net [20]. The standard implementation (depicted in Figure 3.5) is fully convolutional and contains no fully connected layers. In this architecture the input image tile undergoes a sequence of operations, starting with convolutions followed by rectified linear unit (ReLU) activations in the encoding section. The image tile is halved and the number of filters (or feature channels) is doubled for each sequential block. This encoding phase captures context and compresses the input. In the decoding section the network employs sequential transpose convolutions to expand the input to the same size as the output, facilitating localisation. In each sequential decoding stage the tile size is doubled and the number of filters is halved. Skip connections are used to merge features between the

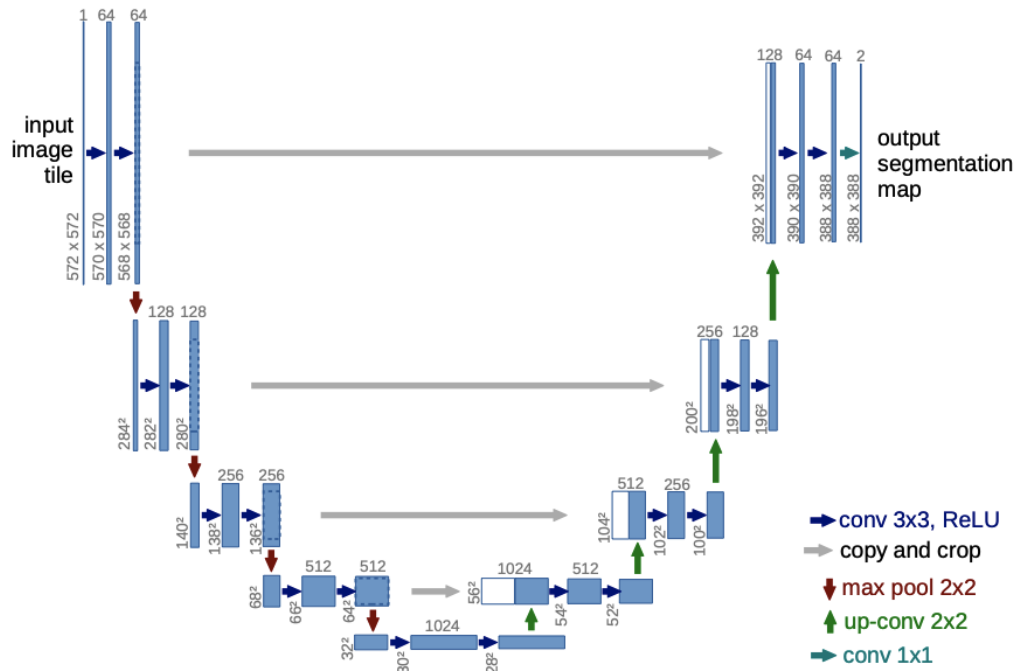


Figure 3.5: Standard U-Net structure [20].

encoding and decoding blocks to provide fine-grained detail or resolution. The architecture provides an output tile size that matches the input tile size, enabling pixel-wise segmentation or classification.

Residual U-Net (ResUnet) [34] is a fully convolutional network that combines features of U-Net [20] and residual learning. It uses U-Net as a foundation and incorporates pre-activated residual blocks or skip connections into the model, substituting some convolutions from the original architecture. A benefit of combining the encoder-decoder and residual ideologies is that U-Net assists with capturing fine-grained details and offers compression through a series of encoder and decoder blocks, while residual learning adds shortcut connections that may assist with potential vanishing gradient problems that can occur in deeper neural networks. Res-Unet combines high-level feature extraction with gradient flow optimisation. Skip connections are strategically positioned between the encoder and decoder layers, facilitating the fusion of high and low-level features, which is particularly useful for localisation, segmentation, or classification.

3.3.2 Filter banks

Filter banks incorporated in this work are briefly introduced here and were chosen specifically for the task of texture analysis for the microplastic classifiers. When choosing filters for texture analysis, it is beneficial to include spot and bar filters at various orientations and scales to accommodate different border applications [35]. Having equally

spaced orientational filters provides periodic feature maps oriented around objects of interest, thus enhancing the ability to capture texture information from various angles. Subsequent tasks like feature selection, dimensionality reduction or other types of compression may be performed at a later stage. Consequently, it is important initially to provide broad coverage across the range of objects or textures of interest without excessive optimisation of the filter selection; this ensures flexibility and adaptability in the later stages of processing.

Gabor filters

These filters, originally introduced by Gabor [36], can be employed in their 2D form as feature extraction filters for images. They comprise a kernel with an oriented sinusoid multiplied by a symmetric Gaussian function [35]. This design allows them to capture both spatial and frequency characteristics, making them useful as texture discriminators in image processing. Notably, Gabor functions have been used to model simple cells in the mammalian visual cortex [37], and the frequency and orientation representations of Gabor filters are believed to resemble those of the human visual system [38].

A 2D Gabor filter in its basic form is parameterised (and can be defined by its impulse response) as follows [39]:

$$g(x, y) = \exp\left(-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2}\right) \cos\left(2\pi\frac{x'}{\lambda} + \phi\right), \quad (3.1)$$

where (x, y) are the spatial coordinates and (x', y') are the rotated coordinates that can be obtained using the given transformations

$$x' = x \cos(\theta) + y \sin(\theta) \quad \text{and} \quad y' = -x \sin(\theta) + y \cos(\theta). \quad (3.2)$$

The remaining parameters include γ , which represents the aspect ratio of the elliptical Gaussian envelope, λ for the wavelength of the sinusoidal wave, ϕ for the phase of the sinusoidal wave, θ for the filter orientation angle in degrees, and σ for the standard deviation of the Gaussian envelope.

A Gabor filter is a tunable filter that is rotationally asymmetric, making it orientation-sensitive. By creating a filter bank with multiple orientations and scales of Gabor filters, textures at different scales and orientations can be simultaneously analysed. This is especially useful for adapting to slight changes in objects detected in a real-world setting. These filters capture localised spatial frequency variations that are helpful for both fine and coarse texture analysis. Each filter can provide responses that serve as a discriminative representation of texture patterns within an image, producing highly

useful features for a classifier.

Root filter set

The root filter set (RFS), as described by Varma and Zisserman [40], is akin to the established Leung-Malik (LM) filter set [41]. It encompasses a total of 38 filters, comprising:

- Anisotropic filters: including the first derivative of Gaussian (edge) and the second derivative of Gaussian (bar), with six orientations and three scales each; and,
- Isotropic filters: consisting of a Gaussian and a Laplacian of Gaussian.

Edge filters detect changes in intensity along the alternate axis — if applied in the x -direction (horizontal), they will detect changes in intensity along the y -direction and highlight vertical edges. Bar filters detect changes in intensity along the alternate axis — if applied in the x -direction (horizontal), they will detect changes in intensity along the y -direction and highlight vertical bar-like or ridge structures. Notably, the isotropic filters, Gaussian and Laplacian of Gaussian, can be condensed into a subset of eight maximum response filters to achieve rotational invariance.

The Gaussian and Laplacian of Gaussian filters are defined in order as follows:

$$G(x, y, \sigma) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \quad (3.3)$$

and

$$\nabla^2 G(x, y, \sigma) = \left(\frac{x^2 + y^2}{\sigma^4} - \frac{2}{\sigma^2}\right) G(x, y, \sigma), \quad (3.4)$$

where σ is the standard deviation scale of the Gaussian envelope, and x and y are coordinates of the filter kernel. The first and second order derivatives of Gaussian filters in the x -direction are defined in order as follows:

$$\frac{\partial G(x, y, \sigma)}{\partial x} = -\frac{x}{\sigma^2} G(x, y, \sigma) \quad (3.5)$$

and

$$\frac{\partial^2 (G(x, y, \sigma))}{\partial x^2} = \left(\frac{x^2}{\sigma^4} - \frac{1}{\sigma^2}\right) G(x, y, \sigma). \quad (3.6)$$

S filters

The S filter bank is based on a dense texture Gabor-like filter that is rotationally invariant [42]. The filter is described as

$$F(x, y, \tau, \sigma) = F_0(\tau, \sigma) + \cos\left(\frac{x^2 + y^2 \pi \tau}{\sigma}\right) \exp\left(\frac{x^2 + y^2}{2\sigma^2}\right), \quad (3.7)$$

where τ is the number of cycles of the harmonic function within the Gaussian envelope of the filter and $F_0(\tau, \sigma)$ is the zero DC component. Although this filter is not used in the base filter bank of this study, it is used as an extended filter set to test the ability of the workflow to rely solely on internal textures when comparing two geometrically similar but different plastic types.

3.3.3 Lensless DIH

We introduce concepts and theory of lensless DIH in a static sample context but they are transferrable to imaging fluid stream analysis. Traditional light microscopy relies on the magnification power of an objective lens to achieve high-resolution imaging. However, this approach has limitations, including a reduced FOV due to the magnified area of interest excluding the rest of the sample area from view. This means that covering a wider working area (or FOV) requires significant processing time. Additionally, the shallow depth of field further restricts the system's ability to process samples with a limited focal range.

Technological advancements have facilitated both the development of image sensors with smaller pixel-pitch sizes and improved the processing capabilities of computers and embedded devices [22]. These advancements have created favourable conditions for implementing lightweight, cost-effective, and simplified hardware designs coupled with image reconstruction algorithms and digital image processing [43,44]. Lensless DIH takes advantage of these advancements and uses them to address the challenges presented by lens-based microscopy. Numerical reconstruction techniques replace an objective lens to achieve focused imaging on the object of interest [45]; this creates a virtual or numerical lens of sorts. The technique relies on a coherent or partially coherent light source and an image sensor, so it is compact and lightweight [46]. For further information, readers are directed to an extensive Fourier optics resource [11].

Although different configurations of this method exist, this work provides information on an on-chip configuration where the light source to object distance (z_1) is substantially larger than the object to image sensor distance (z_2) [43,44]. The close proximity of the object and sensor means that the distance z_2 is usually no more than a few millimetres,

leading to its “on-chip” descriptor; the distance z_1 is generally a few centimetres [22]. An example of this configuration is shown in Figure 3.6. By eliminating the need




Figure 3.6: Basic on-chip DIH setup.

for an objective lens, lensless DIH allows for a wider FOV compared to traditional microscopy [46]. Using the entire image sensor plane as the FOV means that modern image sensors can range from $5 \times 10^{-4} \text{ m}^2$ to $5 \times 10^{-2} \text{ m}^2$. In contrast, an objective lens typically offers a FOV range from $0.1 \times 10^{-5} \text{ m}^2$ to $1 \times 10^{-5} \text{ m}^2$ (for similar resolution).

Angular spectrum of plane waves

As described by Goodman [11], we can formulate the angular spectrum method with a plane wave approximation. Due to the on-chip configuration of the FSA, the distance z_1 between the light source and the object is substantially larger than the distance from the object to the image sensor z_2 , which allows for the Fresnel plane wave approximation to be employed. Therefore, the Fresnel diffraction integral approximates the phase and amplitude variations across the wavefront.

We illustrate an ideal model using coherent collimated illumination and consider a spatially sensitive image sensor that is noiseless with infinitely small pixels. The light wave complex amplitude incident on an image sensor plane with spatial coordinates (x, y) can be described by the superposition of the reference wave A_{ref} and the scattered wave

from the object $a_{obj}(x, y)$:

$$U(x, y) = A_{ref} + a_{obj}(x, y). \quad (3.8)$$

The complex amplitude A_{ref} is a non-spatially varying strong uniform plane wave; in contrast, the complex amplitude $a_{obj}(x, y)$ is weakly scattering and spatially varying. Image sensors only recover the intensity distribution of the light wave, which means that the absolute phase is lost. The intensity incident on the image sensor may be described as [11]

$$I(x, y) = |A_{ref} + a_{obj}(x, y)|^2 = |A_{ref}|^2 + |a_{obj}(x, y)|^2 + A_{ref}^* \cdot a_{obj}(x, y) + A_{ref} \cdot a_{obj}^*(x, y), \quad (3.9)$$

where $A_{ref}^* \cdot a_{obj}(x, y)$ is proportional to the scattered wave, $|A_{ref}|^2$ is a constant, and $*$ is the complex conjugate of the complex magnitude. Both $A_{ref}^* \cdot a_{obj}(x, y)$ and $|A_{ref}|^2$ are desired terms whereas $A_{ref} \cdot a_{obj}^*(x, y)$ and $|a_{obj}(x, y)|^2$ are undesired artefacts due to the intensity image not containing phase information. The $A_{ref} \cdot a_{obj}^*(x, y)$ term is commonly referred to as the twin-image artefact — an out-of-focus representation of the desired in-focus term. Similarly, if $A_{ref} \cdot a_{obj}^*(x, y)$ were to instead become the desired term and be propagated accordingly, $A_{ref}^* \cdot a_{obj}(x, y)$ would become an undesired and thus an out-of-focus term. The reconstruction of the object improves when the light scattered by the object is significantly less than the reference wave $A_{ref} \gg a_{obj}(x, y)$, in which case just the defocused term $A_{ref} \cdot a_{obj}^*(x, y)$ is undesired and affects the complex distribution at the image sensor plane.

To propagate from the image sensor plane to the object, we can employ the angular spectrum method [11]. This provides a Fourier optics approach to modelling propagation and diffraction of coherent light inside a homogeneous medium. A given complex optical wave field on the plane $z = 0$ is described as $U(x, y; 0)$. The angular spectrum of $U(x, y; 0)$ is given by its 2D Fourier transform:

$$A(f_x, f_y; 0) = \int \int U(x, y, 0) \cdot \exp[-i2\pi(f_x x + f_y y)] dx dy, \quad (3.10)$$

where $i^2 = -1$.

The transfer function, $H(f_x, f_y; z_0)$ of wave propagation acting as a linear space-invariant

system is given by

$$H(f_x, f_y; z_0) = \begin{cases} \exp \left[i2\pi \frac{nz_0}{\lambda_0} \sqrt{1 - \left(\frac{\lambda_0 f_x}{n} \right)^2 - \left(\frac{\lambda_0 f_y}{n} \right)^2} \right] & \sqrt{f_x^2 + f_y^2} < \frac{1}{\lambda_0} \\ 0 & \text{otherwise,} \end{cases} \quad (3.11)$$

where λ_0 is the optical wavelength in a vacuum and n is the refractive index of the medium.

The angular spectrum of the optical field on a given parallel plane $z = z_0$ can be calculated by

$$A(f_x, f_y; z_0) = A(f_x, f_y; 0) \cdot H(f_x, f_y; z_0). \quad (3.12)$$

Finally, the optical wave field on the plane $z = z_0$ is obtained by the 2D inverse Fourier transform of its spectrum:

$$U(x, y; z_0) = \int \int A(f_x, f_y; z_0) \cdot \exp [i2\pi(f_x x + f_y y)] df_x df_y. \quad (3.13)$$

3.3.4 Polarisation and birefringence

Anisotropic objects exhibit optical activity that depends on the polarisation state of the incident light or, similarly, the object's orientation with respect to the incident light. Polarisation-sensitive sensing is helpful in the microplastic classification problem, as the detection of optically anisotropic features inherent in synthetic polymers can help reduce the reliance on traditional intensity-derived features that may be less robust [12]. However, as particles in suspension travel through the microfluidic channel, their orientations change. This may lead to misclassification if polarisation features are only captured at a certain orientation because the classifier might be unable to match the sample feature against its training features. It is therefore advisable to obtain a full orientation birefringence map of the sample [14]. In the context of this project, it is necessary to detect and register particles' optical activity in situ for all possible 2D orientations. Particle reorientation during flow is not feasible, highlighting the importance of capturing birefringent optical activity across various particle orientations.

Since the fluid stream analysis involves dynamic samples in motion, using temporal-based polarisation filtering is not feasible. Instead, spatially-based polarisation filtering becomes necessary. These hardware considerations can be combined with a software backend to resolve 2D birefringent activity (including retardance and slow axis orientation) [14].

Birefringence formulation

The following important polarisation definitions from Chipman et al. [23] are used in this study:

- Eigenpolarisation: a wave polarisation state that is the same polarisation state when entering and exiting a system or optical element, but where changes in amplitude or absolute phase can occur. Eigenpolarisations are eigenvectors of the corresponding Jones or Mueller matrix.
- Retardance (Δ): the optical path difference between eigenpolarisations, commonly specified in radians.
- Slow axis (θ): eigenpolarisation with a larger optical path length, which is analogous to the mode with a higher refractive index in a material, commonly specified in radians.
- Fast axis: eigenpolarisation with a smaller optical path length, which is analogous to the mode with a lower refractive index in a material, commonly specified in radians.
- Birefringence: material property that creates retardance associated with propagation, which can apply to transmission, reflection, diffraction, or scattering.

Formulating birefringence information into quantitative measurements was studied by Shribak and Oldenbourg [14] for use with temporal polarisers and traditional image sensors. A similar technique was recently applied to biomedical applications using a PFA sensor instead of multiple temporally separated images to quantitatively resolve polarisation birefringence information [15]. The transition from temporal to spatial-based techniques enables a frame to capture four linear polarisation orientations, facilitating the calculation of birefringence information in a single shot.

We employ this spatial technique in the FSA and now describe the required components and theory. To model the system using Jones calculus, the resulting output field is calculated using a light source and three elements: a left-hand circular polariser (generating polarised light), an object under interrogation, and a spatially-aware PFA sensor. A linear combination of Jones matrices [19] implements these elements. Each Jones matrix can be described modularly and sequentially combined to form the final system. The convention in this paper is from the viewpoint of the receiver. As such, starting

with the first component, a left-hand circular polariser is described as

$$\mathbf{E}_{\text{LHCP}} = \frac{E_0}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}, \quad (3.14)$$

where E_0 is the amplitude of the incident light.

The birefringent sample is modelled as the Jones matrix

$$\mathbf{M}_{\text{object}} = \begin{pmatrix} \cos\frac{\Delta}{2} + i\sin\frac{\Delta}{2}\cos 2\theta & i\sin\frac{\Delta}{2}\sin 2\theta \\ i\sin\frac{\Delta}{2}\sin 2\cos 2\theta & \cos\frac{\Delta}{2} - i\sin\frac{\Delta}{2}2\cos\theta \end{pmatrix}, \quad (3.15)$$

where Δ is the retardance magnitude and θ is the azimuth or slow axis orientation.

The Jones matrices for four different orientations of linear polarisers can be represented as

$$\mathbf{M}_{\text{detector}} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, \quad \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \quad \text{or} \quad \frac{1}{2} \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}, \quad (3.16)$$

for orientations 0° , 45° , 90° , or 135° , respectively.

The product of Jones matrices that describe the final electric field is

$$\mathbf{E}_{\text{out}} = \mathbf{M}_{\text{detector}}\mathbf{M}_{\text{object}}\mathbf{E}_{\text{LHCP}}. \quad (3.17)$$

As an example, the product of Jones matrices that describe the final electric field for a 0° linear polariser is

$$\mathbf{E}_{\text{out}}^0 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \cdot \begin{pmatrix} \cos\frac{\Delta}{2} + i\sin\frac{\Delta}{2}\cos 2\theta & i\sin\frac{\Delta}{2}\sin 2\theta \\ i\sin\frac{\Delta}{2}\sin 2\cos 2\theta & \cos\frac{\Delta}{2} - i\sin\frac{\Delta}{2}2\cos\theta \end{pmatrix} \cdot \frac{E_0}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}. \quad (3.18)$$

The image sensor only captures the intensity of the electric field and can be described for the 0° linear polariser as

$$I_0 \propto |\mathbf{E}_{\text{out}}^0|^2 = \mathbf{E}_{\text{out}}^0 \cdot \mathbf{E}_{\text{out}}^{0*}. \quad (3.19)$$

Relating the captured intensity to birefringence measurements is treated as follows:

$$\begin{aligned}
I_0 &= \frac{I}{2}(1 - \sin\Delta.\sin2\theta), \\
I_{45} &= \frac{I}{2}(1 + \sin\Delta.\cos2\theta), \\
I_{90} &= \frac{I}{2}(1 + \sin\Delta.\sin2\theta), \\
I_{135} &= \frac{I}{2}(1 - \sin\Delta.\cos2\theta),
\end{aligned} \tag{3.20}$$

where I is the intensity of the incident illumination source, and I_0 , I_{45} , I_{90} , and I_{135} are the respective 0° , 45° , 90° , and 135° oriented linear polarisation intensity channels from the image sensor.

We can introduce the following auxiliary variables:

$$\begin{aligned}
A_1 &= \frac{I_{90} - I_0}{I_{90} + I_0} = \sin\Delta.\cos2\theta, \\
A_2 &= \frac{I_{45} - I_{135}}{I_{45} + I_{135}} = \sin\Delta.\cos2\theta.
\end{aligned} \tag{3.21}$$

Finally, retardance (Δ) and azimuth or slow axis orientation (θ) is computed as

$$\begin{aligned}
\Delta &= \arcsin\sqrt{A_1^2 + A_2^2}, \\
\theta &= \frac{1}{2}\arctan\frac{A_1}{A_2}.
\end{aligned} \tag{3.22}$$

3.3.5 Motion blur

When an object or particle is present at the analyser's observation window, it can be captured by the image sensor. It is important to ensure that the analyser supports a suitable flow rate to function effectively as a high throughput device. However, there is a trade-off: increasing the flow rate can result in objects traversing multiple pixels during a single frame exposure, thus generating motion blur artefacts. One solution is to operate the image sensor at a shorter exposure time, but this approach is constrained by the camera's minimum required exposure time. Moreover, decreasing the exposure time might require an increase in the intensity of the light source, given that the exposure takes place over a limited time and therefore only has access to limited incident photons.

Another strategy to mitigate the motion blur is to increase the flow channel depth, as this would reduce the mean flow velocity while maintaining the desired flow rate. However, increasing the channel depth can negatively impact the holography reconstruction quality of the analyser, and the assumption of a small object-to-image sensor distance

(z_2) might not hold. A deeper channel may also increase the likelihood of particle occlusions, as multiple objects with the same x, y coordinates but at different z coordinates (i.e. depths) in the flow channel could obscure each other from image sensor acquisition. An alternative is to reduce the exposure time by pulsing the illumination source; this benefits the analyser system by overcoming the image sensor constraints and is a suitable workaround to achieve shorter effective exposure time and reduce motion blur.

A simplified model can be generated to calculate the effect of motion blur. This model ignores the effects of any other lighting and assumes that the particle flows in a straight line, perpendicular to the optical axis, and is imaged by a global shutter image sensor. We can calculate the maximum exposure time (T) required to achieve motion blur of less than a pixel in a given scenario:

$$T = \frac{P}{2V}, \quad (3.23)$$

where V is the particle's velocity in meters per second, P is the pixel pitch of the image sensor in meters and T is the exposure time of the camera sensor in seconds. We set the blur value to half the time it takes to traverse one pixel to align with the Nyquist-Shannon sampling theorem.

As an example, if one works with a required flow rate (Q) of 20 mL/minute and a flow channel depth (z) of 500 μm and a channel width (w) of 7.07×10^{-3} m, we can calculate the particle velocity. Assuming plug flow fluid transport, we solve for the particle velocity (V) as follows:

$$V = \frac{Q}{A}, \quad (3.24)$$

where A is $w \times d$.

If we aim to limit blur to less than one pixel and allow for a 20% margin, an exposure time of at most 14 μs is needed with sampling considerations. Current polarisation-sensitive image sensors offer minimum exposure times of the order 30 μs to 100 μs . Therefore, implementing a light source pulsing circuit to create a reduced effective exposure time environment would be beneficial. The consideration of overdriving the light source might also be necessary given the short pulse duration and limited photon budget.

3.4 Software development and workstation specifications

General code was developed in MATLAB[®]. Neural network modelling was developed and performed in Python and Keras Tensorflow.

The following digital workstation was used for this study:

Intel[®]Core[™] i7-12700 3.6 GHz, NVIDIA[®]GeForce[™] RTX 3050 8 GB, 32 GB RAM DDR4.

General software environments:

Microsoft Windows 11 Pro, MATLAB[®]/ Simulink[®]2022a, Keras Tensorflow 2.8.0, NVIDIA[®]CUDA[®]11.3.58, Python 3.9.7, ImageJ 2.3.0.

PCB design: KiCad version 7.0.

CAD and 3D printing: Autodesk Fusion 360, Prusu Slicer, Prusa Mini + 3D printer.

3.5 Conclusion

This chapter detailed the project approach which provided the developmental foundation for the FSA in Chapter 4, the classifier in Chapter 5, and the pipeline or workflow in Chapter 6. Furthermore, we provided theoretical support for the development of each subsystem. In the next chapter, we move to the first stage of development, which is the complete implementation and analysis of the FSA.

Chapter 4

Imaging fluid stream analysis

Flow cytometers lack detailed morphological information about particles and the ability to simultaneously accommodate particles of varying sizes. Imaging flow cytometry addresses these limitations, but such devices are often costly, large, and have limited throughput. This chapter introduces an FSA developed for particle analysis. It incorporates DIH and spatially aware polarisation sensing. The FSA is configured in a lensless layout to facilitate high-throughput sample analysis and is equipped with sample segmentation capabilities for isolating particles within the fluid stream. This work includes demonstrations and performance testing of our device. We assess the numerical propagation capabilities of DIH and conduct a resolution test using a standard resolution target. Additionally, we determine the FSA's ability to resolve birefringent particles by validating it with a liquid crystal depolariser. The segmentation performance is evaluated using samples of marine phytoplankton and birefringent microplastics. A calibration procedure with two types of scientific microspheres (including one unseen polymer type and size) is also presented. This study demonstrates the versatility and adaptability of the FSA, which can be configured for various sensing requirements.

4.1 Preliminaries

Traditional optical fluid stream analysis has typically been performed using devices such as flow cytometers for cell analysis or particle analysers for general material analysis. These devices often rely on coherent excitation and several band-pass photodetectors to obtain distinct spectral signatures of fluid contents [47]. Regularly used in the biomedical field for cell counting and analysis, these techniques face limitations when dealing with complex mixtures of heterogeneous objects. Since cells are analysed in a single-file manner, these devices are less suitable for concurrently analysing diverse cell sizes and

cannot provide detailed information about cell or object shape and geometry [47]. As an analyser is configured for a specific sample type, reconfiguration is required if different sample sizes are interrogated.

This study investigates the development of an analyser for both biological and synthetic particles within a marine context. Furthermore, the size range of interest demands an approach that can accommodate various object sizes at a given time. Based on these considerations and the intersection of cell and synthetic material analysis, we adopt the term for our device as a “fluid stream analyser” (FSA) which is developed in this chapter. Imaging-based solutions tend to be computationally demanding, which has been an obstacle in this domain. However, leveraging recent advancements in computational capabilities alongside machine learning and computer vision tools provides a pathway for making this technique more accessible and practical.

Traditionally, imaging flow cytometers are configured with objective lenses that lead to a limited flow rate and depth of field [16]. These devices are costly and primarily offer transmission-mode (and sometimes fluorescence) intensity measurements. As an alternative option, the device presented in this chapter offers a low-cost implementation of an FSA that is lensless for high throughput analysis. This is combined with holographic techniques, allowing for greater field depth. In addition, this system is demonstrated with polarisation-sensitive abilities due to its linear PFA image sensor that is coupled with an illumination source and software backend to compute birefringence parameters.

This chapter starts with the proposed analyser development that covers the hardware configuration of the FSA which includes the flow cell, optical components, and supporting electronics. Next, the software workflow is detailed by covering the image preprocessing routine, followed by implementing the DIH and birefringence resolving workflow. Particle segmentation from background is then studied in which a supervised learned encoder-decoder network is deployed. We used marine phytoplankton and polymer samples to create a small dataset for network training. The implementation delivered favourable segmentation performance. We demonstrate the resolving ability of the system through the use of a resolution test chart, as well as a visual example. The birefringence workflow is verified by using a liquid-crystal depolariser and demonstration on polymer samples. Example outputs of the supervised segmentation model are displayed and discussed.

The issue of motion blur and its reduction is then covered, in which the custom-developed light pulse engine (LPE) is compared against the standard non-assisted illumination method. The ability of the FSA to perform accurate measurements is a priority and

this includes the 2D projected area of particles. A calibration protocol for the FSA was implemented and performed with PE scientific microspheres. We assessed the calibration with two different sample types including an unseen polymer type and diameter (PMMA, $53\ \mu\text{m}$ to $63\ \mu\text{m}$).

The FSA appears to therefore demonstrate versatility and adaptability, which can be configured for various sensing requirements. It is particularly suited for resource-constrained environments and offers a low-cost alternative to traditional flow-based microscopy.

4.1.1 High-level overview of propagation




Figure 4.1: Polarisation propagation. Light propagates from the excitation source as randomly polarised light (RPL). It then passes through the circular polariser to produce circularly polarised light (CPL). Next, it illuminates an optically anisotropic sample (in this example) to become elliptically polarised light (EPL), and finally this light is captured by the PFA. The PFA shows the four different linear polarisation orientations per a super-pixel.

The high-level operation of the FSA is detailed to give the reader a sense of the important components in the signal chain. This chapter develops these high-level concepts into realising the FSA. This layout is incorporated into the flow cell to analyse moving fluids. Figure 4.1 shows the basic polarisation propagation. Working through the signal chain, we commence with the illumination source, namely a laser diode that is pulsed by a custom-developed excitation module. The laser diode is mated to a single-mode fibre optic line, producing partially coherent light. This generated light enters a left-handed

circular polariser, converting it into left-handed circularly polarised light. This circularly polarised light then propagates to the sample (which is within a flowing fluid), where its polarisation state may either remain unaltered or change to elliptically polarised light if the material is optically anisotropic. Finally, the light reaches the image sensor, which is spatially sensitive to four different linear polarisation directions. The measured intensity is then processed through our developed workflow to extract polarisation-sensitive measurements.


4.2 Materials and methods

4.2.1 FSA


The basic components of the analyser comprise the image sensor and the flow cell. The rest of the components and associated software were designed to complement and support these core elements. A polarisation-sensitive camera (MER2-503-36U3M-POL, Daheng, China) was selected due to the birefringent samples investigated in this work; it was also the most affordable unit with polarisation sensitivity that could be found in the market to fulfil the low-cost aspect of the project. This camera has a monochrome PFA sensor (IMX264MZR, Sony) and uses a complementary metal-oxide-semiconductor array with four linear polariser orientations. The sensor has an 11 mm diagonal format with a pixel pitch of 3.45 μm , offering a 2448×2048 pixel resolution. Communication is through general purpose input/output and USB 3.0 for data acquisition, configuration, and trigger control. Other sensor options may be employed depending on the specific application. The FSA has been tested with standard monochrome and RGB sensor variants.

The prototype flow cell, sourced from Disa Scientific, was configured with a channel depth of 500 μm . Normally, copper channel shims are stacked to set the channel depth by placing them between the two observation windows. However, in this specific case, where the image sensor is rectangular and the observation window is circular, the shims were engineered to match the FOV dimensions of the image sensor, i.e. the sensor plane. These optimised shims now serve two purposes, firstly to set the required depth of the fluid channel, and secondly to ensure that the fluid only traverses the image sensor's active area and not around it. This, in combination with the lensless configuration, further supports a high sample throughput.

Particles passing through the flow cell are present at varying depths within the channel; this may lead to occlusions if particles are present at the same (x,y) spatial location but at different depths (z) especially at high particle concentrations. An alternative



(a) FSA overhead. 1. Peristaltic pump. 2. Light pulse engine. 3. Fibre-coupled laser diode (with polariser and mountings inside). 4. Additional spacer cylinders. 5. Main flow cell.



(b) FSA alternative view. 6. Fluid under test entry point to flow cell. 7. Mid-line of the flow cell, location of the fluid channel inside the flow cell (where fluid is illuminated and imaged). 8. Fluid under test exit point. 9. Image sensor and mountings.

Figure 4.2: The realised FSA from two different views indicating important components and areas.

is to use a shallower channel depth, but that would restrict the achievable flow rate due to motion blur and introduce limitations on the maximum object size fractions for analysis. The base flow cell initially had a source-to-object distance of 3.5 cm. However, this distance was insufficient to meet the $z_1 \gg z_2$ requirement and maintain the approximate (or near) unit object magnification assumption. Additionally, providing sufficiently uniform illumination to the sensor's active area was a concern. Four spacer cylinders, each measuring 3.5 cm in depth, were milled out of polyoxymethylene acetal

and added to the in-line configuration to increase the overall source-to-object distance (z_1). The spacer cylinder stack can be seen in Figure 4.2(a). The FSA is presented in Figures 4.2(a) and 4.2(b) with important areas and parts of the analyser being indicated.

Light source and mounting

A partially coherent laser diode was chosen for this study, primarily due to its low cost and sufficient coherence properties. The DIH application requires suitable spatial coherence from the light source to produce strong diffraction rings, making the laser diode an appropriate choice. Although recent studies have explored the use of LEDs with spectral (bandpass) and spatial (convex mirror) filters or pinholes to improve their coherence properties in holography settings [16, 48], a single mode 405 nm laser diode (SLD3233VF, Sony) was preferred for this study. The wavelength was selected based on the low-cost commercial availability of the diode and the extinction ratio curve of the polarisation image sensor, in which wavelengths with a higher extinction ratio are preferred.

The laser diode was coupled to a 3.0 μm diameter single-mode fibre (460HP, Nufern) and terminated on an FC-PC connector. This connector was mated to a single-mode coupler (RS-536-7718, RS) and fixed to a laser-cut stainless-steel end cap (304 SS, 2 mm) to create a stable mount on the flow cell. The fibre coupling of the laser diode provides partial spatial and temporal coherence as only the fundamental mode propagates through the fibre.

Excitation driver

A stable light source was necessary for analysing moving fluids, and a custom driver circuit (Figure 4.3) was designed and fabricated to support the laser diode. The circuit, termed the light pulse engine (LPE), serves two primary tasks. Firstly, it provides a controlled constant injection current to ensure the laser diode outputs a stable, partially coherent beam. Secondly, it offers an adjustable pulse width triggered by the rising edge of the image sensor's exposure window activation. The adjustable pulse width mitigates motion blur that may ordinarily be encountered when running the laser diode in continuous wave operation, and where the minimum supported exposure time of the image sensor is not short enough. The components for this prototype were selected based on meeting the low-cost aspect of this project and being in stock from the PCB fabricator (JLCPCB, China); the circuit diagram can be found in Appendix B.

The excitation driver unit employs a dual monostable multivibrator (74HCT4538, Nexperia) in a non-retriggerable configuration. The driver circuit is designed to output a




Figure 4.3: LPE circuit illustration.

pulse width that is shorter than the minimum exposure time supported by the image sensor. The resulting effect is a shorter equivalent or effective exposure time since the flow cell chamber lacks ambient light and relies solely on the laser diode source for sample illumination. This is especially useful for high flow rates where motion blur becomes more likely as particles may traverse multiple pixels during a single exposure. Reducing the effective exposure time offers particles less opportunity to cross multiple pixels during exposure, resulting in sharper images.

Generating an extremely short pulse reduces the time that the laser diode is active and decreases the number of incident photons per exposure to illuminate particles. The laser diode was overdriven during its short duty cycle by implementing a stable and adjustable current source circuit, which aims to counter the effect of the limited photon issue. For additional excitation applications, we included three driver channels. Each driver channel has its own voltage-controlled current source (transconductance amplifier) that comprises a high-speed operational amplifier (AD8042, Analog Devices) and NPN-bipolar junction transistor (SS8050, Onsemi). We use a CMOS switch (ADG715, Analog Devices) to activate individual excitation channels, controlled through an I²C connection to an ESP-32 development board. Custom code was written for the development board to control the unit. The LPE was fitted with additional capacitors to source the necessary current on demand to overdrive each light source. The resulting effect is the generation of enough photons to maintain the required intensity for effective imaging whilst reducing motion blur simultaneously.

Verification of the pulse width shortening circuit was performed with an oscilloscope (DSO-X 1102A, Keysight Technologies) and the signal trace is shown in Figure 4.4. The bottom signal (1, green) is an input pulse and the top signal (2, orange) is the




Figure 4.4: Oscilloscope trace from LPE analysis.

output pulse from the circuit. The rising edge of the signal (such as the camera exposure start) triggers the pulse start and the output signal ends after a set duration (here approximately $7\ \mu\text{s}$). The exposure window signal continues for a longer time (here approximately $100\ \mu\text{s}$) and the output pulse will not start until it detects another rising edge from the exposure trigger waveform. The three excitation channels may be used to support additional use cases such as multi-band excitation. This may be useful if the unit is used in an RGB-type excitation and image sensor environment for further spectral feature analysis, or multi-modal applications such as multi-band fluorescence and polarisation combinatory imaging. The unit is also suitable for pairing with less coherent imaging setups using LEDs. The LPE included a PWM-controlled driver for a peristaltic pump that used a complementary Darlington configuration (MJD122, Onsemi) also controlled by the ESP-32.

3D printing, polarisation equipment, and FSA costing

Custom fittings were designed for the project to achieve accurate alignment and mounting of the image sensor and light source. These components were 3D-printed to encourage the rapid prototyping nature of this project. A 3D-printed polariser mount (Figure 4.5(b)) was created to link the output of the single-mode coupler to a 15 mm glass zero-order left-hand circular polariser (CP42HE-#19-927, Edmund Optics). Since the image sensor was run in a lensless configuration, the original C-mount lens mounting and sensor enclosure were removed; a custom camera mount was created to cradle the image sensor. See Figure 4.5 (c) and 4.5 (d) for the designs of the mountings and




Figure 4.5: Custom designed components.

cradles. The 3D prints were produced in polylactic acid filament on a low-cost printer (MINI+, Prusa).

The complete hardware cost for the project is \$1 290 and is designed to be paired with a laptop as a proof of concept. This is significantly cheaper than commercial variants of non-polarisation sensitive imaging flow cytometers (upwards of \$40 000) [16] and supports the low-cost aspect of the project. This cost includes all of the fittings, the image sensor, and the illumination source. The polarisation image sensor (at \$800) is the largest cost driver and the total system cost could be significantly reduced in bulk quantities. A full cost breakdown can be found in Appendix A.

4.2.2 Software workflow development

This section covers the software designed to interact with the hardware detailed in the preceding section. Since this chapter focuses on the development of the holographic polarisation-sensitive FSA, it places attention on the tasks of image capture, preprocessing, base feature extraction, and segmentation of objects present within a fluid sample. Figure 4.6 shows the output channel development of the FSA from the initial captured hologram. The base feature extraction includes the DIH methods to numerically propagate the imaged hologram as well as the birefringence measurements from the polarisation-sensitive hardware. A single image capture may encompass multiple objects, which is a consideration particularly relevant for the case involving diverse marine and synthetic polymer specimens within a designated size range. Further feature extraction and classification of these objects is considered in Chapter 5, while the processing pipeline up to the point of completing object segmentation is considered here.

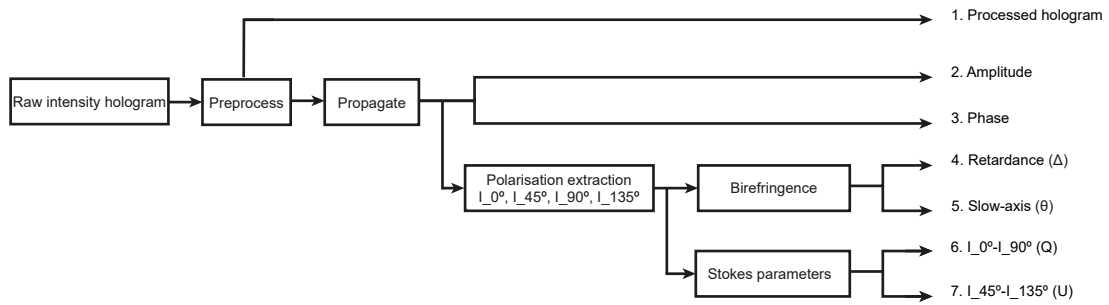


Figure 4.6: Processing workflow of the FSA. The input frame is processed and expanded to seven output channels. The output channels can either be used for qualitative observation or further downstream quantitative processes.

Preprocessing

The intensity image captured by the sensor may exhibit speckle artefacts due to the partial coherence of the source, as well as aberrations caused by static particles present in the flow channel or along the optical path. Since the device was designed to be deployable for multiple capture events with limited maintenance between each run, we wish to reject as many unwanted artefacts and as much noise as possible. A preprocessing algorithm was therefore implemented to improve the SNR. Corrections to an image under analysis may be required due to artefacts and noise. In microscopy settings, background correction or normalisation is often performed by capturing two images: one image is the static background frame (captured before introducing an object of interest into the scene), and the other image is the pertinent frame with the object of interest included. For a background correction procedure, the background frame is subtracted from the pertinent frame and a standard uniform mean intensity value may be added back. On the other hand, for a background normalisation procedure the pertinent frame may be divided by the background frame and then normalised by adding back a standard uniform mean intensity.

The described standard correction techniques are unsuitable for an automated fluid stream imaging workflow, as a recent object-free image frame is not always available. Instead, a simple algorithm was implemented for correction. Firstly, it averages the 90 previous frames to produce a mean background image. Next, this frame is subtracted from the current frame to produce the background corrected frame. Finally, the mean intensity of the background frame is uniformly added back to the newly corrected image to produce the resulting output. Naturally, any moving objects present in the previous 90 frames are averaged out by those frames, leading to minimal presence in the background frame. Any static objects in the 90 preceding frames will, however, create a strong component in the background image of a similar intensity, providing a reference for

what needs to be corrected in the frame and assisting in removing such interference.




Figure 4.7: Image correction routine. Scale bar in all figures indicates $50\ \mu\text{m}$.

An example of the image correction technique can be seen in Figure 4.7, where PE spheres were passed through the detector. The diffraction rings and particles shared by the raw and background holograms are visibly removed from the post-processed output image. Looking closely at the images, the speckle patterns present in the raw and background holograms appear similar. This suggests that the unwanted pattern may be due to the coherent source reflecting off debris and surface imperfections. These artefacts are maintained throughout the average frame capture and therefore the correction process yields a considerably cleaner image in the post-processing result.

Since we consider objects of interest to always be in motion when passing through the flow cell, this correction routine is also beneficial when an object detector or classifier is used in a follow-on stage from the FSA. When the corrected image is used for object detection purposes, it may be less prone to false detections from artefacts that are near-stationary but may vibrate. Additionally, the corrected image provides a more stable input to a classifier, especially in the presence of variations in illumination intensity. The variation in an object's scattering, and by extension transparency properties, results in objects that may have regions of intensity greater or less than the background mean intensity. This can make segmentation more of a challenge.

4.2.3 DIH

Following the approach outlined by Goodman [11], the angular spectrum method is incorporated into our workflow to enable the numerical reconstruction of plane wave propagation within the wavefront. This methodology aligns effectively with the on-chip lensless configuration employed in this research, and since the object-to-sensor distance

(z_2) is substantially less than the source-to-object distance (z_1), a near-unit object magnification is achieved. Consequently, the image sensor's pixel pitch determines the analyser's spatial resolution. The propagation function is supplied with the preprocessed hologram as an input and the following parameters: propagation distance, wavelength of the excitation source, dimensions of the input frame, and pixel pitch of the image sensor. The function uses two separate Fourier transforms to obtain the resulting complex propagation output.

Determining the object propagation distance for numerical focusing requires knowing the object-to-sensor distance beforehand, using an autofocus criterion or setting a coarse propagation metric. Depending on user requirements, two propagation options were implemented:

- **Individual object propagation:** an object can be selectively brought into focus by cropping the object and subsequently applying an autofocus metric calculation that measures the sparsity of the gradient [49]. This approach demands greater computational resources as each object must undergo an iterative individual autofocus routine, which involves multiple propagation steps but ensures object sharpness.
- **Global propagation distance:** since we do not know the exact object-to-sensor distance but do know the working range of the flow channel, a uniform global propagation distance can be set. This is typically equal to the mid-depth distance, which in this instance is 2.25 mm. This distance is uniformly applied to all objects present and is considered to be a coarse propagation of sorts. One drawback of this method is that objects at either the minimum (2.00 mm) or maximum (2.50 mm) depths within the channel may exhibit less sharp details, and this would be especially applicable for fine features closer to the sensor resolution limit. Nevertheless, this approach offers greater computational efficiency.

4.2.4 Polarisation sensitive birefringence imaging

Developing the FSA for birefringence sensing requires the integration of a subsystem block to process the propagated intensity image and create two output channels to accommodate the birefringent retardance (Δ) and slow axis orientation (θ). We compute two further channels to include Stokes parameters 1 and 2 (Q and U). They are defined as [50]:

$$\begin{aligned} Q &= I_0 - I_{90}, \\ U &= I_{45} - I_{135}. \end{aligned} \tag{4.1}$$

The propagated intensity output image from the DIH stage is supplied as an input to this subsystem. To access subsets of the 2D array and perform grouping based on their polarisation orientation, a small script was created to divide the array into four separate 2D arrays each corresponding to a distinct polarisation orientation (0° , 45° , 90° , and 135°), aligned with the image sensor's configuration. Consequently, a minor spatial displacement exists among the four pixels comprising these orientations. This shift is however considered insignificant relative to the minimum size of the sample objects of interest and registration is not required.

Given that the excitation source produces left-hand circularly polarised light and a linear PFA image sensor is employed, a subsequent script applied the Jones calculus method for the computation of relative birefringent retardance and slow axis orientations across the pixels detected by the sensor [14]. As this method requires four unique linear polarisation orientations to output one pixel's birefringence retardance and slow axis calculation, the resultant output image undergoes bilinear interpolation to revert to its initial resolution.

The outcomes of this process yield two distinct periodic 2D maps measured in radians: a retardance map and a slow axis orientation map. The slow axis map undergoes further refinement to eliminate any artefacts originating from instances where the computed retardance value falls below a threshold of 0.1 radians. This adjustment prevents the computation of a slow axis value in scenarios where no particles are present — and is therefore not physically meaningful — or where the birefringent signal is too weak to produce a reliable outcome. The two additional Stokes channels (Q and U) are computed and also added to the channel stack. We upsample all channels by a factor of two using a bilinear kernel, to restore the original image pixel dimensions.

4.2.5 Segmentation

Segmentation is helpful as many particles of interest may be present in a moving fluid that is imaged by the FSA. The goal is to extract these object patches for use in various downstream processes in the pipeline, including visual inspection or classification. In this study, the segmentation of marine phytoplankton and synthetic polymer samples from background is of interest. Segmenting both of these sample types concurrently in a single frame as foreground is necessary, even if they exhibit different intensity characteristics. We implemented a supervised (neural network) approach. This choice was due to the operational requirements of the FSA, which excluded optimisation-based or unsupervised-trainable methods as they were either computationally demanding or less automated and therefore unsuited to the application. Rule-based thresholding methods were also investigated but were deemed unsuitable.

Procedure

We used a circulatory configuration to move particles in solution between a holding vessel and the FSA. The plastics considered in this study (PE, PET, PP, and PS) are considered non-polar hydrophobic synthetic polymers [51], so mixing these samples with water to form a solution required consideration as each type has a different density and they tend to bind to surfaces. The phytoplankton samples were somewhat easier to work with but still required a well-mixed solution. The following protocol was adopted when dealing with any sample type: the samples were placed in a beaker (Glassco, India) with laboratory water and a small amount of dish soap (0.1% v/v) and ethanol (0.1% v/v) was added to act as a surfactant. The soap and ethanol help reduce hydrophobic interactions that cause the particles to bind together or cling to the FSA and associated equipment. The beaker was placed on top of a magnetic stirrer (MSH10, Labcon South Africa) to agitate the solution and mix the particles in the solution despite the differing sample densities. Using a peristaltic pump (NKP-DCS-10B, Kamoer), the sample in solution was circulated between the beaker and the FSA where they were imaged for analysis.

Trainable supervised segmentation

An encoder-decoder approach was chosen because it is common for multiple objects to be present in a scene and a fine-grained pixel-wise inference supports more detailed shape measurement analysis for logging purposes.

We tailored the base structure to better suit the fluid analysis segmentation requirements. Although pre-trained models provide a transfer learning approach that leverages larger datasets and offers benefits when used appropriately, the unique characteristics of the FSA's output base features make them well-suited to a "from scratch" network implementation. Transfer learning has been well applied in the general microscopy arena [52] and could be useful in future studies using the FSA with limited channel inputs. The seven channels delivered from the FSA are: preprocessed hologram, mid-depth propagated intensity, mid-depth propagated phase (unrecovered), birefringence retardance, birefringence slow axis orientation, Stokes parameter 1, and Stokes parameter 2.

Dataset for trainable supervised segmentation

A small dataset was created to provide representative examples of foreground and background regions for segmentation network training. This enables the network to learn from the examples and apply its model to new, unseen data. Synthetic polymers and marine phytoplankton were used as foreground objects for the segmenter. The marine

phytoplankton samples included *Isochrysis galbana* and *Chaetoceros calcitrans*, while the synthetic polymer samples comprised PE, PET, PS, and PP. These foreground objects exhibit regions of both high and low intensity and it is necessary for the network to generalise effectively and segment various types of unseen foreground objects.

A primary objective of this segmentation process is to train the network to distinguish background regions based on texture. This enables the network to accurately segment unseen foreground objects, even when they differ from objects used to train the network, particularly in terms of shape. Many imaged objects have unique shape and intensity properties, making it necessary for the network to recognise the background texture concept to identify new foreground objects effectively.

The FSA was configured in the standard circulatory configuration and a collection of 20 image sequences was acquired, including phytoplankton and synthetic polymers. These capture routines were conducted separately for different classes; for instance, marine phytoplankton were captured in a distinct session from the synthetic polymers and each synthetic polymer class was captured individually. Each image sequence comprised 200 frames and 50 images were selected for the segmentation dataset, including ten images of each object class. These images had dimensions of 2448×2048 pixels and contained multiple objects within each frame. Standard preprocessing procedures were applied to each sample image and run through the holography and birefringence routines to obtain the resulting seven-channel stack per frame ($2448 \times 2048 \times 7$ pixels).

The Matlab image labeller tool was used to manually annotate the regions of interest in the images to generate the ground-truth labelled masks in a pixel-wise manner. The propagated intensity channel was selected from the seven-channel stack to locate the object area of interest for the labelling process. A value of zero was assigned for the background pixels and one to five was assigned to specific foreground object classes: PE, 1; marine, 2; PS, 3; PET, 4; and PP, 5. This extended label indexing was done in preparation for the upcoming classification chapter. However, in this instance, once the training masks had been created, the pixels with class values greater than one were all set to one for this binary segmentation training task.

Augmentation

Typical augmentation for 2D image inputs is applied globally and includes geometric transformations, colour processing and intensity transformations [53]. Given the monochrome environment, our dataset does not require colour processing but intensity adjustments may be made. In terms of geometric transformations, reflection, rotation,

and translation may be performed, as would normally be done with non-holography intensity problems. However, shear transformations are unsuitable as a shear translation at the hologram plane does not translate to a shear transformation at the object plane and therefore does not create a physically meaningful relation.

The following augmentations were applied to the training dataset with a 0.5 probability: vertical and horizontal flip, random rotation 90° , random brightness contrast, and custom function — random propagation (up to $250\ \mu\text{m}$ from channel centre).

The custom propagation augmentation technique was developed to address potential optical misalignments within the system. This function randomly propagates the intensity channel up to a distance of $250\ \mu\text{m}$ from either side of the mid-depth position within the flow cell channel. The rationale behind this approach was as follows: while the training label mask and raw hologram remain constant, minor discrepancies in the optical system may exist. For instance, if the image sensor is slightly further away from the flow channel glass, a mid-depth propagation routine might yield suboptimal performance because it propagates to an area that is slightly out of focus, potentially leading to less accurate segmentation. Instances where the boundary regions of the propagated image exhibit poorly defined edges may challenge a segmenter and lead to greater uncertainty. However, with the proposed approach, the network can learn the relationship between the label mask and the raw hologram, enhancing the reliability of segmentation even in the presence of minor discrepancies. These augmentations were applied exclusively to the training dataset following the dataset split and no augmentation was performed on the validation and test sets.

Network implementation

In this project, the background region often contributes a significant portion of the captured image due to the large field of view provided by the lensless configuration. This is especially true when smaller objects are present. For example, one such object might have an area of $50\ \mu\text{m} \times 50\ \mu\text{m}$ ($2.5 \times 10^{-9}\ \text{m}^2$) compared to the active area of the image sensor, which is $5.96 \times 10^{-5}\ \text{m}^2$. This means that this example object is approximately $1/23840$ of the sensor's active area. However, in the developed dataset, multiple objects are present in a single shot.

The 2448×2048 pixel input image was divided into smaller patches of 128×128 pixels to improve computational efficiency. These patches served as inputs to the model and assisted in reducing the computational burden associated with a larger single frame. Furthermore, patches containing only background pixels, devoid of any object pixels,

were excluded from the dataset during model training. This exclusion provided balancing assistance to the dataset, as in general the frames were dominated by background.

Our proposed implementation modified certain hyperparameters and functions within the ResUnet [34] model. As mentioned, the input stack from the FSA was divided into smaller patches. The kernel size of 3×3 was maintained, as was the layer depth of seven. In particular, we set the number of input channels to five: defocused raw hologram, propagated intensity hologram, phase (wrapped), birefringence retardance, and birefringence slow axis orientation. The remaining two Stokes parameter channels from the FSA were not used. In our adjusted model, we set the input filter number to 128 and used a five-channel input to accommodate the FSA stack with a patch size of 128×128 . In other words, we concatenate the five channels from the FSA into a single matrix and feed the network with patches sized $128 \times 128 \times 5$ pixels (which includes each chosen FSA channel).

Instead of a binary cross-entropy (BCE) loss function as used in U-Net, a combination of Dice [54] and focal loss [55] was employed in this model. Both the Dice and focal loss functions effectively enable networks to learn from imbalanced datasets, especially in scenarios with significant background components as is expected for this application. An equal weighting was given to these functions. Dice loss (DL) and focal loss (FL) are presented as follows:

$$\text{DL} = 1 - \frac{2 \cdot p \cdot y}{p^2 + y^2} \quad (4.2)$$

and

$$\text{FL} = -\alpha \cdot (1 - p)^\gamma \cdot \log(p), \quad (4.3)$$

where p represents the predicted probability of the positive class (class 1), y represents the true binary label (1 for positive, 0 for negative), α is the scaling parameter, and γ is the focusing parameter.

A sigmoid output activation function was used since this network was trained as a binary segmenter for background (BG) and foreground (FG). The sigmoid function maps the input to a probability estimate output (0,1) and is represented as follows:

$$\sigma(x) = \frac{1}{1 + \exp(-x)}. \quad (4.4)$$

Before training, the dataset was normalised and split into training (60%), validation (20%), and test (20%) sets. The augmentation routine was applied only to the training set. We set the class weightings to the inverses of their respective proportions to assist with the background and foreground class imbalance.

The Adam optimiser [56] was used, along with a batch size of 16. The resulting model was trained for 34 epochs — a value determined by an early stopping condition. We trained an additional three network implementations to provide a comparative measure of performance with the model already presented: a vanilla U-Net structure with BCE loss function and 5-channel input, a vanilla U-Net structure with Dice and focal loss function and 5-channel input, and a ResUnet structure with Dice and focal loss function and 3-channel input. The base U-Net implementation results in a similar number of learnable parameters but a different number of filters and depths due to the lack of residual blocks; this implementation used 64 filters in the first layer to make up the training parameter difference. An early stopping condition determined each additional model's epoch number.

Performance metrics

Standard machine learning metrics such as accuracy and loss were not used to evaluate the pixel-wise model's segmentation performance. Instead, the metrics IOU and F1-measure were used. This choice was due to the dominant background region, which can often skew metrics like accuracy and report misleading figures. By using IOU and F1-measure, which consider true positives, false positives and false negatives, the evaluation provides a more accurate representation of the model's segmentation performance. IOU is represented as

$$\text{IOU} = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}}, \quad (4.5)$$

where TP are the true positive predictions, FP are the false positive predictions and FN are the false negative predictions.

F1-measure is represented as:

$$\text{F1-measure} = \frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}}. \quad (4.6)$$

4.3 Results

In this section, we first present the results from the segmentation network training. Next, we present our experimental demonstrations where the FSA is analysed in its ability to perform both holographic and birefringence resolving tasks. A segmentation study follows this. Finally, we cover the calibration of the unit.

Table 4.1: Binary segmentation results.

Model	Mean IOU	Mean F1	IOU BG	IOU FG	F1 BG	F1 FG
ResUnet (proposed)	94.56	97.14	99.38	89.74	99.69	94.59
ResUnet 3-channel	94.43	97.07	99.37	89.49	99.68	94.45
U-Net DL+FL	90.94	95.07	99.01	82.88	99.50	90.64
U-Net BCE	88.74	93.75	98.52	78.97	99.25	88.25

4.3.1 Trainable segmentation results

The results from the trainable segmentation are displayed in Table 4.1. The vanilla U-Net with BCE loss exhibited less competitive performance than the Dice and focal loss implementations (all other models), and may suggest some form of bias due to the dominant background regions that encouraged a preference to predict the background class. A performance improvement was noted for all metrics when the same-sized parameter U-Net model was used with Dice and focal loss. When assessing our proposed implementation, further improvements in all metrics were noted. As a form of ablation study, the 3-channel ResUnet network that excluded the two birefringence channels saw a minor performance decrease due to their exclusion. This suggests that the network can sufficiently learn to segment without using the additional birefringence channels, which may be excluded for improved inference speed at the sacrifice of minor performance.

4.3.2 Experimental demonstration

DIH

We demonstrate the lateral resolving abilities of the FSA by imaging a USAF 1951 resolution test target (HCM01-63R, Dongguan Hongcheng Optical Products). The test target was placed between the image sensor and the illumination source (with the FSA channel glass removed due to the test target size being larger than the flow channel). The acquired hologram was propagated to its focus distance (1.28 mm) using an auto-focus criterion [49]. The propagated intensity image (Figure 4.8) is displayed with no background correction routine as this was a static target and we wish to demonstrate the pure resolving abilities of the system. We cropped the orange bounded area for closer inspection. Figure 4.9(a) shows the cropped region before propagation and Figure 4.9(b) shows the same region after propagation. The propagated image shows more resolvable detail but quantification is necessary.

Upon further inspection of Figure 4.9(b) we note some elements in groups 5 and 6 are visually resolvable. These smallest groups appear resolvable by inspection and are selected for further analysis. The high-intensity transparent regions on the target substrate are

considered background regions, whilst the low-intensity regions are the chrome elements printed on the test target. For an analysis to be performed, an intensity profile is created along the perpendicular direction of the elements under inspection. We apply the following criterion to consider an element being resolved: the central peaks (high-intensity regions) must rise over 27% above the dip intensity (chrome printed regions) [11].

Looking at the intensity profile of selected group 5 elements (G5-E4, G5-E5, and G5-E6) in Figure 4.10, we observe two distinct peaks that separate each middle element from their outer neighbour on each side; the rise is above the required threshold and therefore meets the resolving criterion. If we now look at group 6 elements (G6-E2 and G6-E3), we observe that G6-E2 is resolved, however G6-E3 does not satisfy the criterion. G6-E2 corresponds to approximately $7.0\ \mu\text{m}$ and G6-E3 corresponds to approximately $6.2\ \mu\text{m}$. We take the mean value $6.6\ \mu\text{m} \pm 0.8\ \mu\text{m}$ as the resolving ability of the FSA.

The results agree with the image sensor's expected resolution ability and our optical setup. The influence of twin image noise and no preprocessing routine does impact the observed curves. We would expect more regular transitions in the absence of additional interference.




Figure 4.8: Resolution test target — full frame (propagated).

To further demonstrate the DIH propagation imaging, a sample of synthetic polymers was mixed in solution and passed through the FSA. The intensity image acquired from the image sensor exists as a single 2448×2048 2D array, though the figures presented




Figure 4.9: Resolution test target analysis — cropped orange region from Figure 4.8.




Figure 4.10: Intensity profile resolution test target analysis. (a) Successfully resolved elements. (b) Point of failure.

throughout this example are cropped parts of this larger original frame. This array undergoes the standard preprocessing and mid-depth propagation routine (Figure 4.11), resulting in a processed and propagated intensity array that is suitable for either visual analysis or for a subsequent stage in the workflow.




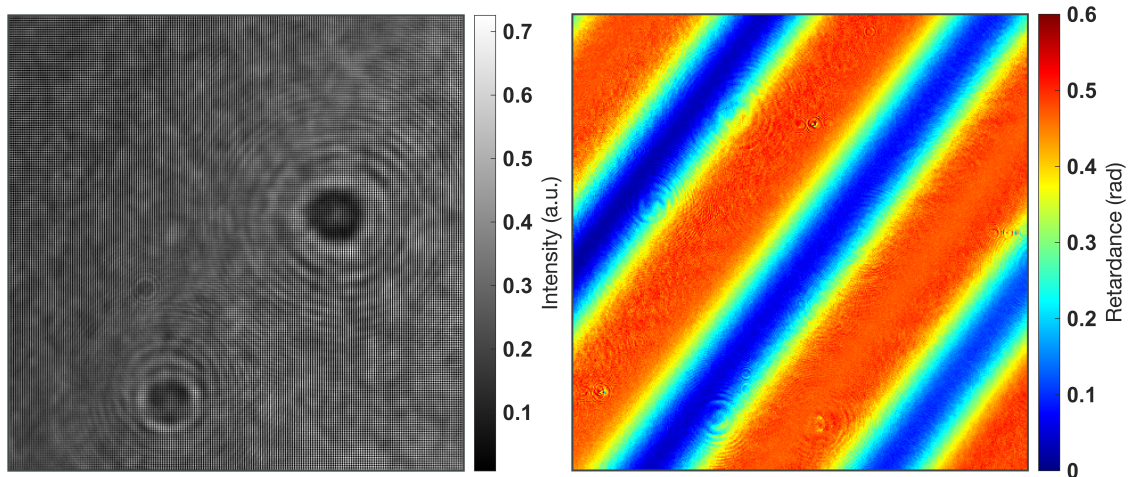
Figure 4.11: Preprocessed propagated mid-depth intensity of microplastic samples (low-intensity regions).

Birefringence imaging

To verify the quantification of birefringence, we imaged a birefringent depolariser (DPP-25B, Thorlabs). This liquid crystal polymer depolariser is sandwiched between two glass plates and has a pattern of varying retardation that has been imprinted in the liquid crystal polymer. When illuminated with a linearly polarised monochromatic source, its polarisation state will become spatially varied. The periodic retardation is expected to vary between 380 nm and 430 nm. We used a 650 nm laser diode (SLD1133VS, Sony) for this experiment due to the anti-reflective coating on the depolariser and the reference data available from the manufacturer. Figure 4.12(a) shows a cropped portion of the captured raw intensity frame with no general periodic variation in intensity. Figure 4.12(b) shows the computed birefringence retardance of the same raw intensity frame. The computed retardance shows the spatially periodic variation in retardation as one moves diagonally across the frame and this agrees with the manufacturer's specifications.

The birefringence computation provides a quantitative measure of optical anisotropy that can be used for visual inspection when using the FSA in a visual capacity. Alternatively, these measurements can serve as inputs for machine learning algorithms geared towards particle classification.

To further demonstrate the polarisation-sensitive birefringence imaging, we use the same



(a) Raw intensity.

(b) Birefringence retardance.

Figure 4.12: Depolariser analysis at 650 nm.

example shown in the propagation section of some synthetic polymers mixed in solution and passed through the detector. The two-channel output (retardance and slow axis) can be combined for visualisation to generate a unified map, as depicted in Figure 4.13. This map was created by employing a hue and saturation colour space transformation, subsequently translated back into the RGB colour space. Birefringence retardance corresponds to saturation, while slow axis orientation refers to hue.




Figure 4.13: Pseudo-coloured retardance and slow axis orientation (cropped).

The advantage of using the birefringence data with the hue and saturation visualisation technique is its ability to display both metrics in one image. Furthermore, the birefringence computation suppresses both optically isotropic particles and background intensity with a notable reduction in twin-image interference. This aids users in isolating the quantifiable birefringence activity. However, as is evident from the output image, specific pixels within the borders of the sample specimens exhibit minimal to no birefringence which is indicated by near-zero intensity pixels. We select a smaller particle





Figure 4.14: Birefringence processing of PET. (a) Preprocessed hologram. (b) Propagated intensity. (c) Pseudo-coloured birefringence retardance and slow axis orientation. (d) Pseudo-coloured birefringence retardance and slow axis overlay/fused with propagated intensity. Scale bar in all figures indicates $50\ \mu\text{m}$. Colour scale legend available in Figure 4.13.

from the original hologram to demonstrate the processing workflow and illustrate the lower regions of anisotropy (Figure 4.14). The workflow moves from left to right as the preprocessed hologram (Figure 4.14(a)) is propagated (Figure 4.14(b)) and processed for birefringent activity (Figure 4.14(c)), which is then fused with the propagated intensity image (Figure 4.14(d)). Some edge pixels and a few of the centre pixels suggest low anisotropy. Although these pixels may have legitimately low activity regions, observing clusters of a few pixels is recommended to determine overall presence of activity. Similarly, if these pixels were to be used for feature vector generation, a single pixel may not provide a sufficiently dependable input. Grouping a few pixels within a neighbourhood or using filters for further feature extraction would therefore be advisable.

Segmentation

We provide cropped frames (2448×2304) as a reference example to demonstrate the wide FOV for a single-shot multi-particle application (Figure 4.15). The scale bars present illustrate the ability of the FSA to capture many particles of the size range of interest in a single frame for high throughput processing. In Figure 4.15(a) the raw hologram is shown before the preprocessing routine has taken place to correct and remove any static particles or artefacts. In Figure 4.15(b) the output from the network

is displayed, bearing in mind that any particles below the size threshold of interest can be removed with morphological processing. Although the input image was partitioned into patch sizes of 128×128 and fed to the segmentation network, the output patches were reassembled to make up the larger frame with pixel-wise segmentation.




(a) Raw hologram frame.

(b) Network segmentation output.

Figure 4.15: Comparison of the raw hologram at beginning of pipeline and the resulting segmentation output (cropped).

If we select a challenging segmentation example and pass it through the network-based approach (Figure 4.16), we obtain a relatively acceptable result. It is observed that the high and low intensity diffraction rings that surround each particle have little effect on the segmentation performance which suggests that particle measurements should be reliable for downstream logging purposes. Additionally, if these object pixels were to be sent for a foreground classification task further down in the pipeline it may result in fewer background pixels being sent to the classifier, thus reducing the likelihood of performing inference on a non-category pixel.



(a) Preprocessed intensity frame.

(b) Segmented frame.

Figure 4.16: Image segmentation when applying workflow for network-based segmentation (cropped).




Figure 4.17: Propagated intensity and predicted output patches (128×128) pixels. Scale bar in all figures indicates $50 \mu\text{m}$.

A selection of ten examples is shown in Figure 4.17 to demonstrate the learnable segmentation performance qualitatively. Our proposed implementation with the 5-channel input was used to generate these predictions. The model has not encountered these examples during training and includes synthetic polymer and phytoplankton samples. These examples were taken straight from the output of the segmenter with no further post-processing being performed. We chose these examples for display as they are generally more challenging segmentation instances of particles in a fluid stream. We found that the smaller particles of only a few pixels were often more difficult to segment, whereas the larger objects presented less of an issue. Many of the examples used in this section illustrate the performance near the lower particle size cut-off ($50\ \mu\text{m}$) to demonstrate the segmentation abilities and weaknesses. Looking at all of the examples, the spatial awareness of the network is evident. The ability of the segmenter to associate the background texture as well as the surrounding diffraction rings (low and high intensity regions) as non-particle elements is useful in our application.

This study considers microplastics from $50\ \mu\text{m}$ to $1000\ \mu\text{m}$ in length to be of interest, and the diameter of the PE spheres range from $63\ \mu\text{m}$ to $75\ \mu\text{m}$ (to provide a reference of the working range). Figure 4.17(a) presents an example of a sphere that has passed through the detector, with its corresponding segmentation output (Figure 4.17(b)) being reasonably performed. In Figure 4.17(c), a microalgae sample is shown. This example was selected to illustrate the presence of both light and dark intensity areas within the region of interest, which makes it more challenging to accurately segment compared to the previous microsphere example. The segmentation result (Figure 4.5(d)) shows effective segmentation of both regions. In Figure 4.17(k), a plastic sample is shown and the resulting output (Figure 4.17(l)) displays suboptimal segmentation. While the model has successfully found the general region of interest, the boundary areas have not been delineated accurately and could be improved upon.

The remaining examples were deemed reasonably segmented and are provided for further analysis. Notably, the fine small particles near the resolution limit in Figures 4.17(g), (i), and (s) were set as background pixels by the segmenter as it was trained to consider them as background. These fines were below the detection size threshold of this project and were therefore considered to be artefacts that are equivalent to background noise. The training data used to create the model had equivalently sized objects that fell under the general background label. These examples demonstrate the segmentation capability to allow these fines to pass through the segmentation process without contributing to false positive detection results.

Motion blur

Observing or resolving finer details of the particles in motion is necessary for further visual or machine-learning activities in the pipeline. Since motion blur can be an issue in a fluid stream setting, we performed a stepped flow rate test with PE spheres and three flow rates (20, 40, and 60) mL/minute to determine the impact of motion blur with and without using the LPE.

The image sensor’s frame rate (five frames per second) and exposure time (80 μ s) were set and held constant for all sequence captures. When testing the system without the LPE, we supplied a constant current (50 mA) to the 405 nm fibre-coupled laser diode with the resulting beam being the continuous wave output. When testing the system with the LPE, the LPE’s pulse width was set to 7 μ s and the constant current source (1 250 mA) was used to overdrive the same fibre-coupled laser diode but this time in a low-duty cycle pulsed configuration.



(a) Continuous wave operation.

(b) LPE operation.

Figure 4.18: Motion blur test (40 mL/minute).

Image sequences were captured for both configurations at the three flow rates and all sequences underwent the standard preprocessing routine. We show an example (Figure 4.18) of the propagated intensity channel for each configuration at one flow rate (40 mL/minute). It is noticed that the LPE (Figure 4.18(b)) frame shows a smoother background texture whereas the continuous wave operation (Figure 4.18(a)) shows “wave-like” ripples. We inspected the other two flow rate sequences and they presented similar patterning, but when a static sample was tested the ripples were found to be absent. It is possible that flow-induced vibrations were introduced when delivering fluid with the

peristaltic pump. The flow-induced vibrations are not apparent when using the LPE configuration as it effectively minimises the integration time of the image sensor.




Figure 4.19: Propagated intensity images of PE spheres using continuous excitation. (a) 20 mL/minute. (b) 40 mL/minute. (c) 60 mL/minute. Scale bar in all figures indicates 50 μm .




Figure 4.20: Propagated intensity images of PE spheres using the LPE. (a) 20 mL/minute. (b) 40 mL/minute. (c) 60 mL/minute. Scale bar in all figures indicates 50 μm .

Upon closer inspection, three examples of each configuration are provided to view the effect of the stepped change in flow rate. The particle flow direction in all of these examples was from the top to the bottom of the page. The continuous wave operation (Figure 4.19) shows that as the flow rate increases, the particle becomes more elongated and loses its circular appearance even though a global shutter image sensor was used in this study. This is an issue for the accuracy of shape measurement performance and the smearing of the object also means that any further derived characteristics such as birefringent texture might be less useful for a classifier later in the pipeline. The LPE operation (Figure 4.20) shows the particles retaining their circular projected area at all flow rates, and even at the lowest flow rate there is noticeably less smearing compared to its non-assisted counterpart. Furthermore, the LPE-assisted frames' non-object regions

show finer detail with more defined diffraction rings and less noticeable flow-induced vibrations.

4.3.3 Calibration

It is important that the detector yields meaningful quantitative information about the particles that have passed through the flow channel. Although the maximum distance between the object and the image sensor is significantly less than the distance between the light source and the object, the theoretical unit magnification approximation needs to be validated appropriately. Spheres are widely employed as calibration standards for imaging and general flow cytometry applications prior to their use. These calibration spheres are available in various materials and size fractions. They are typically characterised by either a diameter range, with a specified percentage of particles falling within that range (e.g., 50 μm to 60 μm with 80% of particles within this range and unknown distribution), or a mean diameter with a standard deviation (e.g., 50 $\mu\text{m} \pm 5 \mu\text{m}$). Such specifications allow for accurate calibration of the FSA.

Although we consider near unit magnification of the sample due to the short object-to-sensor distance, the thickness of the glass viewing discs in the flow cell suggests it would be prudent to implement a calibration protocol to ensure measurement accuracy. Additionally, the system has two sources that may contribute to measurement inaccuracy: firstly if the image sensor is not precisely aligned with the glass plate due to manufacturing tolerances of the image sensor housing and other optical alignment issues, and secondly if the developed segmentation network introduces inaccuracies in the segmentation process. To deal with this, we imaged calibration spheres of known size and applied a correction factor in the processing pipeline to account for a difference between the expected value and the observed value. This process should be performed following the FSA assembly or any adjustments to the segmentation network.

The calibration process simply involved passing calibration spheres in solution through the FSA and analysing several of them. PE spheres were added to a solution of laboratory water, with 0.1% v/v dishwashing soap and 0.1% v/v ethanol acting as a surfactant. The solution was placed in a beaker and pumped into the flow cell with the peristaltic pump. Image data was captured by the FSA, which subsequently performed the pre-processing and finally the segmentation. The segmented data objects, including the 2D projected area, were logged and recorded. The manufacturer specifies the diameter range of 63 μm to 75 μm with more than 90% of particles within this range. Since these particles are not supplied with a coefficient of variation, we could not consider the size distribution to be normal and rather made a conservative assumption of a uniform

Table 4.2: PE sphere measurement run.

Measurement	Count	Mean	Median	Standard deviation	90th percentile
1	233	290.14	288.04	64.39	370.57
2	263	287.37	286.52	61.33	358.56

Mean, median, standard deviation, and 90th percentile are given in units of pixels.

distribution with the expected size being the midpoint.

The 63 μm to 75 μm diameter particles cover 261 pixels to 371 pixels of 2D projected area at unit magnification and acceptable segmentation. We coarsely filtered the range of analysed particles from 200 pixels to 432 pixels so that any fines or outliers did not skew the process. From the logged data, 200 spheres were used for calibration and the mean area was calculated. We use the ratio of the actual and expected area as a calibration factor. This factor was applied at the end of the processing pipeline as an adjustment to the projected 2D area that is logged. In other words, this calibration does not impact the region that is segmented by the workflow and is purely for measurement logging accuracy.

Calibration assessment

We first tested the FSA pipeline on a new measurement of 100 frames of PE spheres, done in duplicate. Particles such as fines with an area of less than 125 pixels were excluded from the analysis. The particle area data was used to create a normalised histogram, which approximates a probability density function (PDF) based on the experimentally measured data. The following approach is used to generate the measured PDF:

$$f_i = \frac{c_i}{B \cdot w_i}, \quad (4.7)$$

where f_i is the i^{th} histogram bin, w_i is the bin width, c_i is the number of elements in the bin, and B is the number of bins in the histogram. The normalised histograms for both acquisitions are presented in Figure 4.21 and provide an approximate particle size distribution for the PE spheres. Both histograms demonstrate peaks in the 296 to 308 bin with the tendency to taper off on either side, and are more in line with a weak normal distribution. Further analysis was performed and is presented in Table 4.2. The two measurement runs yield closely related mean and median values. The observed mean for both measurements is within the specified size range from 261 pixels to 371 pixels, although it is below the midpoint value. We also calculated the 90th percentile value for the observed data since it specified that 90% of spheres should fall into the size range. Both measurements yield 90th percentile values that fall into the expected size

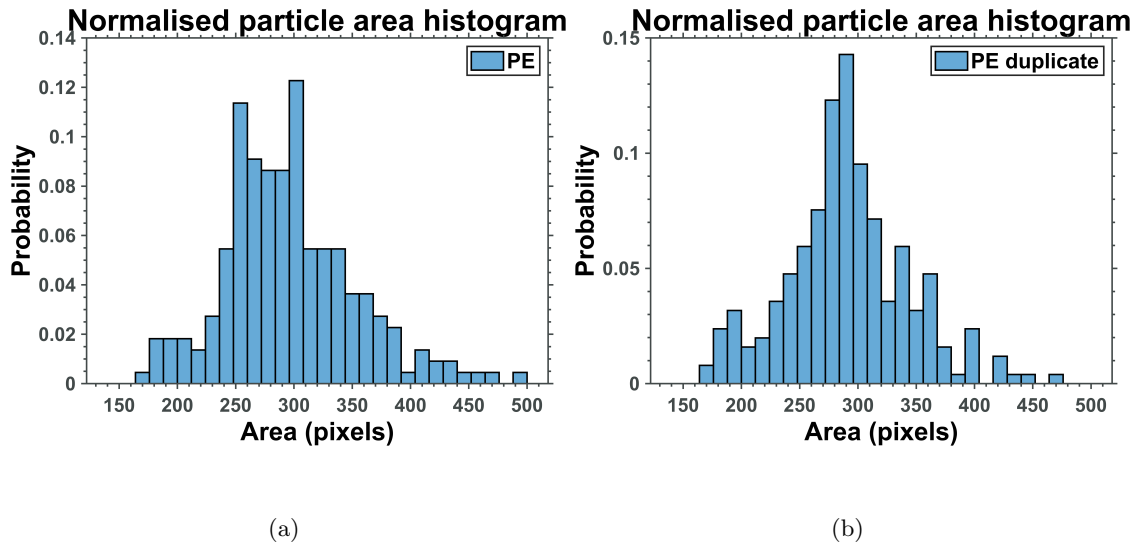


Figure 4.21: Normalised histogram (PDF approximation) PE. (a) PE initial run. (b) PE duplicate run.

Table 4.3: PMMA sphere measurement run.

Measurement	Count	Mean	Median	Standard deviation	85th percentile
1	853	227.54	221.92	48.14	256.12
2	786	232.01	225.72	53.44	258.85

Mean, median, standard deviation, and 85th percentile are given in units of pixels.

range, thus suggesting agreement with the specified data and calibration procedure.

We also performed an additional experiment with PMMA spheres (Cospheric, USA). This sphere class is considered unseen because it is a different type of plastic and the samples themselves have not been used to train the network. This test was used to further ascertain if the network can use the background holographic texture to discern background and foreground classes as well as provide a measurement of the calibration performance. Similar to the PE spheres, these PMMA scientific calibration spheres were also provided with a known but different diameter range ($53\ \mu\text{m}$ to $63\ \mu\text{m}$) and a percentage of spheres (85%) within this range. The same acquisition process was followed and the normalised histograms for both capture runs are displayed in Figure 4.22.

The histograms for both runs illustrate different distributions to the PE runs performed earlier. It is noticed that the distributions are packed considerably tighter and are less spread out than the PE runs. When further assessing the statistical data in Table 4.3 it is observed that the standard deviation is smaller than the PE run, supporting the difference in histogram appearance. The mean for both PMMA measurements falls




Figure 4.22: Normalised histogram (PDF approximation) PMMA. (a) PMMA initial run. (b) PMMA duplicate run.

close to the mid-band (223 pixels) of the expected area range (185 pixels to 261 pixels) and the calculation for the 85th percentile is within the expected range, which agrees with the specifications. The ability of the network to segment this unseen type of plastic supports the addition of further object types for background and foreground segmentation, enabling the network's use as a form of object detector.

4.4 Discussion

Phase recovery is not covered in this study due to the weakly and strongly scattering variations in synthetic polymers [10], which make phase unwrapping algorithms unpredictable. Since the particles flow in solution, static particle multi-height iterative phase recovery methods are also impractical, but deep learning [57] or multi-wavelength excitation phase recovery methods could make for interesting future studies. Deep learning approaches have already made their mark in the autofocus and phase recovery domains [57]. However, they typically require training on meticulously curated datasets that are often tailored to very specific object types that can sometimes lead to hallucinations when presented with a more exotic example of the same object type. This has not been considered in this work but again offers an opportunity for further investigation. It is thought that the LPE would be useful for other use cases where the pixel pitch sizes are smaller than the example presented in this work, or if the flow channel depth is decreased. These constraints would lead to a shorter effective exposure time requirement for the same given flow rate.

4.5 Conclusions

In this chapter, we developed the FSA, a low-cost digital holographic imaging tool that includes polarisation-sensitive sensing and can extract 2D birefringence features. The device is intended for fluid stream studies and includes an LPE to reduce motion blur. A trainable segmentation routine was also developed. The analyser was assessed regarding holographic propagation, birefringence sensing, segmentation, motion blur, and calibration. We feel that the device provides an opportunity for further development in imaging-based sensing in many environments (not solely our microplastic use case).

The next chapter uses the FSA to sample fluid streams of both microplastic and phytoplankton samples to develop a dataset. This dataset is used to train classifiers for microplastic analysis in a recreated marine setting. Both learned feature extraction and filter bank feature extraction are explored to select the most suitable candidates.

Chapter 5

Computer vision

Classification of microplastics presents a challenge in that they are found in differing shapes and sizes, leading to unique samples. This complexity is further compounded when attempting classification within a fluid stream, as these microplastic samples often coexist with debris and microorganisms, such as phytoplankton, particularly in marine applications. To address these challenges, we leverage both the hardware and software components of the FSA, which were introduced in the previous chapter, to facilitate the development of multiple microplastic classifiers. A small dataset is created for the purpose of classifier training including plastics and phytoplankton. Using polarisation-sensitive imaging generates essential birefringent features required for the classifier models. Our approach encompasses a comprehensive examination of learned feature extraction and filter bank feature extraction techniques, which serve as inputs for training various classifiers. Among these classifiers, the customised ResUnet implementation delivered the highest classification performance and achieved the lowest FPR in the context of binary marine and microplastic classification. A sequential routine of filter banks, supervised dimensionality reduction, and a random forest model contributed to the best-performing multi-class polymer classifier. This work effectively supports the approach that birefringent features are useful in enhancing the classification of microplastics in imaging-based fluid stream analysis.

5.1 Introduction

This chapter develops methods for classifying microplastic samples amongst marine matter. The classification process uses the FSA presented in Chapter 4 as a base. The techniques developed aim for polarisation-sensitive machine learning driven classification of microplastic particles. However, these methods are not intended to replace traditional

microplastic classification techniques. Instead, they complement them by providing quick and cost-effective classification for resource-limited settings or when used as a prescreening tool. Should more detailed analysis be required, samples can be further processed and analysed through existing laboratory-based methods.

The FSA developed in Chapter 4 has a lensless holographic imaging system that can sample flowing fluid and spatially resolve intensity and birefringence features. These features serve as inputs to the classification pipeline developed in this chapter. Two pathways are explored to obtain a suitable classifier: one through manual feature development and the other through learned feature extraction. Manual feature development relies on a more hand-crafted approach to generate and extract features for classifier training, which often includes the use of filter banks. In contrast, learned feature extraction may rely primarily on neural networks to automatically extract features. Both pathways have their strengths and weaknesses, which are thoroughly examined in this chapter.

The suggested workflow may vary depending on the specific use case, such as the ability to classify microplastic particles as a distinct class from marine microorganisms or classifying different types of synthetic polymers separately from marine microorganisms. Furthermore, the adopted approach may be influenced by the available computing resources, processing duration restrictions, classification accuracy, and limits of detection. This chapter's classifier development task treats the segmentation step as a preliminary preprocessing stage in the classification workflow that has already taken place. Consequently, the metrics assessed are limited to the classification performance and not the segmentation performance, which was dealt with separately in Chapter 4.

The chapter commences by covering the creation of a small dataset comprising marine phytoplankton and plastic samples. Subsequently, implementing a learned feature classifier in the form of an adapted encoder-decoder U-Net-style structure is assessed, demonstrating promising results for the binary (microplastic and marine) classification problem.

Filter bank manual feature extraction is then explored, implementing both Gabor and MR filter sets to create a high-dimensional pixel-wise feature vector. Feature selection and supervised dimensionality mapping were applied to streamline the expanded feature set, and this process resulted in a more suitable feature vector length for training a machine learning classifier. Multiple machine learning models were tested for the multi-class polymer classification problem, with the random forest classifier showing promise.

For the purposes of the binary (microplastic and marine) classifier it was felt that the

learned feature approach was most suited to this application, whilst for the multi-class synthetic polymer classifier the filter bank feature extraction route with a random forest was preferred. The results suggest that the selected classifiers are capable of performing classification on the binary (microplastic and marine) problems as well as multi-class polymer problems, and that the birefringence-resolving capabilities of the FSA provide a suitable input to this workflow.

5.2 Materials and methods

5.2.1 Dataset

The preceding chapter created a small base dataset comprising synthetic polymers and marine phytoplankton samples for the purpose of training a supervised binary segmenter. Since microplastic and phytoplankton datasets are not readily available, a larger dataset is created to train classifiers effectively. Given this study's exploratory nature, the developed dataset is limited to a proof of concept approach but encompasses sufficient examples representing each class to facilitate training. It is acknowledged that this dataset is not extensive and there is a recognised need to expand its content over time to accommodate a broader and more diverse sample set.

Samples

Marine phytoplankton cultures of *Chaetoceros calcitrans* and *Isochrysis galbana* were obtained from the DFFE Seapoint Marine Research Aquarium. These cultures served as a comparison class of non-plastic particles or objects commonly found in marine environments that fall within or below the size range of interest. This comparison class, or negative class, enables evaluation of the classifier's performance by offering alternative objects that aid in identifying instances of misclassification.

Since only two phytoplankton cultures were used in this study, it is highlighted that the approach creates a simplified marine-like environment. The aim is not to replicate the complex and diverse phytoplankton composition typically found in regular seawater samples; rather, the intention is to establish a controlled setting that allows for effective validation of outputs in response to known inputs without the risk of potential microplastic contamination. This approach provides a reliable foundation for assessing the performance of the developed workflow as a proof of concept.

Micronised synthetic polymer powder samples (Shijiazhuang Tuya Technology Ltd., China) included PET, PP, PE, and PS. Additionally, scientific calibration microspheres

of both PE and PMMA (Cospheric, USA) were obtained for dataset development and experiments. The micronised synthetic polymer powder samples exhibited varying shapes and sizes, ranging from under 1 μm to over 2 mm. When visually inspecting these samples with the FSA, it was observed that the particle size distribution varied considerably among different plastic types. While achieving a collection of synthetic polymers with the same size distribution but diverse shapes is challenging, a step toward enhancing comparability amongst the samples' size ranges was necessary. Additionally, some synthetic polymer samples contained many very fine or small plastic particles that were near or below the image sensor's resolution limit. These tiny particles created an extremely turbid environment and overwhelmed the image sensor without contributing objects within the desired size range for imaging. As a result, a dry screening mechanical sieving process was employed to refine and enrich these plastic samples, aligning them more closely with the size fractions of interest.

Stainless steel laboratory test sieves (Endecotts Ltd., UK) with aperture sizes of 25 μm , 53 μm , and 105 μm were used for the dry screening process along with a vibratory shaker (Analysette 3 Pro, Fritsch GmbH., Germany). The screening process ran three five-minute program intervals at an amplitude of 0.1 mm that were performed separately for each synthetic polymer type. After the screening process, samples measuring 25 μm and below were set aside, while the particles ranging from 25 μm upwards were grouped into equal parts by mass. While this process significantly improved the availability of particles of interest in the processed sample batches, it was noted that fine particles were still present, albeit in reduced quantities. This occurrence can be attributed to the electrostatic nature of the micronised plastic particles, causing them to adhere to various parts of the metal sieves, particularly the walls of the sieve bands. Consequently, these remaining fines would be retained in the refined sample batches, potentially contributing to background noise levels. This suggests that these smaller particles may contribute to the particle count distribution being positively skewed as the lower size range may contain many particles. Please refer to Appendix A for a comprehensive description of the full dry screening procedure.

Dataset creation

Once the samples were within the size range of interest, they could be used to generate a dataset for classifier training. The following points regarding the dataset creation process are deemed important for reproducibility:

- **Flowing samples:** Images were captured while samples were in motion; static sample imaging was not employed to generate the dataset.

- **Solution preparation:** Sample solutions for each object class were prepared by firstly creating a saline solution (3.5% salinity and laboratory water) to which 0.1% v/v ethanol and 0.1% v/v dish soap were added, acting as a surfactant. Finally, object class samples were added.
- **Marine and non-microplastics:** Creating individual classes for microalgae, debris, and other particulates would be an intractable task due to the substantial variations among marine organisms and their states of health. Therefore, we have categorised these entities under a single, all-encompassing “marine” class, which includes all particles except synthetic polymers.
- **Sample classes:** The dataset included the following classes: PET, PE, PP, PS, and marine. Individual solutions were prepared for each class type so that they could be imaged through the FSA without contamination from other classes. The phytoplankton present in the “marine” class consisted of mixed *Chaetoceros calcitans* and *Isochrysis galbana* samples, which were imaged together as they belonged to the same object class in the view of the classifier.
- **Solutions and concentrations:** Each synthetic polymer class had a concentration of 0.32 g/L, while the marine class was a 1:1000 dilution from each stock. Since the objects imaged in the acquisition runs are used to generate samples for the dataset only, the intention here is not to specify the exact suitable concentrations and dilutions for dataset creation but rather to state what was chosen for our experiments. The stock dilution factor is purely based on what was observed when studying the raw stock and thus will vary depending on the cell concentration of the culture. We did not want extremely sparse frames, which would require capturing many frames to include sufficient objects. Similarly, we aimed to avoid having too many objects in the frame so that the image sensor was not overwhelmed.
- **Capture sessions:** Each class was run through the detector in separate image capture sessions to preserve the known ground truth. Multiple objects of the same sample class were present in a single frame and a collection of these frames formed a session. Each image session comprised 200 frames with standard image dimensions of 2448×2048 pixels.
- **FSA settings:** The image sensor settings remained consistent with the previous chapter, namely, five frames per second with the LPE generating $7 \mu\text{s}$ pulses to reduce motion blur. The flow rate was set to 20 mL/minute and the standard preprocessing procedure was applied to each capture session.

- **Maintenance and handling:** The flow cell, piping, and glassware were dismantled and cleaned in soapy water between each sample class run. Cotton laboratory overalls were worn during the procedures to avoid any (synthetic) fibre contamination from clothing.

Images were selected from the capture runs to represent each object class and formed a total set of 90 images that provided the foundation for the small dataset generated in this study. Each of these 90 images had seven primary feature channels as generated and delivered by the FSA, namely: preprocessed hologram, mid-depth propagated intensity, mid-depth propagated phase, birefringence retardance, birefringence slow axis orientation, Stokes parameter 1, and Stokes parameter 2.

For generating the ground-truth labelled masks, Matlab Image Labeller was used to manually label the regions of interest in the images, assigning values of 0 - 5 for background, PE, marine, PS, PET, and PP, respectively, for the object or foreground class in a pixel-wise manner. The propagated intensity channel was used as the reference image for the labelling process.

Given that the dataset follows a pixel-wise labelling approach, it exhibits certain levels of class imbalance. Nonetheless, subsequent phases of classifier development offer suitable opportunities to address this imbalance by selecting a subset of the dataset or generating synthetic data through augmentation techniques to supplement lacking classes.

5.2.2 Learned feature classifier

We continue with the pixel-wise approach for this classification application, in a similar fashion to what was implemented for the segmentation block of the FSA in Chapter 4. This encoder-decoder approach was chosen because it is common for multiple objects to be present in a scene. Therefore, a fine-grained pixel-wise inference supports more detailed shape measurement analysis for logging purposes. The network was trained from scratch as pre-trained backbones are not readily available for the non-traditional seven-channel FSA input used in this work, even though transfer learning has been well applied in the general microscopy arena [52].

In approaching the binary classification task of categorising marine and synthetic polymer particles, we investigated many variants of the encoder-decoder scheme. The residual U-Net (ResUnet) [34] is a fully convolutional network that combines features of U-Net [20] and residual learning. It uses U-Net as a foundation and incorporates pre-activated residual blocks or skip connections into the model instead of the convolutions found in the original U-Net.

Our proposed implementation modifies certain hyperparameters and functions within the ResUnet model. The input stack from the FSA was divided into smaller patches rather than a single large frame as it is more computationally efficient. The kernel size of 3×3 was maintained, as was the seven-layer network depth. In our adjusted model, we set the input filter number to 128 and used a seven-channel input to accommodate the FSA stack with a patch size of 128×128 . In other words, we concatenated the seven channels from the FSA into a single matrix and fed the network with patches sized $128 \times 128 \times 7$ pixels (which included each FSA channel). The network layout is illustrated in Figure 5.1 and the network structure and feature extraction are presented in Table 5.1.




Figure 5.1: Layout of adapted ResUnet classifier.

The original ResUnet used mean-squared error as a loss function, however, a combination of Dice [54] and focal loss [55] were employed in this model in softmax form. Similar to the segmentation problem noted in Chapter 4, these loss functions are useful in scenarios with class imbalance in datasets and are given a 50/50 weighting. Categorical Dice loss (CDL) and categorical focal loss (CFL) are defined as follows:

$$\text{CFL} = - \sum_{i=1}^N (\alpha_i \cdot (1 - p_i)^\gamma \cdot \log(p_i)), \quad (5.1)$$

$$\text{CDL} = 1 - \frac{2 \sum_{i=1}^N p_i \cdot g_i}{\sum_{i=1}^N p_i^2 + \sum_{i=1}^N g_i^2}, \quad (5.2)$$

where N is the number of classes, γ is the focusing parameter, α_i is the balancing factor for each class i , p_i is the predicted probability over class i , and g_i is the ground truth label.

Table 5.1: Adapted ResUnet structure.

	Level	Convolutional layer	Feature size	Stride
Input			$128 \times 128 \times 7$	
Encoder	1	1	$128 \times 128 \times 128$	1
		2	$128 \times 128 \times 128$	1
	2	3	$64 \times 64 \times 256$	2
		4	$64 \times 64 \times 256$	1
	3	5	$32 \times 32 \times 512$	2
		6	$32 \times 32 \times 512$	1
Bridge	4	7	$16 \times 16 \times 1024$	2
		8	$16 \times 16 \times 1024$	1
Decoder	5	9	$32 \times 32 \times 512$	1
		10	$32 \times 32 \times 512$	1
	6	11	$64 \times 64 \times 256$	1
		12	$64 \times 64 \times 256$	1
	7	13	$128 \times 128 \times 128$	1
		14	$128 \times 128 \times 128$	1
	Output		15	$128 \times 128 \times 3$

A softmax output was used since this network is trained as a three-class output: background, synthetic polymer, and marine. The softmax output assigns a probability estimate to each class and is represented as follows:

$$s(v_i) = \frac{\exp(v_i)}{\sum_{j=1}^N \exp(v_j)}, \quad (5.3)$$

where (v_i) represents the vector of raw outputs from the neural network, i denotes the i^{th} entry of the softmax output vector, and N is the number of classes.

The dataset exhibited a distribution of 58% synthetic polymer and 42% marine pixels. To account for this imbalance, class weightings were applied to the model and the class weights were calculated as the inverses of their respective proportions.

Before training, the dataset was normalised and split into training (60%), validation (20%), and test (20%) sets. The augmentation routine developed earlier in this work was applied only to the training set. The Adam optimiser [56] was used, along with a batch size of 16. The resulting model was trained for 45 epochs — a value determined by an early stopping condition.

For testing and model evaluation purposes, it is assumed that the segmentation process has already occurred and that only the two class classification results are of interest. This decision was made to facilitate accurate comparisons with filter bank feature extraction methods later in this work as the inclusion of background segmentation in this section would introduce inconsistencies in these comparisons.

A traditional U-Net was used as a base implementation for performance comparison against the ResUnet model. The U-Net implementation results in a similar number of learnable parameters, but naturally a different number of filters and depths. This implementation uses 64 filters in the first layer to make up the training parameter difference. The learned feature approach was only implemented for the binary marine and polymer classification problem as initial exploration into a learned feature multi-class polymer classifier yielded poor results.

5.2.3 Filter bank classifier

Birefringent texture

If we consider the multi-class synthetic polymer classification problem and examine Figure 5.2, which displays the mid-depth intensity of four different synthetic polymer samples processed with the FSA, it is observed that the images reveal differences in geometric appearance, particularly in terms of random shapes rather than interior object intensity. While these shapes might prove useful for specific filters like edge detectors, the edges themselves may vary randomly from one sample to the next. It is possible that the edges themselves exist due to a sample-unique degradation process and thus are possibly not unique to a specific polymer type. Given that this study prioritises the use of interior features over global geometric measurements, our attention now turns to polarisation-assisted feature development. Examining Figure 5.3, we find




Figure 5.2: Intensity mid-depth propagation of synthetic polymer samples, where objects are the low-intensity regions in the frame. (a) PS. (b) PE. (c) PET. (d) PP. The small particle fines mentioned in the sample refining process are visible in (b). Scale bar in all figures indicates 50 μm .

that the birefringence retardance channel exhibits clear optical anisotropy. It is well-established that different synthetic polymers display varying degrees of birefringence owing to their manufacturing processes. Moreover, micron-resolution imaging combined with polarisation-sensitive analysis offers insights into the sub-micron molecular arrangement of the lattice, even when it falls below the resolving capabilities of the imaging

system [12]. Given the presence of birefringent activity in these samples, as well as




Figure 5.3: Retardance channel of the same four samples from Figure 5.2. (a) PS. (b) PE. (c) PET. (d) PP.

noting that the manufacturing process of polymers leads to crystalline or amorphous structures [12], it is conceivable that the birefringence channels may contain periodic or stochastic texture features that could prove valuable for feature extraction when subjected to appropriate image filters.

The FSA delivers baseline feature channels, and we wish to generate further derived features by using them as an input to generate multiple outputs. The input signal is convolved with different filters and is decomposed into sub-bands. This filter bank is also referred to as an analysis bank [35, 58]. The resulting decomposed sub-bands reveal low-level features that a machine-learning classification algorithm may find useful.

In the case of the synthetic polymer samples, an expectation is that periodic features or stochastic structures may be present in the birefringence data. In that respect, filters for texture-based analysis are appropriate for this use case. However, texture is a global characteristic so relying on filters to gain insight into low-level neighbourhood features that may describe part of the texture is vital and should form a significant contribution to the feature vector. Since the segmentation process has already been performed by the FSA further up in the processing pipeline, these filters are primarily intended for internal feature extraction and less emphasis is placed on external edge detection. Moreover, this is not considered a texture segmentation problem but rather a texture representation and classification problem.

Out of the synthetic polymers considered in this study, PE, PET, and PP may exist in crystalline or amorphous structures, or a combination of the two (semi-crystalline). In the PE domain, HDPE may have a higher degree of crystallinity when compared to LDPE. PS is generally primarily amorphous but can be modified to exhibit some degree of crystallinity [12]. As far as features and textures are concerned, a combination of filters is required to extract a periodic and stochastic view of texture, especially on the

birefringence channel considering the order and random textural arrangement of pixels.

It is established that the plastic particles present in this study will exhibit non-uniform shapes or sizes. The marine phytoplankton samples are expected to be more uniform in shape but with varying dimensions, especially for *Chaetoceros calcitrans* in terms of chain length. Naturally, plastics as well as non-plastic debris (considered marine class) will vary in both shape and size. During a particle's movement through the FSA, its orientation will vary as it traverses the flow channel and as such the image sensor has no control over which orientation it will be presented with during this process.

Filter bank implementation

The filter bank approach for feature extraction and the design of the pixel-wise classification workflow was not developed for speed but rather as a proof of concept for an acceptable standard of synthetic polymer classification. The choice of filters for feature extraction, therefore, was not based on computational demands but rather on which are the most appropriate for the texture features.

The input frame from the FSA is a multi-channel image due to the holography and polarisation principles applied to generate digitally propagated birefringent features; these channels are referred to as primary or base feature channels. Three of these channels (mid-depth propagated intensity, birefringent retardance, and birefringent slow axis orientation) are used for further feature extraction with the analysis bank. We therefore refer to these derived features as extended features. This creates a higher-dimensional feature space compared to applying the same filter set to a single-channel monochrome image. Training a classifier in such a higher-dimensional space may pose challenges and performing classification in a lower-dimensional feature space is preferred, thereby necessitating methods for dimensionality reduction.

We implement two filter banks for feature extraction: a Gabor and an MR set.

The Gabor filter bank broadly covers a range of wavelengths (λ), orientations (θ), and standard deviations (σ). The aim is to provide general coverage and not be overly tuned. The selected parameters for the filter bank are as follows: $\lambda = [10\ 20\ 30\ 40]$, $\theta = [0\ 45\ 90\ 135]$ degrees, and $B = [0.5\ 1\ 1.5\ 2]$, where $B = \frac{1}{2\pi\sigma}$.

The MR filter bank [40] is employed with its complete set of 38 filters in this study. The authors of the filter bank implemented a 49×49 kernel size for their study, as they used 200×200 pixel texture patches. However, in our study some particles were considerably smaller than this kernel size. As a result, the kernel size was adjusted to 13×13 through experimentation and the following parameters were used for these

Table 5.2: Feature list

Feature number	Input	Filter bank
1-7	Fluid stream analyser base features	None
8-71	Propagated intensity	Gabor
72-135	Birefringent retardance	Gabor
136-199	Birefringent slow axis	Gabor
200-237	Propagated intensity	RFS
238-275	Birefringent retardance	RFS
276-313	Birefringent slow axis	RFS

Gabor bank: $\lambda = [10\ 20\ 30\ 40]$, $\theta = [0\ 45\ 90\ 135]$ degrees, and $B = [0.5\ 1\ 1.5\ 2]$.

RFS bank: kernel size [13], scales [1 2 4], $\theta = [0\ 60\ 120\ 180\ 240\ 300]$ degrees, and rotationally invariant filters [2].

anisotropic filters: scales [1 2 4] and orientations [0 60 120 180 240 300] degrees.

We commence the construction of a feature space using the aforementioned filters. The resulting analysis bank is outlined in Table 5.2. This feature bank aims to offer comprehensive coverage for texture analysis, with subsequent application of feature selection and dimensionality reduction techniques to refine it. Consequently, not all the features in this list are used in the final workflow.

Feature selection

Filter-based feature selection was performed through a Pearson correlation analysis; Pearson correlation measures the normalised linear relationship or covariance between two variables.

The Pearson correlation coefficient (c) is calculated as follows [59]:

$$c = \frac{\sum_i (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_i (X_i - \bar{X})^2 \sum_i (Y_i - \bar{Y})^2}}, \quad (5.4)$$

where c is the Pearson correlation coefficient, X_i is the value of the variable X for the i th data point, \bar{X} is the mean of variable X, Y_i is the variable Y for the i th data point, and \bar{Y} is the mean of variable Y.

The dataset was used in conjunction with the analysis bank to create an extended feature dataset, which was then examined for correlations. A correlation matrix was generated from the feature dataset which, combined with a predefined correlation threshold (set at 0.8), identified highly correlated features. The correlation score served as a criterion to select and discard the most redundant features. The feature selection served as a first-pass attempt in the greater dimensionality reduction campaign and is considered a broad approach to removing only very highly linearly correlated features.

The feature selection process identified 171 features from the original list of 313 features for retention. All seven primary feature channels provided by the FSA were selected. Additionally, out of the 102 filters in the analysis bank, each applied to the mid-depth intensity, retardance, and slow axis channels. The following numbers of derived features were retained for each primary or parent feature:

- Mid-depth propagated intensity derived features: 39
- Birefringent retardance derived features: 69
- Birefringent slow axis orientation derived features: 56.

This suggests greater linear redundancy when applying the analysis bank to the mid-depth intensity channel on our dataset. However, when applying the bank to birefringent primary features one can extract richer, less linearly correlated features.

Dimensionality reduction

Initial exploration using PCA successfully accounted for 95% of the variance by reducing the feature set to five eigenfeatures. However, the performance was substandard when this mapping was further tested in classifier training and real-world experiments. The issue appeared to be the linear nature of the PCA mapping, which failed to establish well-defined relationships among the features, particularly in cases where non-linear relationships were present. Instead, we employed uniform manifold approximation and projection (UMAP) for graph-based supervised non-linear dimensionality reduction [60].

We employed a non-parametric training approach for UMAP. Although the algorithm is not entirely deterministic, the authors suggest that it is near deterministic. We used a supervised approach to assist in mapping the higher-dimensional feature space to a lower-dimensional feature space that can be clustered based on class. The resulting model subsequently performs future mappings on new, unseen, and unlabelled data. This dimensionality reduction model occupies the position between the output of the selected features (primary and analysis bank) and the classification model input within the workflow.

We first performed mapping with reference to the multi-class polymer classification problem. The input data used to train the model was balanced to prevent potential bias towards a specific class. This approach ensured that the training routine had adequate examples from each class, enabling it to establish meaningful distance relationships for cluster mapping. The input features were normalised using a standard scoring approach. For the training set, a random sample of 20 000 feature vectors was selected from each

class of synthetic polymer, resulting in a total training set of 80 000 labelled feature vectors each with a length of 171.

A supervised learning approach was used for this model, where class labels were provided. The criteria for hyperparameter selection was to maximise inter-class cluster distances and minimise computational overhead, with the generated feature vector mapping being suitable as an input into traditional machine learning classifier models like random forests, k-nearest neighbours (KNN)s, or SVMs. The chosen hyperparameters were a five-dimensional vector mapping, ten nearest neighbours, and a minimum distance of 0.7. These parameters were selected to balance feature mapping quality and support efficiency as computing additional nearest neighbours is computationally demanding. Additionally, balancing global and local distance relationships is important for sensitivity to outliers and reducing overfitting risks.

After the model was constructed, an additional set of 160 000 unseen samples was drawn, in which each synthetic polymer class was equally represented. This sample set (devoid of their class labels) was processed by the constructed model. The generated feature mapping transformed the original 171-dimensional feature space into a reduced five-dimensional space. These newly selected samples were intentionally chosen to be independent of the earlier sample data used to construct the model; this approach was undertaken to reduce the risk of classifier overfitting and support robustness in the forthcoming stage of the workflow.

Although the mapping was performed in five-dimensional space, Figure 5.4 displays three of these dimensions, revealing clusters for the four different polymers that were generated with unseen data. The class labels were added after the mapping process for display purposes and were not used to generate this mapping. It is evident that some polymers (and their clusters) have tighter distributions than others, raising questions about the influence of noise. All five dimensions were interrogated but without loss of generality this figure is used for discussion.

Suppose a cluster is extremely small and tight; it may lead to model overfitting. Conversely, if the cluster is too large and close to other clusters, it might be susceptible to influence from non-class particles, resulting in misclassification. Hyperparameters such as the minimum distance span, the distance metric, and the number of nearest neighbours considered are important for tuning purposes. However, as can be seen from the plot, different polymers have different spreads and therefore hyperparameters were selected to accommodate the differing distributions fairly.

This also prompts a discussion about post-classification processing and how the workflow




Figure 5.4: Dimensionality reduction four-class (PE, PP, PS, and PET), showing three of five dimensions.

should handle classification output. Using an ensemble-type approach to take the class mode of the model’s output for a given object patch could lead to more reliable inference, especially when encountering noise.

We also generated a model for the two-class (polymer and marine) problem with the filter bank methodology. This served as a contrast to the ResUnet neural network implementation discussed in the previous section. A similar approach to the multi-class polymer-only dimensionality reduction was followed, but this time mapping to a three-dimensional space was suitable.

For the supervised training set, we randomly selected 10 000 feature vectors from each synthetic polymer class, combining them into a single polymer class of 40 000 vectors. Additionally, 40 000 marine feature vectors were randomly selected, resulting in a total training set of 80 000 labelled feature vectors to construct the model. The same hyperparameters were used as for the multi-class polymer approach.

After constructing the model, we drew an additional set of 100 000 unseen samples, including 50 000 from the single polymer class and 50 000 from the marine class. Figure 5.5 shows a two-dimensional orientation of the 3D mapping of these 100 000 unseen and unlabelled samples. The visualisation suggests that these two classes each have prominent clusters, but they share significant overlap in between.

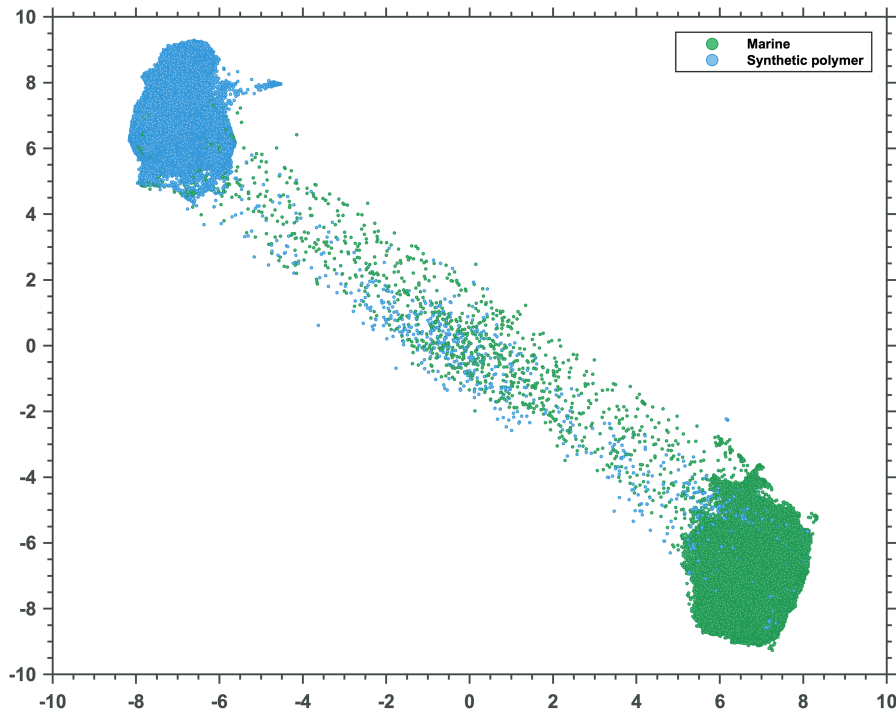


Figure 5.5: Dimensionality reduction for two classes (marine and polymer).

Following the mapping process, the class labels were reintroduced to these lower dimensional datasets and were available for use with supervised learning approaches.

Training classifiers

The resulting lower dimensional feature set provides a more suitable input to a classical machine learning model. We batch-ran the training of multiple supervised learning classifiers and optimised hyperparameters to determine the most suitable model. Given the nature of the dataset, it was felt that a random forest would be a useful classifier in this instance. Random forests are particularly well-suited to this dataset because they use random splits in trees; this characteristic makes them less susceptible to overfitting and enhances their ability to generalise, effectively handling variations in the data. The reduced risk of overfitting can be particularly advantageous for the marine class, given the limited dataset size. In the two-class classification scheme, marine particles encompass varying phytoplankton and debris, which poses a challenge for classifiers with rigid boundaries. In this context, a random forest classifier may be less sensitive to outliers and noise than models such as an SVM.

Examining the clusters again in Figure 5.4, it is evident that PS and PET clusters are small and tight, which may make them prone to overfitting and sensitive to noise in new data. Additionally, boundary regions where the inter-class cluster distance is the shortest also pose a risk. Compared to the clusters generated in the marine-polymer task,

multi-class polymer classification is a more difficult problem. Therefore, a consistent preprocessing protocol and normalisation routine are important for new data processing, and using an ensemble-type scheme to take the mode of classification output might prove beneficial for the workflow.

The classifiers were modelled separately for two problems: a microplastic and marine classifier to compare with its learned feature neural network counterpart and a multi-polymer classifier. Two dimensionality-reduced datasets were available from the previous section: a three-dimensional 100 000 feature vector set for the binary marine-polymer problem and a five-dimensional 160 000 feature vector set for the multi-class polymer problem. Both sets were normalised using the standard-score approach.

Binary (microplastic-marine) classifier

The dataset of 100 000 three-dimensional feature vectors was used to train binary classifiers to predict marine and synthetic polymer classes. The pool of feature vectors was equally divided between the two classes and subsets were created by partitioning the dataset into training and test sets using an [80, 20] distribution with five-fold cross-validation. Several traditional machine learning models were trained; each model used a 30-iteration Bayesian optimisation routine to minimise the objective function in aid of appropriate hyperparameter selection.

Four class (multi-polymer) classifier

The dataset of 160 000 five-dimensional feature vectors was used to train multi-class classifiers to predict PE, PET, PP, and PS samples. The pool of feature vectors was equally divided amongst the four classes and subsets were again created by partitioning the dataset into training and test sets using an [80, 20] distribution with five-fold cross-validation. Several traditional machine learning models were trained and the same iterative optimisation routine was used for hyperparameter selection as was used previously in the binary classifier section.

Hyperparameters and tuning

The tuning of the classifier model's hyperparameters was limited to the standard settings during training and no additional tuning of the model's decision boundary was performed after training. This decision aimed to avoid overfitting and maintain the model's ability to generalise effectively to real-world data, and moreover to ensure the general reproducibility of our methods.

Table 5.3: Learned feature two-class classifier results (mean).

Model	Accuracy	Precision	Recall	F1-score	FNR	FDR	FPR
ResUnet	98.66	98.57	98.67	98.62	1.33	1.43	1.33
U-Net	98.18	98.17	98.06	98.11	1.94	1.84	1.94

Table 5.4: ResUnet classifier results — individual classes.

Class	Accuracy	Precision	Recall	F1-score	FNR	FDR	FPR
Polymer	98.66	99.12	98.59	98.85	1.41	0.88	1.25
Marine	98.66	98.02	98.75	98.38	1.25	1.98	1.41

5.3 Results

Learned feature results

The mean results in Table 5.3 show that the ResUnet network style implementation yields a modest advantage compared to the baseline vanilla U-Net setup in terms of all performance metrics. Table 5.4 details the individual classes in the ResUnet implementation, and shows that its FPR is 1.25% when considering the polymer class. This metric is important as we want to limit any false detections of the polymer class and thus it should be as low as possible. The U-Net implementation (Table 5.5) obtained an FPR for the polymer class of 2.57%. The accuracy and precision results are similar for each class. This suggests that the classifiers did not favour one class over another and therefore shows limited bias. These results are further discussed in the analysis section and examined in the mixed-media experiments in the next chapter. This learned feature approach appears to offer an appropriate solution to the binary classification problem with the customised ResUnet implementation being put forward as the selected model.

Filter bank feature extraction — binary results

Table 5.6 displays the model performance, based on the 20 000 feature vector test set. The random forest and KNN implementations demonstrated the most favourable performance metrics, but all of the classifiers produced competitive results. Tables 5.7(a) and 5.7(b) present the confusion matrices for the two best-performing models (random forest and KNN) where again the performance is comparable. Detailed confusion matrices and

Table 5.5: U-Net classifier results — individual classes.

Class	Accuracy	Precision	Recall	F1-score	FNR	FDR	FPR
Polymer	98.18	98.23	98.70	98.46	1.3	1.77	2.57
Marine	98.18	98.10	97.43	97.76	2.57	1.90	1.30

Table 5.6: Feature extraction — two-class classifier results (mean %).

Model	Accuracy	Precision	Recall	F1-score	FNR	FDR
KNN	97.06	97.08	97.06	97.06	2.94	2.92
Random forest	97.03	97.07	97.05	97.05	2.95	2.93
Tree	96.88	96.94	96.88	96.87	3.12	3.06
SVM	96.80	96.82	96.81	96.80	3.19	3.18
Discriminant	96.50	96.31	96.71	96.51	3.50	3.50
Naive Bayes	96.83	96.87	96.83	96.82	3.18	3.13

Table 5.7: Confusion matrices two-class with TPR on diagonal: random forest (left), KNN (right).

		Predicted %				Predicted %		
			SP	Marine			SP	Marine
Actual %	SP	96.10	3.90		Actual %	SP	95.97	4.03
	Marine	2.00	98.00			Marine	1.85	98.15

hyperparameters for all models can be found in Appendix A.

The hyperparameters obtained from the Bayesian optimisation routine for the top two performing classifiers are as follows:

- KNN: distance metric, Chebyshev; number of neighbours, 216; distance weight, inverse.
- Random forest: number of learners, 44; number of splits, 283.

The FPR results (in percent) when considering polymer as the positive class are as follows: tree, 1.27; naive Bayes, 1.69; KNN, 1.85; random forest, 2.00; SVM, 2.34, and discriminant, 3.71. While the KNN achieved the highest accuracy among the classifiers, it did not attain a suitable FPR and was therefore not selected to be put forward as a candidate. The lowest FPR was achieved by the tree model at 1.27% and this result is comparable to the 1.25% FPR achieved by the learned feature approach. However, the general mean metrics in terms of accuracy, precision, recall, and F1-score were all lower than the learned feature approach, thus showing less promise.

Filter bank feature extraction — four class results

Based on the test dataset, the performance metrics of the six different model implementations are displayed in Table 5.8. The hyperparameters obtained from the Bayesian optimisation routine for the top two performing classifiers are as follows:

- KNN: distance metric, Minkowski (cubic); number of neighbours, 268; distance weight, squared inverse.

- Random forest: number of learners, 496; number of splits, 12747.

Upon reviewing the classifier performance metrics in Table 5.8, it can be seen that the random forest model achieved the best overall classification accuracy of 95.03%, closely followed by the KNN classifier. The naive Bayes model sits at the bottom of the classification accuracy ranking at 91.53%. It was hoped that the SVM classifier would achieve stronger performance. However, it only presented accuracy figures in the mid-band of the accuracy ranking, possibly due to its sensitivity to noise or outliers.

Figure 5.6 displays the confusion matrices for all of the classifiers. If we consider the performance of different polymers with the different classifier models, it is observed that PE exhibits the highest classification accuracy in each of the models and also has the lowest variance across the models. PET, on the other hand, had the highest variance, whereas both PP and PS have a variance in between.

When examining the misclassification performance of individual polymers across various models, we observe that misclassifications were relatively evenly distributed among all classes for random forest (Figure 5.6(a)), tree (Figure 5.6(e)), and KNN (Figure 5.6(f)) models. However, the SVM model (Figure 5.6(b)) exhibited a higher frequency of misclassifications, with PET and PS often being mistaken for PP. The discriminant model showed a reasonably balanced distribution of misclassifications, except for PS which was frequently classified as PET. On the other hand, the naive Bayes model (Figure 5.6(d)) displayed the highest number of misclassifications, particularly with PP and PET being frequently classified as PS. This suggests that an ensemble approach of implementing multiple different classifiers could achieve better overall performance and may be worth investigating in future studies.

Upon reviewing the confusion matrix for the random forest ensemble (Figure 5.6(a)), it is evident that PS has the lowest performance with a true positive rate of only 86.3%. A significant portion of PS false negatives occurred when a PS particle was misclassified as a PP particle. This suggests a certain textural similarity between these two plastics, leading to confusion in their classification. The false negatives of PP are more evenly distributed among the other three classes and a similar pattern is observed with PS. Both PP and PET exhibit similar classification performance, with true positive rates of 85.6% and 86.6%, respectively. On the other hand, PE shows the best performance with a true positive rate of 94.8%, but it also displays an equal distribution of false positives among the other three plastics.

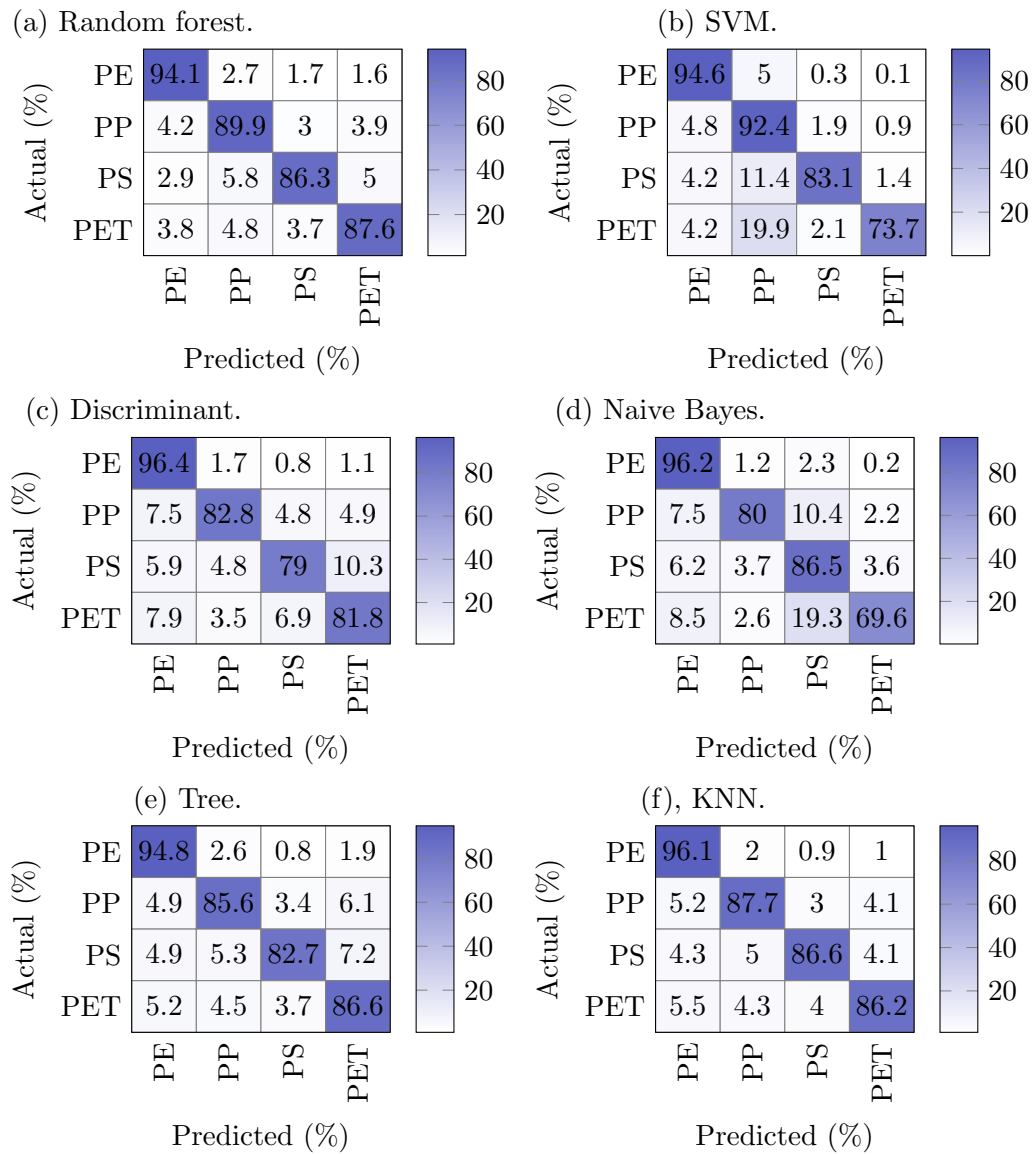


Figure 5.6: Classifier confusion matrices showing the TPR on the diagonal. (a) Random forest. (b) SVM. (c) Discriminant. (d) Naive Bayes. (e) Tree. (f) KNN.

Table 5.8: Feature extraction four-class classifier results (mean).

Model	Accuracy	Precision	Recall	F1-score	FNR	FDR
Random forest	95.03	90.08	90.06	90.04	9.94	9.92
KNN	94.58	89.29	89.17	89.13	10.83	10.71
Tree	94.94	86.33	94.76	87.37	12.59	12.46
SVM	92.96	87.84	85.93	86.04	14.07	12.16
Discriminant	92.47	85.17	84.94	84.83	15.06	14.83
Naive Bayes	91.53	84.42	83.06	82.96	16.94	15.58

Table 5.9: Random forest classifier results — individual classes.

Class	Accuracy	Precision	Recall	F1-score	FNR	FDR	FPR
PE	96.06	90.44	94.20	92.28	5.80	9.56	3.32
PP	94.51	88.23	90.05	89.13	9.95	11.77	4.00
PS	94.89	91.75	87.44	89.54	12.56	8.25	2.62
PET	94.65	89.91	88.54	89.22	11.46	10.09	3.31

5.4 Discussion and analysis

In the learned feature implementation, an attempt was made to transition from binary synthetic polymer and marine classification to a five-class classification exercise using encoder-decoder methods. However, this performed poorly. Given that the synthetic polymers in this study exhibit varying shapes, pure geometric measurements are believed to be less critical to a feature extractor and, rather, textures may be more useful. Improving performance in this context may require an expanded dataset, enhanced resolution through the development of software super-resolution, or the use of lens-based imaging to capture finer details (at the expense of analyser throughput). This presents an opportunity for future studies to explore further development in this domain.

Since the learned feature ResUnet model offered the lowest FPR for the polymer class, it is recommended as the appropriate candidate for the binary classification task. Furthermore it outperformed the other classifiers in terms of all other performance metrics. While adjusting the classifier’s decision boundary to achieve a lower false positive rate is possible, we opted to keep the decision boundary unchanged, where the class achieving the highest probability is the class output for a pixel. To ensure reproducibility, the decision was made to present the workflow and classifier performance in a raw and unadulterated state. Should one wish to optimise based on the FPR, either the decision boundary could be shifted or an offset applied to the logging output of the workflow. If an offset were to be applied, a percentage of the total positive counts could either be added or subtracted from the measurement log and applied to the negative class. However, this remedy only partially satisfies more complex scenarios such as considering individual particle areas and measurements over and above a raw particle count.

While the filter bank feature extraction workflow performed fairly well in the binary task of classifying synthetic polymer and marine particles, it exhibits slower processing compared to its neural network counterpart. More importantly, the FPR for polymer was higher in all instances and similarly the mean metrics were all less suitable in the classical approach. Now, if the focus is shifted towards being able to classify specific classes of synthetic polymers efficiently, it would make sense to use the neural network

model as a first stage to separate marine and plastic particles and the second stage purely for plastic classification. The benefit to this is that particles classified as marine do not require the extensive filter setup and overhead of the filter bank workflow. Hence, these particles can be swiftly processed and discharged without further action. Considering that marine particles might outnumber the sparse presence of microplastics in a real-world scenario, this approach may be a more resource-efficient choice. The method also facilitates the isolation of marine particles. This effectively creates a clear boundary for the polymer classification scheme, which has a strategy tailored specifically to polymer classification (considering the textural intricacies inherent to this task) and can be further developed for additional microplastic classification without affecting the marine and polymer decision boundary.

Although a pixel-wise approach was adopted in our classifiers, a shift to a global object feature vector could be integrated with our design. This could improve classification inference time and may be investigated in a future study where the trade-off between inference time and classification performance can be studied. Additional statistical approaches to texture analysis through computing object patch metrics such as local binary patterns, co-occurrence distributions, autocorrelation, power spectrum, Laws' texture energy, and Markov random fields are available resources for future studies. Texture compression techniques in the form of textons and histogram matching, as described in Varma and Zisserman's work [40], also offer potential for future studies. While being suited to a multi-polymer classification problem, sufficient object pixels are required for reliable inference. Unfortunately the minimum size of our object patches would not provide enough pixels, making this approach impractical. Enhancing FSA resolution through super-resolution research could render this implementation more attractive.

5.5 Conclusions

In this chapter, we used the FSA and crafted classifiers for microplastic analysis. A small dataset including marine phytoplankton and microplastics was created. Both learned feature and filter bank feature classifiers were investigated and implemented. In the binary (microplastic and marine) classification task, a customised ResUnet learned feature network was selected as the most suitable choice. For the multi-class polymer problem, the filter bank feature proposal was the most promising candidate. This filter bank approach incorporated filter banks, supervised dimensionality reduction, and a random forest classifier. The amalgamation of these subsystems is developed in the next chapter and comprehensive mixed-media experiments are conducted to assess this end-to-end workflow.

Chapter 6

Pipeline

Conventional microplastic detection methods are often labour-intensive, costly, and offer limited sample throughput. A processing pipeline was established which is presented for automated end-to-end microplastic classification in a fluid stream. Our method integrates the FSA and classifiers developed in the previous chapters to produce a pipeline or workflow and protocol to support microplastic testing beyond standard classification results. A simplified marine environment was created and spiked with varying concentrations of microplastics to allow for workflow evaluation. The assessment encompassed testing with mixed media samples, including binary microplastic-marine classification, a five-class classification involving marine and individual microplastic types, and the classification of geometrically similar polymer calibration spheres composed of two different plastic types. The results obtained support the FSA and classification techniques developed as a proof of concept; this suggests that our workflow has the potential for further development as an analyser suitable for resource-limited, time-sensitive, or point-of-sampling testing.

6.1 Preliminaries

Automated detection and classification of microplastic particles in a fluid stream remains challenging. Existing microplastic detection techniques often require human involvement, which impacts the ability to support a high throughput analysis. Samples either need significant pretreatment steps or certain tasks in the processing chain are human-aided and could therefore benefit from automation. We wish to move towards automating the detection and classification of microplastics in a fluid stream in a low-cost manner.

This chapter establishes and presents a processing pipeline for end-to-end microplastic

classification in a fluid stream. Our method puts together a pipeline and protocol to support the FSA and classifiers developed in the previous chapters. This was done to provide a more holistic assessment of the system.

The developed FSA and classifiers show utility but must be more formally assessed. In general, existing studies on microplastic classification using optical imaging methods often conclude their experiments by presenting classification results on their designated test datasets. These test datasets typically consist of sample images captured concurrently with the training data sample images for their models, and the analysis is often limited to a focused approach closely aligned with a well-defined ground truth. While this approach is commendable, it suggests the need for additional real-world experiments to ascertain the methods' broader applicability to polymer classification.

This motivation led to the pipeline development in our study. We further introduced a hybrid-classifier workflow following the positive results that were obtained from both a learned and filter bank feature approach in Chapter 5; the hybrid workflow includes both an encoder-decoder and a random forest in a sequential approach. This workflow accepts an image sequence stream from the FSA developed in Chapter 4 and produces class predictions along with key shape metrics for logging purposes.

A simplified marine environment was created and spiked with varying concentrations of microplastics. The workflow's performance was assessed through mixed media samples testing, including the binary (microplastic-marine) and five-class (marine and individual microplastic type) classification problems. This chosen approach enables the comprehensive testing of the proof of concept FSA through end-to-end experiments. Using mixed media experiments in a simulated marine environment allows for a move toward assessing the system's performance in more diverse and realistic scenarios that still have well-defined ground truths.

We also address a unique scenario where we trial a PMMA and PE spheres workflow. These spheres exhibit similar morphology but differ in birefringent characteristics. Here, we implement a condensed workflow with the target to classify these particles based on their interior texture and without the influence of any shape characteristics.

The remainder of this chapter starts by introducing the experiments and detailing the implementation of the comprehensive end-to-end workflow. This is followed by discussing the establishment of a simulated marine environment, which includes using two phytoplankton cultures. The general techniques and procedures adopted to ensure the reproducibility of our methods are covered. Next, we cover the experimental design adopted for our workflow assessment in which three experiments are performed. This

is followed by the results, where we evaluate the workflow's performance as a binary classifier for marine and polymer classes, yielding promising results and a low FPR. We then assess multi-class polymer classification. Finally, the classification of geometrically similar polymer calibration spheres of two different plastic types (PMMA and PE) is interrogated.

Results suggest that our workflow may be useful for further development of an analyser for resource-limited, time-sensitive or point-of-sampling testing and support the proof of concept for the FSA and classification techniques.

6.2 Workflow

This chapter aims to develop a workflow that amalgamates the FSA and microplastic classifiers alongside further structures to provide an end-to-end classification of microplastic particles. The workflow must be appropriately designed, implemented, and trialled to determine suitability. A simulated marine environment and a protocol were established to support the experiments.

We developed three experiments to assess the workflow and refer to them as follows:

- **Experiment 6A:** Combined microplastic analysis in a marine environment.
- **Experiment 6B:** Multi-class microplastic analysis in a marine environment.
- **Experiment 6C:** PE and PMMA sphere comparison.

The aim of Experiment 6A is to establish if the workflow can detect and distinguish microplastic particles from marine particles as a single class, whereas 6B aims to further categorise observed microplastic samples by their individual class types. Finally, 6C aims to classify similarly shaped objects based on their birefringent texture in which we study PE and PMMA spheres.

We start by designing the workflow (which is adaptable if required). Then we implement the experimental protocol, and finally we carry out the three experiments.

6.2.1 Algorithm implementation

Our complete end-to-end workflow is presented here and this architecture is used for all three experiments (with Experiment 6C using a lightweight implementation). The processing chain from image acquisition to output classification is detailed in Figure 6.1. We make use of three sub-systems from the previous chapters: the developed FSA including both hardware and software elements, the marine-polymer ResUnet network

developed in the learned feature study, and the multi-class polymer classifier from the filter bank feature study.

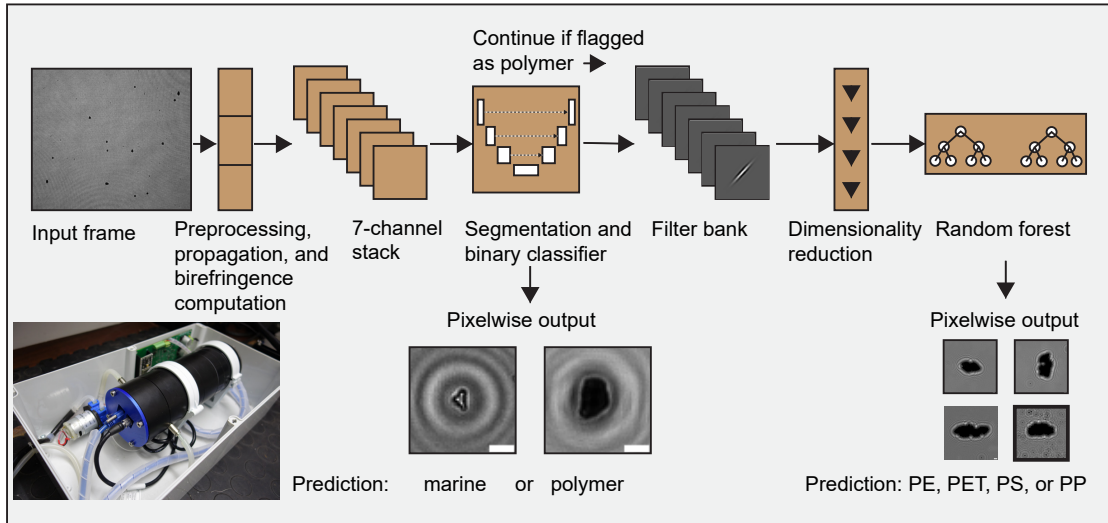


Figure 6.1: Pipeline overview. The input frame acquired by the FSA undergoes the preprocessing and base feature extraction routine. The expanded seven-channel stack is passed to the neural network for segmentation and classification. The particles identified as polymer type are then sent for further feature extraction with the filter bank. This high-dimensional feature vector undergoes dimensionality reduction. Finally, the random forest classifier will perform a pixel-wise prediction on whether the polymer flagged sample is PE, PET, PS, or PP.

The initial monochrome raw hologram input image undergoes preprocessing as described in Chapter 4, where the image is first background corrected. The channels are then expanded to seven through the FSA functionality, which includes using the angular spectrum method to obtain mid-depth propagated intensity and unrecovered wrapped phase. The mid-depth intensity frame is demosaiced to access each of the four linear polarisation orientations and is held in four separate 2D matrices. These matrices are sent for birefringence and Stokes parameter 1 and 2 calculations.

Birefringence retardance and slow axis orientation are calculated using the method described by Shribak and Oldenbourg [14]. However, slow axis orientation values are paired with their corresponding retardance values only if the computed retardance exceeds a certain threshold (to be considered valid). Specifically, instances of slow axis orientation are only chosen when the computed retardance is greater than 0.1 radians. This criterion is applied to prevent the inclusion of artefacts and ensure that the slow axis orientation data is grounded in physically meaningful retardance values. Otherwise, the slow axis result for that pixel will remain at 0 radians. The resulting retardance and slow axis maps are upsampled by a factor of two using a bilinear kernel to restore the original image pixel dimensions.

The seven-channel frame is now directed to the custom ResUnet network for segmentation and classification of both marine and polymer particles. The frame is partitioned into 128×128 patches and fed into the network. The network's output is a pixel-wise classification frame of the same size as the input, labelled 0 for background, 1 for polymer, and 2 for marine. The 128×128 output mask patches are reassembled to create the original frame size.

The image mask obtained will likely contain multiple objects in a frame. This mask undergoes morphological processing to eliminate connected components containing fewer than 125 pixels. This step removes smaller, non-relevant objects and artefacts that fall below the interest threshold, reducing processing overhead by avoiding unnecessary object extraction and classification stages for these objects. In diverse settings, the detector may encounter particles like fines or small objects that are close to the image sensor's resolving limit, especially in more turbid environments. As a criterion for this project, objects with a length exceeding $50 \mu\text{m}$ are of interest to the classifier. Given that this project uses a $3.45 \mu\text{m}$ pixel pitch image sensor, a radius of approximately 7.25 pixels is necessary to accommodate the circular area projection of a minimum-length object ($50 \mu\text{m}$). Thus, a minimum object area of 164 pixels is required for a circular object.

Following the removal of objects falling below the area threshold, each retained connected component in the image mask is assigned a distinct object identifier in a designated data repository, along with their class (polymer or marine). The class is determined by an ensemble voting approach, where the class prediction of each pixel in the object patch contributes a vote. Additionally, dimensions including the projected 2D area, centroid coordinates, and bounding box coordinates are computed and assigned to the identifying tag.

The workflow could terminate at this stage if it is only used as a two-class marine and polymer classifier. Should a five-class classification be required, the next stage of the workflow would classify the specific polymer group of any object patches that are flagged as a general (non-marine) polymer.

A subset of the FSA channels (mid-depth propagated intensity, birefringence retardance, and birefringence slow axis orientation) are supplied as individual inputs to the analysis bank as described in the previous chapter. The selected feature vectors are normalised based on the standardised score approach. The standard deviation and mean values calculated over the training dataset are used in this process.

Each polymer-identified object is cropped through the obtained segmentation mask.

The periphery of the analysed object, stemming from the segmentation process, is more prone to uncertainty than the central pixels. This concern prompted the processing of the object mask by applying a 40% radial erosion, which is limited to the polymer classification process and does not affect the actual recorded object dimensions.

The eroded object patch is now processed to obtain each pixel and assign the appropriate high-dimensional feature vector to its designated pixel; the feature vector is normalised according to the standard score approach. Each feature vector now undergoes a dimensionality reduction through the trained UMAP model to obtain a feature vector of length five.

The random forest classifier developed in Chapter 5 accepts the feature vectors from the pipeline and outputs a class label for the polymer type (1, PE; 2, PP; 3, PS; and 4, PET). The classifier output contains all pixels related to the specific object under analysis. To determine the class of that object, the mode of the output is taken (in case some pixels are misclassified). This class label is then assigned to its appropriate object identifier in the system log. The process is repeated for the next object in the frame, and once all objects in the frame have been logged the workflow continues to the next frame.

6.2.2 Recreated marine environment and protocol

A marine stock was created to simulate a basic marine environment containing marine phytoplankton. The stock was maintained at a salinity level of 3.5‰ and was carefully prepared to minimise the introduction of unwanted microplastic contamination. Phytoplankton cultures of *Chaetoceros calcitrans* and *Isochrysis galbana* were sourced from the Marine Research Aquarium in Cape Town, South Africa to create the stock. The cultures were stored in Schott borosilicate media bottles and the concentration of microorganisms could be adjusted using sterile saline (NaCl and laboratory water) to create a stock representative of approximate mean ocean phytoplankton concentrations.

The two microorganism species used in this study exhibit a size range of 5 µm to 500 µm in length. Notably, *Chaetoceros calcitrans*, in particular, forms long chains or clusters consisting of multiple cells [61], resulting in significant size variations within a single object. The classification approach adopted in this study treats marine microorganisms as a unified class of objects, acknowledging their potential geometrical variations.

A baseline cell count of the acquired cultures was performed to be able to adjust the stock concentration quantitatively. The obtained phytoplankton cultures were analysed with a haemocytometer counting chamber (Helber Z3000, Hawksley) and a bright-field

Table 6.1: Cell count of phytoplankton cultures.

Species	Count 1 / mL	Count 2 / mL	Mean / mL	Std. dev / mL
<i>C. calcitrans</i>	13 437 500	16 562 500	15 000 000	1 608 978.70
<i>I. galbana</i>	22 500 000	20 625 000	21 562 500	29 612.14

microscope (CX-43, Olympus) with a 100X oil immersion objective lens (1-U2B235, Olympus). Counts were performed in duplicate and are shown in Table 6.1. While it is not the intention of this study to accurately recreate a marine environment solely based on two types of phytoplankton, the dilution of the stock concentrations must be reasonable in terms of either what may be encountered when sampling a raw marine solution or when concentrating a marine solution by an appropriate factor. In previous imaging-based phytoplankton studies, raw ocean water samples were concentrated up to 3000 times using trawl net methods and then further diluted, resulting in an effective concentration of 60 times the raw sample. High-concentration samples were defined as containing 2000 to 3000 objects/mL [16]. *Chaetoceros* is a common and abundant phytoplankton genus in typical marine waters. Some studies have reported it to constitute as much as 91% of total phytoplankton in certain areas, with *Chaetoceros* blooms reaching concentrations of up to 31 000 cells/mL [62].

In our experiment, the two obtained cell cultures were diluted in a sterile saline solution to create a 1:1000 dilution. Subsequently, equal parts of *Chaetoceros calcitrans* and *Isochrysis galbana* were mixed to form a solution of *Chaetoceros calcitrans* ($7\,500 \pm 804.49$ cells/mL) and *Isochrysis galbana*, ($10\,781.250 \pm 14.81$ cells/mL). This resulting net cell concentration is expected to be higher than what is typically observed in the wild but falls within the range observed in extreme cases. This concentration level also encompasses scenarios where a standard or typical sample has undergone some pre-concentration steps.

Since we lack a well-defined working range for expected microplastic concentrations in marine waters, we opted for significantly more marine particles than synthetic polymer particles. This approach allowed for effective assessment of the system's performance in low microplastic concentration environments and emphasised the importance of the FPR. Although seawater samples containing phytoplankton could have been directly obtained from the ocean, this was avoided due to concerns about potential contamination with microplastic particles. Ensuring a reliable and uncontaminated ground truth was the priority.



Figure 6.2: Raw cultures. Left: *Chaetoceros calcitrans*. Right: *Isochrysis galbana*.

General techniques and configuration

Upon conducting early investigations with mixed media samples, it was observed that attempting manual injection of samples into the flow cell using syringes or a peristaltic pump led to poor sample mixing and presented the workflow with an unrepresentative sample for analysis. This was primarily due to the varying densities of the different types of plastics and their hydrophobic propensity to adhere to the walls of the system and glassware. As a result, fewer particles could pass through the detector and the plastic proportions were significantly unbalanced. This presented further challenges when attempting stepped concentration processes such as serial dilution experiments, thus making them highly inaccurate in practice.

The adhesion made creating a well-mixed solution that could effectively be circulated between the flow cell and a holding container difficult. Moreover, if the solution is not well circulated, comparison against a ground truth leads to inaccuracy. The manual sample injection approach was abandoned in favour of developing a specific protocol to address the challenges posed by adhesion and to ensure effective fluid circulation within the system. This protocol was consistently applied to all flow cell experiments in a recirculating configuration. The following steps were taken to ensure optimal experimental conditions:

- **Surfactant addition:** To reduce particle adhesion to surfaces, a mixture of surfactants was added to each litre of the stock solution. This included dishwashing soap (0.1% v/v) and ethanol (0.1% v/v). These additives helped reduce the tendency of particles to stick to the surfaces within the system.

- **Circulation and mixing:** A 600 mL borosilicate beaker served as the containment vessel for the flow cell and it was placed on a magnetic stirrer (MSH10, Labcon South Africa) operating at approximately 150 revolutions per minute. A peristaltic pump (NKP-DCS-10B, Kamoer) maintained a continuous flow. This pump pushed the solution (at a constant flow rate of 20 mL/minute) from the beaker through the flow cell's inlet, allowing it to pass through the imaging system before being returned to the beaker. This controlled circulation ensured consistent and controlled testing conditions throughout all experiments.
- **Stock media holding configuration:** A holding vessel, namely a 1000 mL borosilicate beaker that is independent of the circulatory setup described above, was placed on a second magnetic stirrer (MSH10, Labcon South Africa) operating at approximately 150 revolutions per minute to keep the stock solution in a mixed state (whether solely marine, polymer, or a hybrid marine-microplastic mix as used in some experiments). Fluid was retrieved from this beaker with the peristaltic pump and transferred to a 500 mL measuring cylinder to perform fluid dilution when needed.
- **Cleaning:** After each processing run and between different experiments a thorough cleaning was done. The flow cell, pipes, and glassware were cleaned using soapy water to remove any residue that might interfere with subsequent runs.

Although the synthetic polymer samples used for these experiments were from the same batch originally procured, only virgin samples were used in the experiments and no samples were reused. The plastic samples used in each dataset acquisition or experiment were discarded after each data capture session and fresh samples from the batch were introduced for subsequent data capture sessions.

6.3 Materials and methods

The established workflow was examined through a series of experiments to assess the performance when encountering varying concentrations of microplastics in a marine simulated environment. These experiments were conducted on separate dates from when the original sample dataset was captured, constructed, and used to train the classifiers used. Furthermore, specifically for Experiments 6A and 6B, the phytoplankton culture samples used were obtained from a different, independent batch. This approach aimed to expose the workflow to a more real-world environment with the possibility of the phytoplankton being at different stages in their development, leading to geometric differences between samples. This could help expose any biases that may have been

overlooked.

6.3.1 Experiment 6A: Materials and methods

Combined microplastic analysis in a marine environment

The workflow's performance was evaluated with regard to its ability to differentiate between microplastics and marine microorganisms. Specifically, microplastics (PP, PE, PS, and PET) were grouped as one class and marine microorganisms (*Chaetoceros calcitrans*, *Isochrysis galbana*, and debris) were combined as another class, forming a two-class classifier. A serial dilution approach was adopted where the concentration of marine stock is kept constant whilst the microplastic concentration is varied. A well-functioning classifier should detect little change in marine particle concentrations and the microplastic concentrations should vary in step with the serial dilution process.

The experiment began with an initial measurement of the stock solution, with no microplastics added. This measurement served as the baseline or "zero plastic" reference. The marine stock solution was placed in a separate magnetic stirrer vessel. Initially, a volume of 200 mL of the stock solution was withdrawn from the vessel and transferred into the beaker on the circulatory setup. The system was left to circulate the media for two minutes as a settling period, after which 100 frames were recorded. This was carried out in duplicate.

Subsequently, serial dilution measurements were executed, involving the introduction of microplastics into the solution. To achieve a solution with a microplastic concentration of 0.08 g/L, equal quantities of the four plastic types were weighed using a scale (AS220.R2, RADWAG). While the PP, PS, and PET samples were used from the micronised powder samples, which are of various mixed shapes and sizes, the PE sample was an equal blend of both the PE micronised powder sample and the spheres of a defined shape and size. The reasoning for mixing two sample types of PE is that it allows the classifier to be tested on non-defined shapes as well as uniform spherical smooth shapes. The smooth spheres may provide an additional challenge when classifying against smooth round phytoplankton marine objects.

These plastic samples were added to the standard marine stock that was already present in the circulatory system configuration. This procedure creates a solution containing the highest desired microplastic concentration.

Each serial dilution was executed in duplicate and involved the following steps:

- A measurement run consisting of capturing 100 image frames was initiated. Following the completion of each capture cycle, 100 mL of the solution from the circulatory configuration beaker was directed into a separate measuring cylinder and subsequently discarded.
- To replenish the circulatory configuration beaker with fresh solution, 100 mL of marine base stock that was devoid of any plastic particles was added from the second independent magnetic stirrer holding vessel. This addition resulted in the beaker retaining its original fluid volume of 200 mL, while the concentration of plastic particles was approximately halved compared to the previous measurement.
- After introducing the fresh solution, a waiting period of two minutes was observed. This interval allowed for the solution to circulate effectively and ensure uniform distribution.
- The next batch of 100 frames was captured for analysis of the new solution composition, completing the measurement cycle.

Although this serial dilution process aimed to halve the plastic concentration at each stage, residual static plastic particles might still adhere to surfaces in both circulatory and holding configurations. However, despite this limitation, this approach offered a viable means of systematically diminishing the plastic concentration in the solution over successive measurement runs while maintaining a consistent marine concentration.

6.3.2 Experiment 6B: Materials and methods

Multi-class polymer analysis in a marine environment

The workflow's performance was evaluated regarding its ability to make predictions in the multi-class polymer problem. Specifically, microplastics, including PP, PE, PS, and PET, were treated as unique classes while marine microorganisms (*Chaetoceros californicus*, *Isochrysis galbana*, and general debris) were grouped as a separate class resulting in a five-class classification system.

Four sub-experiments involving serial dilutions were conducted to assess the detector's effectiveness. In these experiments, the concentration of a specific plastic of interest was systematically varied through serial dilution while maintaining a constant concentration of the marine stock and the remaining three baseline plastics.

An able workflow should exhibit minimal fluctuations in detected marine particles and non-selected plastic concentrations throughout these sub-experiments. Conversely, the concentration of the plastic of interest should track the changes induced by the serial

Table 6.2: PE and PMMA microsphere comparison.

Polymer	Size range (μm)	Particles within size range (%)	Density (kg.m^3)
PE	63 to 75	90	960
PMMA	53 to 63	85	1 200

Particle size specified by their diameter.

dilution process. This approach ensures the detector can differentiate accurately between the various plastic types and marine microorganisms, even when subjected to varying concentrations and complex mixtures.

The experimental setup followed the same procedures and used the same equipment as presented in the marine-polymer two-class run, with one difference: the baseline stock concentration now included three non-selected plastic types in addition to the marine stock. This baseline marine-microplastic stock had a microplastic concentration of 0.04 g/L that was set by equally weighing three non-selected baseline plastic types and mixing them with the standard marine stock.

Considering the added presence of baseline microplastics, and that the investigation of the marine-microplastic problem had been interrogated in the previous section, we opted for a lower concentration in the baseline marine stock to avoid overwhelming the image sensor. A concentration equal to 25% of the standard marine stock concentration was used, except in the case of the PP mixed media experiment where a 50% stock dilution was performed in error. The serial dilution was initialised for each experiment with a concentration of 0.16 g/L for the microplastic of interest.

6.3.3 Experiment 6C: Materials and methods

PE and PMMA sphere analysis

Given the task of classifying similarly shaped objects, this study assessed the workflow's ability to differentiate between two types of plastics with similar geometry as a Bernoulli trial. This experiment used two types of synthetic polymer spheres (Cospheric LLC, Santa Barbara, USA): (i) PE (CPMS-0.96, 63 μm to 75 μm , Lot#100923-4); and (ii) PMMA (PMPMS-1.2, 53 μm to 63 μm , Lot#1961-306-3). The spheres were used to train a classifier using a stripped-down version of the established workflow. However, this classifier was exclusively trained on the binary pair of plastic types, one of which (PMMA) has not been used in the dataset or classifier training before and may also serve to show the pipeline's ability to support the addition of further plastics on which to train.




Figure 6.3: PMMA and PE sphere propagated intensity and birefringence retardance channels. (a) PMMA propagated intensity. (b) PMMA birefringence retardance. (c) PE propagated intensity. (d) PE birefringence retardance. Scale bar in Figures 6.3(a) and 6.3(b) indicates $50\ \mu\text{m}$.

Samples of PMMA spheres were passed through the FSA and imaged. These particles underwent the standard preprocessing procedure and were labelled pixel-wise as when the original project dataset was constructed. Existing labelled PE sphere data was available from prior sample acquisition. A visual inspection was also performed. Looking at the focus intensity channels of both types of plastics in Figures 6.3(a) and 6.3(b) it is noticed that the focus intensity channels of both plastics look similar, with minor geometrical differences. Now if we look at the retardance map in Figures 6.3(c) and 6.3(d), a significant difference is visible as the birefringent activity is higher in the PE sample. This is expected, as PE should possess a greater optical path difference between the slow and fast axis refractive indices than PMMA.

As shown in Table 6.2, there are differences in the size fractions of the two test particles. This discrepancy poses a challenge, as the output features generated by the analysis bank in the filter bank machine-learning setup can be influenced by the particle size. Larger particles inherently have a greater boundary region with a higher percentage of total area pixels that make up these regions. Given that the area of the circular

projected area increases proportionally to the square of its radius, the peripheral pixels in the outer region of the PE spheres will considerably outnumber those of the PMMA spheres. We adopt four techniques to reduce the risk of the classifier exploiting these size differences for inference.

Firstly, we replace our filter analysis bank in the pipeline with a reduced filter bank comprising only the rotationally invariant Schmid filter set [42], and adjust the kernel to size 3×3 . These kernels focus on more localised regions of the object and are much smaller than the typical object area of between 155 and 320 pixels in this application.

Secondly, we radially erode the observation area of the object by 60% to analyse only the inner region. This eliminates the potential effects of varying boundary pixel ratios between classes. The filter supports are not expected to straddle regions on or over the boundary as well as these inner regions concurrently. Therefore we only use this extreme inner region on the object to create the pixel-wise feature vectors.

Thirdly, we introduce random sampling into the dataset and use a balanced 50/50 split of training data between the two classes. This safeguards against the classifier holding bias or relying on size-related distinctions when generating predictions.

Fourthly, after tuning the hyperparameters, no manual decision boundary adjustments are made after training.

The standard pixel-wise labelling procedure was performed to provide a ground truth for the smaller dataset. The reduced dataset containing PMMA and PE underwent a limited feature extraction process employing the small analysis bank and the standard FSA features only. Subsequently, augmentation was implemented and the feature set underwent a similar dimensionality reduction procedure as the previous classification setups, albeit this time to a 3D mapping. The UMAP model was tested with unseen data. The post-mapping output is illustrated in Figure 6.4, where we have added back the class labels for visualisation. The two synthetic polymers manifest as modest clusters, although certain overlaps exist between them.

Following the classifier implementation, mixed-media serial dilution experiments were performed in the same manner as for the previous mixed-media experiments. The concentrations of the two sample plastics were varied in two separate sub-experiments. Firstly, a constant background concentration of PE was maintained in the solution while the concentration of PMMA was adjusted. Secondly, the background concentration of PMMA was kept constant while the concentration of PE was varied. This allowed for evaluating the classifier's performance in distinguishing geometrically similar plastics under varying concentration conditions.




Figure 6.4: PE and PMMA mapping.

Due to the significant difference in density between the two synthetic polymers, there was a noticeable variation in the distribution of particles when preparing each stock solution. When using the standard magnetic stirrer setup for both stock solutions, the PE spheres tended to accumulate more on the liquid surface. This contrasts the PMMA spheres, which were well dispersed within the solution. This effect was more pronounced in the 1 L beaker stock solution compared to the 600 mL flow cell feed beaker: the 1 L stock beaker typically contained between 600 mL to 1000 mL of stock during different stages of the serial dilution experiment, depending on the specific stage. Running the magnetic stirrer at the same speed for both batches resulted in significantly different mixing profiles for the two synthetic polymers.

On the other hand, the primary agitation stirrer in the flow cell loop had a fixed volume of 200 mL in a 600 mL beaker. This meant that the same mixing speed resulted in more vigorous agitation and provided satisfactory mixing for both synthetic polymers. Nevertheless, issues such as mild adhesion to glass surfaces and piping could still influence particle circulation, even with the surfactant protocol in place.

6.4 Results

In this section, we start by reviewing the results from the combined microplastic analysis experiment (6A) in which the marine concentration was held constant and the synthetic polymer concentration was adjusted in stepped increments. Additionally, the particle size distributions for both classes are analysed. Next, the multi-class experiment (6B)

is assessed in which the stock concentration of marine and polymer particles was held constant whilst an individual polymer type was adjusted in stepped increments. We conclude by running through the performance of the PE and PMMA classifier workflow experiment (6C) where morphologically similar polymers were analysed.

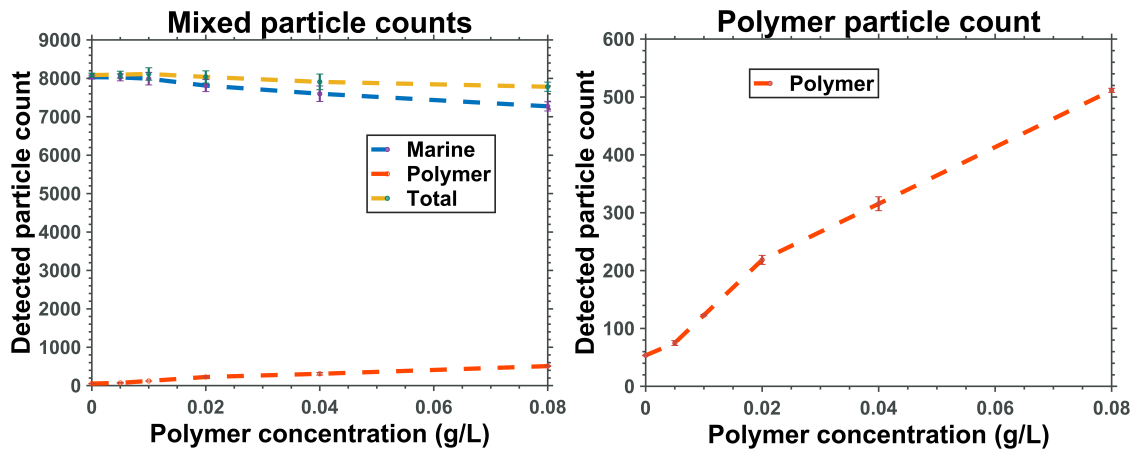
6.4.1 Experiment 6A: Combined microplastic analysis in a marine environment

The combined microplastic analysis experiment adjusted the concentration of a mixed polymer sample whilst the marine concentration was held constant. After completing the data acquisition session, the captured sequences were processed through the workflow.

The recorded number of detected particles (both marine and synthetic polymer) are presented in Figures 6.5(a) and 6.5(b). An approximately linear trend is observed in Figure 6.5(b) as the polymer concentration is adjusted. In Figure 6.5(a), it is evident that the total particle counts exhibit a slight decline as the synthetic polymer concentration is increased. This decrease could be attributed to some larger synthetic polymer particles potentially occluding smaller marine particles at higher concentrations. In other words, a single large synthetic polymer object may occlude many small marine particles at one instance, reducing the overall particle count. On the other hand, if the particle size distributions vary slightly between the two classes, this may also lead to the variation in particle counts. From low to high polymer concentrations, the number of polymer particles constitutes approximately 1% to 7% of the total particle count (combining marine and polymer particles). This suggests the analyser's ability to cope in environments where the ratio of polymer particles to marine particles is low, as may be encountered in the wild. However, when the same experiment assesses the total projected area of the detected particles (Figure 6.6), a contrasting trend emerges indicating an increase in the total projected area with rising microplastic concentration. These outcomes assist in determining the upper threshold of microplastic concentration that a detector like this one can handle effectively. In cases of high concentrations, the sample can be diluted or an image sensor with a larger active area combined with an increased channel width can be employed, thereby offering an expanded field of view within the existing configuration.

In the same experiment, when focusing solely on the synthetic polymer particles as depicted in Figure 6.5(b), the increase in detected polymer area becomes apparent as the synthetic polymer concentration is adjusted in stepped increments.

Reviewing the FPR is important, especially due to the unknown but potentially low



(a) Synthetic polymer concentration adjustment — particle counts. (b) Synthetic polymer concentration adjustment — polymer particle counts.

Figure 6.5: Synthetic polymer concentration adjustment. The error bars indicate the standard deviation and the sample points indicate the mean value of the repeated runs of the experiments.

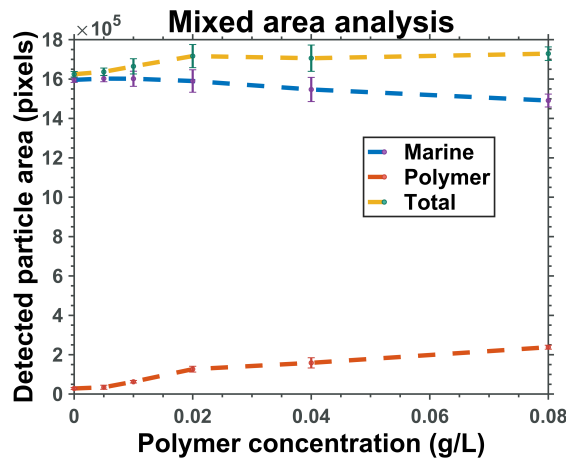


Figure 6.6: Synthetic polymer concentration adjustment — area analysis.

presence of synthetic polymer samples in real-world marine settings. The FPR achieved in this case was measured at $0.66\% \pm 0.01\%$ with the known fixed pure marine sample concentration and a 0.00 g/L synthetic polymer concentration. An additional observation is that in both Figures 6.5(a) and 6.6, the increase in polymer concentration does not lead to an increase in detected marine particles or area. This may be interpreted as the workflow demonstrating limited bias towards predicting marine particles and suggests that the classifier is balanced.

The FPR from the two-class ResUnet classifier confusion matrix was calculated at 1.25% when treating the synthetic polymers as the positive class, and is comparable to the

0.66% \pm 0.01% achieved in this experimental application. This suggests that the model was not overfit and that the augmentation protocol as well as the ensemble approach to the pixel-wise prediction output may have contributed to this improved result. Additionally, this workflow could be applied to more diverse datasets and samples for model development with the potential for further real-world experimentation.

Particle size distributions

The particle area data from the synthetic polymer concentration adjustment experiment was used to create a normalised histogram, which approximates a PDF based on the experimentally measured data. The following approach is used to generate the measured PDF:

$$f_i = \frac{c_i}{B \cdot w_i}, \quad (6.1)$$

where f_i is the i^{th} histogram bin, w_i is the bin width, c_i is the number of elements in the bin, and B is the number of bins in the histogram. The normalised histograms pre-

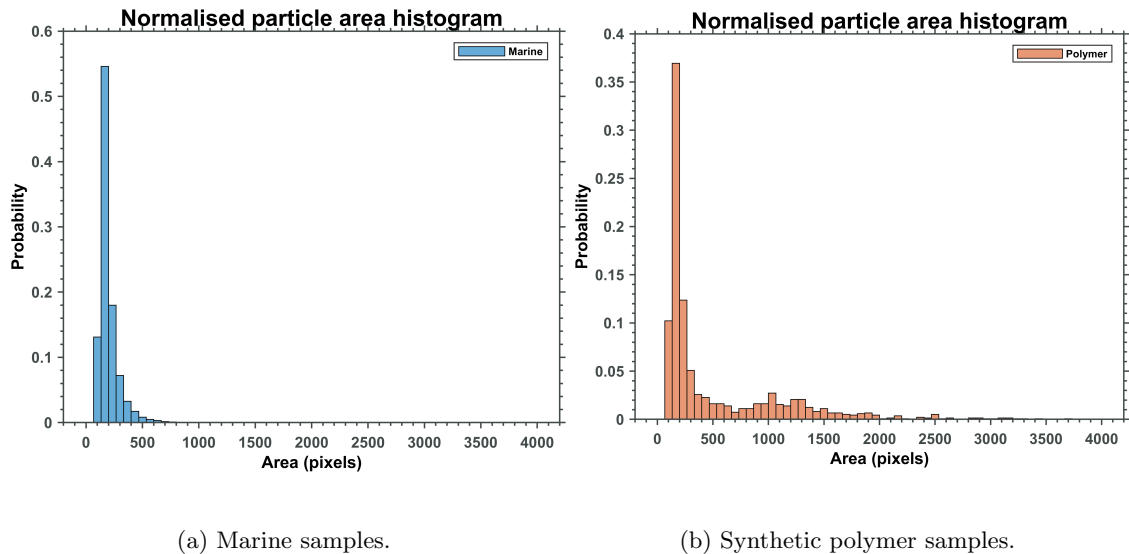


Figure 6.7: Normalised histogram (PDF approximation). Synthetic polymer and marine.

sented in Figure 6.7 provide an approximate particle size distribution for both marine (Figure 6.7(a)) and synthetic polymer particles (Figure 6.7(b)). The distribution of the two object types shows that both exhibit peaks in the 133 to 267 pixels range, and that the marine particles are dominant in a more tightly bound area distribution. It is not anticipated to see many marine particles in larger size fractions due to the small size of the individual cells. This suggests that these larger size fractions would be owing to linked clusters or chains of cells and otherwise non-synthetic polymer particulates. The frequency of marine particles sharply declines beyond 533 pixels. The synthetic

polymer particles also experience a drop-off but, in contrast, maintain a noticeable presence up to the sample area limit of 3 500 pixels. This suggests that the particle size distributions of the two classes are similar, with marine particles contributing more to the lower end of the area range and synthetic polymers occupying the upper end of the area distribution. This common overlap between the two classes allows for meaningful comparisons, providing the classifier with objects with similar area characteristics. However, since the analysis is only on the projected surface area of the object, the actual volume distribution may be a different consideration entirely.

6.4.2 Experiment 6B: Multi-class microplastic analysis in a marine environment

We now look at the multi-class experiment results where a base concentration of synthetic polymers and marine particles was held constant whilst an individual polymer type of interest was adjusted in stepped changes. After completing the data acquisition session, the captured sequences were processed through the workflow. A metric derived from the output log is the total projected area occupied by particles at different concentrations, as illustrated in Figure 6.8 for PET, PS, PE, and PP.

Starting with the PET plot in Figure 6.8(a), it is observed that the stock areas at the lowest and highest spike concentrations are similar and in agreement with expectations. However, a slight increase in the stock concentration is noted during the transition from 0.01 g/L to 0.02 g/L, possibly due to some particle adhesion during the serial dilution process. The concentration of PET particles follows a near linear trend with the stepped changes in concentration, although the transition from 0.02 g/L to 0.04 g/L does not display a significant change, possibly also due to adhesion.

Moving to the PS plot in Figure 6.8(b), a similar approximately linear trend is observed with the stepped increases in PS concentration. However, there is also an increase in the classified stock area, indicating that some PS particles might be misclassified as other polymer particles.

Next, examining the PE plot in Figure 6.8(c), a similar approximately linear trend is present with the stepped increases in PE concentration. The increase in stock area is however smaller, suggesting fewer misclassifications.

Finally, looking at the PP plot in Figure 6.8(d), the weakest linear trend of all four plastics is observed as the PP concentration is varied. PP is the least dense polymer of the four and is more prone to general adhesion, as it is likely to make contact with

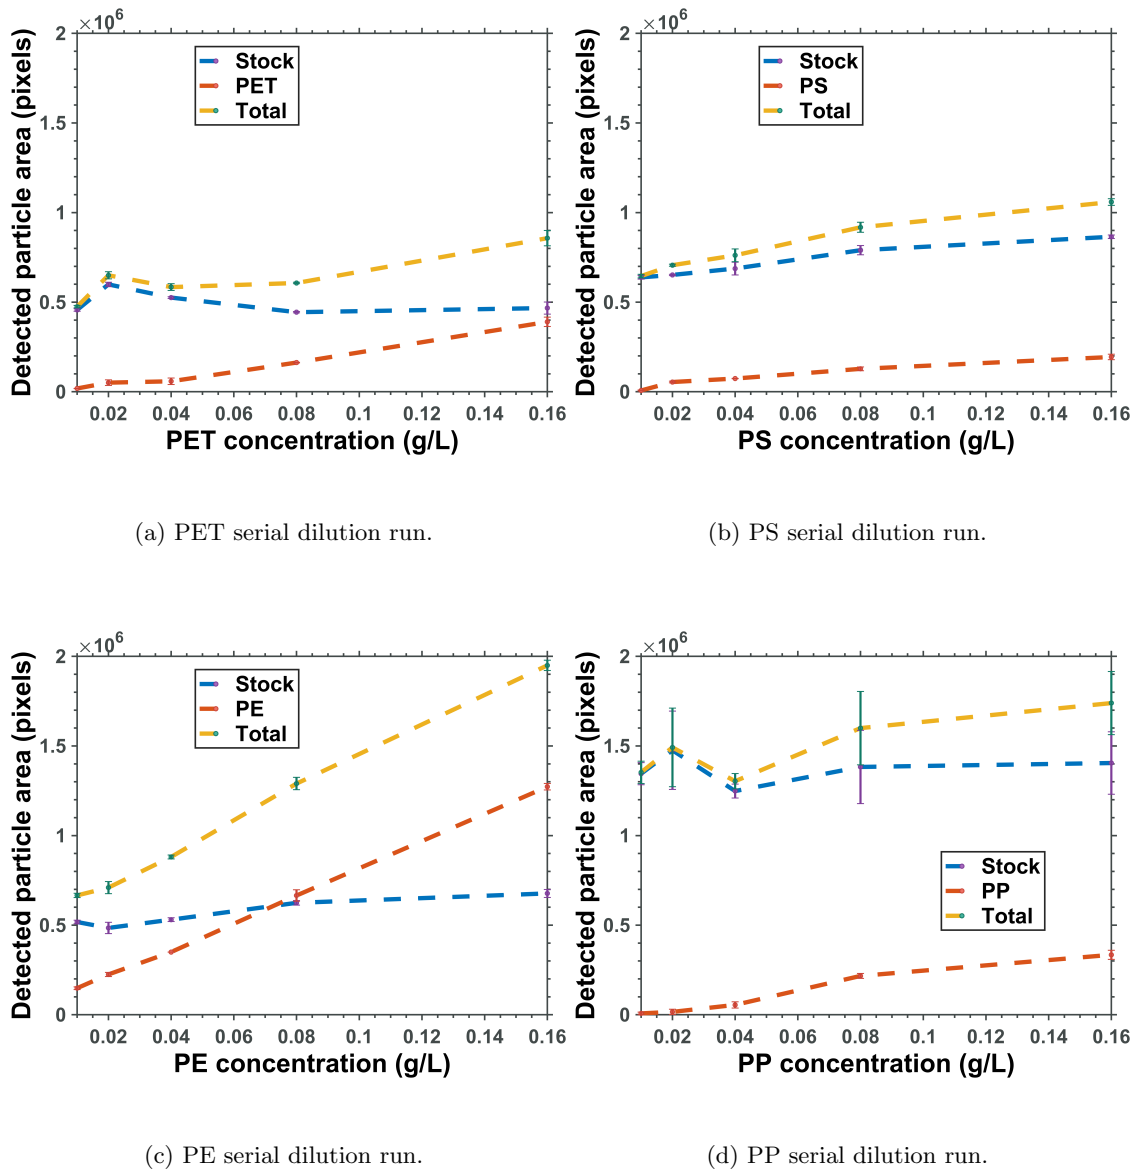


Figure 6.8: Mixed media experiments.

the liquid-free parts of the glassware where it could bind and remove itself from analysis. Additionally, a few misclassifications are present, although the stock area at the maximum PP concentration is only slightly higher than the area at the minimum PP concentration. The spike in the stock at 0.02 g/L could be attributed to the presence of other plastics in the stock causing a concentration spike.

A general observation is that the noted stock area value for a particular experiment is similar at the lowest and highest plastic concentrations, albeit slightly higher at the highest plastic concentration. This suggests that although some misclassifications occur, they are mild. Given the low FPR of the marine-polymer classifier in the workflow, these

Table 6.3: PE and PMMA confusion matrix — model test set.

		Predicted (%)	
		PE	PMMA
Actual (%)	PE	84.1	15.9
	PMMA	23.9	76.1

Table 6.4: PE PMMA confusion matrix live data capture.

		Predicted (%)	
		PE	PMMA
Actual (%)	PE	93.06	6.94
	PMMA	20.59	79.41

discrepancies are likely due to inter-plastic misclassifications that were also present in the confusion matrix of the random forest classifier. Density differences among plastics make some more prone to adhesion to glass surfaces exposed to air and may impact the general composition of the mixture in circulation. Additionally, as we analyse the 2D projected area, differences in volumetric characteristics may exist among plastic classes, impacting the observed area. Unfortunately we cannot further investigate these factors, as we are working with a mass of plastics for sample stock creation.

6.4.3 Experiment 6C: PE PMMA sphere analysis

The assessment of the PE and PMMA classifier workflow is now presented. The same classification learning methodology was applied as from the previous chapter’s filter bank feature extraction process, with a random forest model emerging as the top-performing choice. The random forest results are presented as a confusion matrix in Table 6.3. These results, predicting on a single pixel-wise feature vector input, can be deemed reasonably satisfactory.

In examining the mixed media results in Figure 6.9, it is evident that there is a difference in the base stock particle counts. In Figure 6.9(a) the mean PE base count is considerably lower than the PMMA base count in Figure 6.9(b), which could be due to density and area differences. Both adjusted synthetic polymers PMMA and PE exhibit a roughly linear trend that supports the stepped concentration changes but, as mentioned previously, adhesion issues may slightly impact this trend. The misclassifications in both instances are noticed with the slight change in the polymer that was held constant while the other polymer was adjusted.

Since we have a classifier that makes predictions on a single feature vector and it is a Bernoulli trial, it was anticipated that a performance improvement should be observed when taking the mode of a group of feature vectors that make up an object patch. This

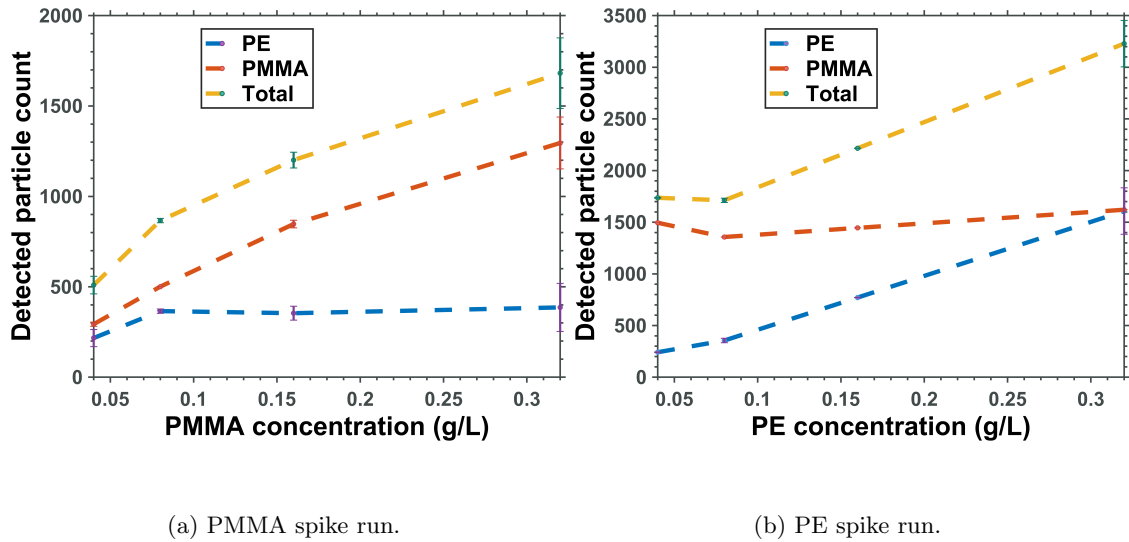


Figure 6.9: Stepped concentration analysis for two synthetic polymers. (a) PMMA concentration is changed whilst PE is held constant. (b) PE concentration is changed whilst PMMA is held constant.

Table 6.5: PE PMMA live data capture counts.

		Predicted	
		PE	PMMA
Actual	PE	389	29
	PMMA	118	455

is considered somewhat of an ensemble approach.

The results of running pure samples through the FSA and workflow from the two separate 100-frame capture runs for each synthetic polymer are displayed in Table 6.4 and Table 6.5. We show objects within a range of 20% above and below the maximum and minimum projected object area. These tables provide insight into the false detections of the workflow when applied to independent flowing samples. We can see that the classification improvement for the PE class is more than twice the improvement achieved by the PMMA class when comparing the original classifier performance (Table 6.3) to the observed mixed media trials (Table 6.4). Although a slight overall improvement in classification performance is noted when shifting to a capture session using the processing pipeline, it falls short of the expected improvement and suggests that the pipeline could be refined to improve resilience to noise. Considering that this is a stripped-down version of the analysis bank, it may benefit from using additional feature descriptors. Alternatively it may benefit from more input pixels in an object patch, as this experiment only considers a few pixels given the radial erosion process.

6.5 Discussion and analysis

6.5.1 General workflow

Since these experiments used virgin plastics, biofouling does not come into consideration. The impact of biofouling on microplastics adds complexity to sample analysis and could make for a useful future study. Particle adhesion proved a challenge in most of the experiments and further work is required to deal with this issue. Although the system is not optimised, the workflow could benefit from investigating the following: performing mixed media tests with reduced filter sets for comparative purposes or building a larger dataset that could make a learned feature multi-class polymer model more accessible. This approach could offer some performance gains. Since it has been observed that the device may offer benefits for low-cost, in-stream analysis of microplastics, the analyser could move to further refinement and field testing. Furthermore, the workflow can be adapted to sensing different sample types, including additional marine microorganisms, which should be encouraged for future studies.

6.5.2 Experiments 6A and 6B: Mixed media

For the mixed media binary experiment, the observed performance was an improvement over the standard classifier results from Chapter 5. This included an improved FPR and, as stated previously, the user could tune the decision boundary depending on their requirements. The histogram comparison showing similar particle distributions for the marine and synthetic polymer samples supports the argument that the workflow can perform classification without relying strictly on size metrics, although it is suggested that future studies could benefit from a more diverse number of phytoplankton cultures. The multi-class polymer experiments produced acceptable results, with some misclassifications that were noted amongst polymer types and not due to marine-polymer interactions. Further work is necessary to reduce the misclassification issues. Generally the stock concentrations at the lowest and highest polymer spikes were similar, which suggests low misclassifications. It was also observed that stepped concentration changes in a polymer of interest led to a concentration change in classified particles of interest, supporting the view that the workflow can distinguish one sample type from another.

6.5.3 Experiment 6C: PE and PMMA spheres

The PE and PMMA workflow task proved challenging in classifying morphologically similar polymers. However, the results support the idea of using filter-based feature extraction to develop a birefringent texture that a classifier may find useful.

6.6 Conclusions

The pipeline presented in this chapter was developed for automated end-to-end classification of microplastic particles. The combination of the FSA and developed classifiers was used to create the pipeline and three experiments were performed to assess its suitability. The results suggest that the work presented in this chapter has potential for further development and supports the effectiveness of the FSA and classifiers.

Chapter 7

Conclusion

In this final chapter, we summarise the main topics presented, provide concluding remarks, discuss future work, and highlight the contributions made throughout this research.

7.1 Summary of main topics

This work details the development, design, and implementation of a fluid stream analysis device dedicated to end-to-end microplastic detection. The analyser detects particles in a fluid stream using a lensless polarisation-sensitive holographic imaging-based approach. The development process included three objectives: the creation of the FSA, the development of classifiers, and the establishment of a robust pipeline.

FSA

The FSA integrates the necessary hardware and foundational software for fluid stream imaging. Outputs from the FSA serve as inputs for the subsequent objectives. The hardware components encompass a PFA image sensor, a custom-designed light pulse circuit (LPE), a prototype flow cell, a fibre-coupled laser diode, an optical polariser, and 3D printed components. Together these components form the FSA, which is a lightweight hardware implementation.

Software development included the integration of the following: an image preprocessing routine, a DIH propagation routine, a birefringence computation procedure, an encoder-decoder network for image segmentation, and control software for the LPE.

Experimental results and demonstrations validate the FSA's capabilities, supporting its ability to numerically propagate to a sample of interest in a fluid stream, execute segmentation, and retrieve birefringent features.

Computer vision classifier

The microplastic classifiers developed for this objective performed either of two tasks: a binary (marine and microplastic) classification, or a multi-class polymer classification. These classifiers use the outputs from the FSA as inputs to perform further feature extraction. Two approaches were explored: a neural network learned feature classifier and a filter bank feature classifier. A small dataset was created including marine phytoplankton and microplastic samples.

We found that a customised encoder-decoder network was most suited for the binary (marine and microplastic) problem. For the multi-class polymer problem, we found that a texture-sensitive filter bank was useful for polarisation-sensitive feature extraction in the case of birefringent plastics. Extracted features underwent selection and supervised dimensionality reduction, resulting in a 5D feature vector for classifier training. A random forest proved to be the most suitable model implementation.

Pipeline

The pipeline developed for this objective amalgamated the FSA and microplastic classifiers alongside further support tasks to provide end-to-end classification of microplastic particles.

This workflow automates accepting an image sequence stream as an input and provides particle class predictions along with key shape metrics for logging. In other words, the algorithm integrates developed subsystems including preprocessing, segmentation, and classification to output particle information. Additionally, we established a protocol and simulated marine environment for comprehensive assessment beyond classifier test results.

The pipeline underwent testing using different types of synthetic polymers and marine phytoplankton, demonstrating its capability to classify microplastics as a distinct class from phytoplankton in a laboratory setting. Furthermore, the system was tested for the classification of individual classes of synthetic polymers, yielding promising results.

7.2 General conclusions

The analyser, protocols, and workflow developed in this work are intended to serve as a proof of concept that supports the advancement toward in situ analysis of microplastic particles. The feasibility of using polarisation-resolved optical sensing for microplastic feature generation was presented in this work and provides a foundational framework for its use. However, it also acknowledges the need for further refinement and validation

in real-world scenarios.

This project aims to present a conceptual framework that can be taken forward to be optimised and integrated into embedded systems for future studies that may require field testing and deployment. It is envisioned that this concept can be further developed as a tool for use in other environments or for various types of studies, not necessarily limited to the applied microplastic classification problem. The interchangeability of the camera system, the customisability of the light source system, and the software code facilitate its adaptability and encourage further development.

Should a future project wish to continue on the microplastic classification task, it may benefit from increasing the number of sample types. This study took an initial step in providing a workflow that had been introduced to four microplastic types and two phytoplankton types. Therefore, broadening the coverage of sample types could make the concept more suitable for real-world deployment. The same workflow and assessment could be followed.

An investigation into using the deep learning space to develop methods to improve lensless image reconstruction based on a raw hologram would also make for an interesting study. This would be especially useful in the microplastic space and it has been successfully used in other environments.

7.3 Contributions

On a final note, we visit the contributions and provide a brief comment. Our key contributions are that we:

- Present an analyser concept that provides a cost-effective, label-free, and automated analyser using DIH and polarisation-sensitive techniques to provide quantitative birefringence-resolved sensing for microplastic classification. This concept may provide the foundation to be developed into a stand-alone prototype.
- Develop a pixel-wise classification scheme for microplastic particles in a fluid stream using DIH and polarisation resolved sensing. The move to a pixel-wise approach deviates from the global object feature vector approaches that are more common in the literature for learned feature classifier implementations.
- Illustrate the appropriateness of using a PFA image sensor for dynamic fluid stream analysis of microplastic particles in a single-shot manner. These sensors are effective when coupled with a software backend to extract full birefringence maps

of particles; this helps avoid a complex and sensitive optical setup. Future development for polarisation-sensitive deployable fluid stream analysers could benefit from using these sensors.

- Highlight lensless birefringent texture of microplastics and its usefulness over standard intensity measurements to train a classifier. While the birefringence of microplastics is a known property and has been used in machine learning classifiers, the lensless imaging domain's use of microplastic birefringence analysis has been relatively unexplored.
- Formulate a pipeline to incorporate birefringence sensing and machine learning to create an end-to-end automated routine. The end-to-end workflow is something that has been lacking in the literature and this pipeline provides some of the glue to assess our models out of isolation. The reconfigurability of the pipeline also makes it suitable to exchange models and functions depending on the analyser requirements.
- Develop a simulated mixed media protocol and test procedure that moves beyond standard classifier model results for microplastics. While this protocol is not exhaustive, it moves in the right direction for research to benchmark the microplastic classification process more thoroughly.

Bibliography

- [1] J. C. Prata, J. P. da Costa, A. C. Duarte, and T. Rocha-Santos, “Methods for sampling and detection of microplastics in water and sediment: A critical review,” *TrAC Trends in Analytical Chemistry*, vol. 110, p. 150–159, 2019.
- [2] A. Faltynkova, G. Johnsen, and M. Wagner, “Hyperspectral imaging as an emerging tool to analyze microplastics: A systematic review and recommendations for future development,” *Microplastics and Nanoplastics*, vol. 1, no. 1, 2021.
- [3] E. Miller, M. Sedlak, D. Lin, C. Box, C. Holleman, C. M. Rochman, and R. Sutton, “Recommended best practices for collecting, analyzing, and reporting microplastics in environmental media: Lessons learned from comprehensive monitoring of san francisco bay,” *Journal of Hazardous Materials*, vol. 409, p. 124770, 2021.
- [4] K. Waldschläger, M. Z. Brückner, B. Carney Almroth, C. R. Hackney, T. M. Adyel, O. S. Alimi, S. L. Belontz, W. Cowger, D. Doyle, A. Gray, and et al., “Learning from natural sediments to tackle microplastics challenges: A multidisciplinary perspective,” *Earth-Science Reviews*, vol. 228, p. 104021, 2022.
- [5] V. H. da Silva, F. Murphy, J. M. Amigo, C. Stedmon, and J. Strand, “Classification and quantification of microplastics (< 100 m) using a focal plane array–fourier transform infrared imaging system and machine learning,” *Analytical Chemistry*, vol. 92, no. 20, p. 13724–13733, 2020.
- [6] R. C. Thompson, Y. Olsen, R. P. Mitchell, A. Davis, S. J. Rowland, A. W. John, D. McGonigle, and A. E. Russell, “Lost at sea: Where is all the plastic?” *Science*, vol. 304, no. 5672, p. 838–838, 2004.
- [7] A. Cózar, F. Echevarría, J. I. González-Gordillo, X. Irigoien, B. Úbeda, S. Hernández-León, Á. T. Palma, S. Navarro, J. García-de Lomas, A. Ruiz *et al.*, “Plastic debris in the open ocean,” *Proceedings of the National Academy of Sciences*, vol. 111, no. 28, pp. 10 239–10 244, 2014.

- [8] M. Valentino, J. Béhal, V. Bianco, S. Itri, R. Mossotti, G. Dalla Fontana, E. Stella, L. Miccio, and P. Ferraro, “Synthetic microfibers discriminated by ai-enabled polarization resolved digital holography,” in *2022 IEEE International Workshop on Metrology for the Sea; Learning to Measure Sea Health Parameters (MetroSea)*. IEEE, 2022, pp. 511–515.
- [9] V. Bianco, D. Pirone, P. Memmolo, F. Merola, and P. Ferraro, “Identification of microplastics based on the fractal properties of their holographic fingerprint,” *ACS Photonics*, vol. 8, no. 7, pp. 2148–2157, 2021.
- [10] V. Bianco, P. Memmolo, P. Carcagnì, F. Merola, M. Paturzo, C. Distante, and P. Ferraro, “Microplastic identification via holographic imaging and machine learning,” *Advanced Intelligent Systems*, vol. 2, no. 2, p. 1900153, 2020.
- [11] J. W. Goodman, *Introduction to Fourier optics*. Roberts and Company publishers, 2005.
- [12] L. Sawyer, D. T. Grubb, and G. F. Meyers, *Polymer microscopy*. Springer Science & Business Media, 2008.
- [13] M. Valentino, J. Béhal, V. Bianco, S. Itri, R. Mossotti, G. Dalla Fontana, T. Battistini, E. Stella, L. Miccio, and P. Ferraro, “Intelligent polarization-sensitive holographic flow-cytometer: Towards specificity in classifying natural and microplastic fibers,” *Science of The Total Environment*, vol. 815, p. 152708, 2022.
- [14] M. Shribak and R. Oldenbourg, “Techniques for fast and sensitive measurements of two-dimensional birefringence distributions,” *Applied Optics*, vol. 42, no. 16, pp. 3009–3017, 2003.
- [15] B. Bai, H. Wang, T. Liu, Y. Rivenson, J. FitzGerald, and A. Ozcan, “Pathological crystal imaging with single-shot computational polarized light microscopy,” *Journal of biophotonics*, vol. 13, no. 1, p. e201960036, 2020.
- [16] Z. Göröcs, M. Tamamitsu, V. Bianco, P. Wolf, S. Roy, K. Shindo, K. Yanny, Y. Wu, H. C. Koydemir, Y. Rivenson *et al.*, “A deep learning-enabled portable imaging flow cytometer for cost-effective, high-throughput, and label-free analysis of natural water samples,” *Light: Science & Applications*, vol. 7, no. 1, p. 66, 2018.
- [17] C. Işıl, K. de Haan, Z. Göröcs, H. C. Koydemir, S. Peterman, D. Baum, F. Song, T. Skandakumar, E. Gumustekin, and A. Ozcan, “Phenotypic analysis of microalgae populations using label-free imaging flow cytometry and deep learning,” *ACS Photonics*, vol. 8, no. 4, pp. 1232–1242, 2021.

- [18] H. S. Auta, C. U. Emenike, and S. H. Fauziah, "Distribution and importance of microplastics in the marine environment: a review of the sources, fate, effects, and potential solutions," *Environment international*, vol. 102, pp. 165–176, 2017.
- [19] R. C. Jones, "A new calculus for the treatment of optical systems i. description and discussion of the calculus," *Josa*, vol. 31, no. 7, pp. 488–493, 1941.
- [20] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*. Springer, 2015, pp. 234–241.
- [21] D. Gabor, "Microscopy by reconstructed wave-fronts," *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 197, no. 1051, pp. 454–487, 1949.
- [22] A. Ozcan and E. McLeod, "Lensless imaging and sensing," *Annual review of biomedical engineering*, vol. 18, pp. 77–102, 2016.
- [23] R. A. Chipman, W. S. T. Lam, and G. Young, *Polarized light and optical systems*. CRC press, 2018.
- [24] J. Behal, M. Valentino, L. Miccio, V. Bianco, S. Itri, R. Mossotti, G. Dalla Fontana, E. Stella, and P. Ferraro, "Toward an all-optical fingerprint of synthetic and natural microplastic fibers by polarization-sensitive holographic microscopy," *ACS Photonics*, vol. 9, no. 2, pp. 694–705, 2022.
- [25] A. Greenbaum and A. Ozcan, "Maskless imaging of dense samples using pixel super-resolution based multi-height lensfree on-chip microscopy," *Optics express*, vol. 20, no. 3, pp. 3129–3143, 2012.
- [26] T. Colomb, P. Dahlgren, D. Beghuin, E. Cuche, P. Marquet, and C. Depeursinge, "Polarization imaging by use of digital holography," *Applied optics*, vol. 41, no. 1, pp. 27–37, 2002.
- [27] T. Liu, K. de Haan, B. Bai, Y. Rivenson, Y. Luo, H. Wang, D. Karalli, H. Fu, Y. Zhang, J. FitzGerald *et al.*, "Deep learning-based holographic polarization microscopy," *ACS photonics*, vol. 7, no. 11, pp. 3023–3034, 2020.
- [28] M. Sonka, V. Hlavac, and R. Boyle, *Image processing, analysis, and machine vision*. Cengage Learning, 2014.

- [29] D. Powers, "Evaluation: From precision, recall and f-measure to roc, informedness, markedness & correlation," *Journal of Machine Learning Technologies*, vol. 2, no. 1, pp. 37–63, 2011.
- [30] Sony, "Sony imx polarized sensors," Feb 2021. [Online]. Available: https://www.sony-semicon.com/files/62/flyer_industry/IMX250_264_253MZR_MYR_Flyer.en.pdf
- [31] Ibidi, "Ibidi channel slides," Aug 2020. [Online]. Available: https://ibidi.com/img/cms/products/labware/channel_slides/S_801XX_Slide_ILuer/IN_801XX_ILuer.pdf
- [32] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of machine learning*. MIT press, 2018.
- [33] A. Alwosheel, S. van Cranenburgh, and C. G. Chorus, "Is your dataset big enough? sample size requirements when using artificial neural networks for discrete choice analysis," *Journal of choice modelling*, vol. 28, pp. 167–182, 2018.
- [34] Z. Zhang, Q. Liu, and Y. Wang, "Road extraction by deep residual u-net," *IEEE Geoscience and Remote Sensing Letters*, vol. 15, no. 5, pp. 749–753, 2018.
- [35] D. A. Forsyth and J. Ponce, *Computer vision: a modern approach*. prentice hall professional technical reference, 2002.
- [36] D. Gabor, "Theory of communication. part 1: The analysis of information," *Journal of the Institution of Electrical Engineers-part III: radio and communication engineering*, vol. 93, no. 26, pp. 429–441, 1946.
- [37] S. Marçelja, "Mathematical description of the responses of simple cortical cells," *JOSA*, vol. 70, no. 11, pp. 1297–1300, 1980.
- [38] B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, no. 6583, pp. 607–609, 1996.
- [39] I. Fogel and D. Sagi, "Gabor filters as texture discriminator," *Biological cybernetics*, vol. 61, no. 2, pp. 103–113, 1989.
- [40] M. Varma and A. Zisserman, "A statistical approach to texture classification from single images," *International journal of computer vision*, vol. 62, pp. 61–81, 2005.
- [41] T. Leung and J. Malik, "Representing and recognizing the visual appearance of materials using three-dimensional textons," *International journal of computer vision*, vol. 43, pp. 29–44, 2001.

- [42] C. Schmid, “Constructing models for content-based image retrieval,” in *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, vol. 2. IEEE, 2001, pp. II–II.
- [43] O. Mudanyali, D. Tseng, C. Oh, S. O. Isikman, I. Sencan, W. Bishara, C. Oztoprak, S. Seo, B. Khademhosseini, and A. Ozcan, “Compact, light-weight and cost-effective microscope based on lensless incoherent holography for telemedicine applications,” *Lab on a Chip*, vol. 10, no. 11, pp. 1417–1428, 2010.
- [44] O. Mudanyali, C. Oztoprak, D. Tseng, A. Erlinger, and A. Ozcan, “Detection of waterborne parasites using field-portable and cost-effective lensfree microscopy,” *Lab on a Chip*, vol. 10, no. 18, pp. 2419–2423, 2010.
- [45] I. Navruz, A. F. Coskun, J. Wong, S. Mohammad, D. Tseng, R. Nagi, S. Phillips, and A. Ozcan, “Smart-phone based computational microscopy using multi-frame contact imaging on a fiber-optic array,” *Lab on a Chip*, vol. 13, no. 20, pp. 4015–4023, 2013.
- [46] A. Greenbaum, W. Luo, T.-W. Su, Z. Göröcs, L. Xue, S. O. Isikman, A. F. Coskun, O. Mudanyali, and A. Ozcan, “Imaging without lenses: achievements and remaining challenges of wide-field on-chip microscopy,” *Nature methods*, vol. 9, no. 9, pp. 889–895, 2012.
- [47] H. M. Shapiro, *Practical flow cytometry*. John Wiley & Sons, 2005.
- [48] Z. Göröcs, D. Baum, F. Song, K. de Haan, H. C. Koydemir, Y. Qiu, Z. Cai, T. Skandakumar, S. Peterman, M. Tamamitsu *et al.*, “Label-free detection of giardia lamblia cysts using a deep learning-enabled portable imaging flow cytometer,” *Lab on a Chip*, vol. 20, no. 23, pp. 4404–4412, 2020.
- [49] Y. Zhang, H. Wang, Y. Wu, M. Tamamitsu, and A. Ozcan, “Edge sparsity criterion for robust holographic autofocusing,” *Optics letters*, vol. 42, no. 19, pp. 3824–3827, 2017.
- [50] H. C. Hulst and H. C. van de Hulst, *Light scattering by small particles*. Courier Corporation, 1981.
- [51] F. Wang, M. Zhang, W. Sha, Y. Wang, H. Hao, Y. Dou, and Y. Li, “Sorption behavior and mechanisms of organic contaminants to nano and microplastics,” *Molecules*, vol. 25, no. 8, p. 1827, 2020.

- [52] V. Cheplygina, M. de Bruijne, and J. P. Pluim, “Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis,” *Medical image analysis*, vol. 54, pp. 280–296, 2019.
- [53] C. Shorten and T. M. Khoshgoftaar, “A survey on image data augmentation for deep learning,” *Journal of big data*, vol. 6, no. 1, pp. 1–48, 2019.
- [54] C. H. Sudre, W. Li, T. Vercauteren, S. Ourselin, and M. Jorge Cardoso, “Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations,” in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings 3*. Springer, 2017, pp. 240–248.
- [55] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.
- [56] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [57] Y. Rivenson, Y. Wu, and A. Ozcan, “Deep learning in holography and coherent imaging,” *Light: Science & Applications*, vol. 8, no. 1, p. 85, 2019.
- [58] P. P. Vaidyanathan, *Multirate systems and filter banks*. Pearson Education India, 2006.
- [59] J. Lee Rodgers and W. A. Nicewander, “Thirteen ways to look at the correlation coefficient,” *The American Statistician*, vol. 42, no. 1, pp. 59–66, 1988.
- [60] L. McInnes, J. Healy, and J. Melville, “Umap: Uniform manifold approximation and projection for dimension reduction,” *arXiv preprint arXiv:1802.03426*, 2018.
- [61] S. A. Spaulding, M. G. Potapova, I. W. Bishop, S. S. Lee, T. S. Gasperak, E. Jovanoska, P. C. Furey, and M. B. Edlund, “Diatoms.org: supporting taxonomists, connecting communities,” *Diatom Research*, vol. 36, no. 4, pp. 291–304, 2021.
- [62] B. Booth, P. Larouche, S. Bélanger, B. Klein, D. Amiel, and Z.-P. Mei, “Dynamics of chaetoceros socialis blooms in the north water,” *Deep Sea Research Part II: Topical Studies in Oceanography*, vol. 49, no. 22-23, pp. 5003–5025, 2002.

Appendix A

Supporting data

A.1 Confusion matrices for filter bank feature extraction

Table A.1: Confusion matrices two-class with TPR on diagonal: random forest (left), KNN (right).

		Predicted %	
		SP	Marine
Actual %	SP	96.10	3.90
	Marine	2.00	98.00

		Predicted %	
		SP	Marine
Actual %	SP	95.97	4.03
	Marine	1.85	98.15

Table A.2: Confusion matrices two-class with TPR on diagonal: tree (left), SVM (right).

		Predicted %	
		SP	Marine
Actual %	SP	95.02	4.98
	Marine	1.27	98.73

		Predicted %	
		SP	Marine
Actual %	SP	95.95	4.05
	Marine	2.34	97.66

Table A.3: Confusion matrices two-class with TPR on diagonal: naive Bayes (left), discriminant (right).

		Predicted %	
		SP	Marine
Actual %	SP	95.34	4.66
	Marine	1.69	98.31

		Predicted %	
		SP	Marine
Actual %	SP	96.71	3.29
	Marine	3.71	96.29

A.2 Dry screening procedure

Adapted from the Centre for Minerals Research, University of Cape Town.

Pre-inspection:

- Inspect electrical parts and connections for defects. Apply a lock-out procedure before the inspection. Inspect all mechanical parts such as belts, locking screws,

and guards for defects.

- Ensure that the correct PPE is worn at all times
- Ensure all equipment and surrounding areas are free of any traces of another sample.
- Ensure that all samples are labelled correctly.

Procedure:

- Ensure that no holes are present and that the screens have no tears.
- Check that the screens are not clogged.
- Stack the selected screens and a pan in the sieve shaker machine with the biggest aperture at the top and the smallest at the bottom.
- Pour the dry sample onto the top screen.
- Set the time and vibration to the required amounts.
- At the end of the screening period remove the screens and the bottom pan.
- Weigh the sample on each screen and record the screen size and mass.
- Clean the screens after use in an ultrasonic bath and air dry.
- Calculate the mass distribution by size.

A.3 Costing of the FSA

Table A.4: Costing of the FSA.

Item	Price (USD)
Polarisation image sensor	800
Laser diode and fiber line	160
Machining prototype flow cell, copper shims, glass channels, and end cap	180
Circular polariser	40
3D printer filament	20
Peristaltic pump, piping, and mountings	20
LPE circuit	70

Total: \$1 290

Total converted to rands (1 USD = ZAR 18.56): ZAR 23 942.40

Appendix B

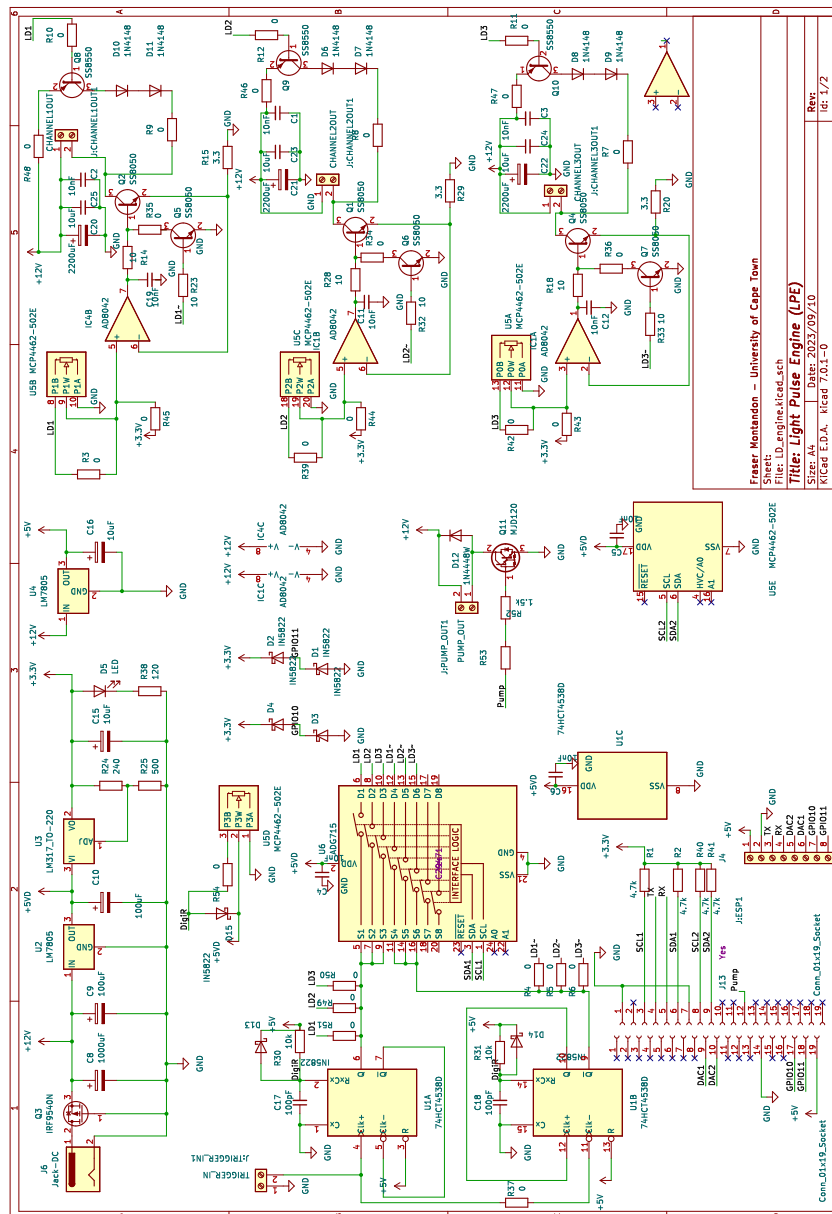


Figure B.1: LPE circuit.