

University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Anthony Lindley (2023), "Design, Fabrication and Characterisation of a Low-Cost, Acoustically Focussed Imaging Flow Cytometer for Automated Analysis of Phytoplankton", University of Southampton, School of Ocean and Earth Science, PhD Thesis, pagination.

COVID Statement

Striking in the middle of my PhD journey, the COVID-19 pandemic caused unprecedented disruptions to research activities worldwide. The constraints posed by the pandemic and associated lockdowns had a profound effect on my initial research plan, which was heavily reliant on experimental, biological work. The unforeseen closure of laboratories through multiple lockdowns hindered my ability to conduct the experimental studies I had planned, and I was faced with the challenge of adapting my research approach under these unforeseen circumstances.

I was determined to make the most of this unexpected turn of events, so embraced the opportunity to explore new research directions that would enrich my thesis and showcase my adaptability as a doctoral researcher. Although my background is in Environmental Science, COVID has presented the opportunity for me to invest substantial time and effort into learning many engineering concepts from scratch, and to apply them in my doctoral project.

The challenges presented by the COVID-19 pandemic have, without a doubt, tested my resilience and adaptability as a researcher. However, these circumstances have also served as a catalyst for personal and academic growth, pushing me to explore different domains and expand my skillset. I believe that the journey I have undertaken throughout the course of my PhD, in the face of such adversity, is a testament to my dedication and determination to produce novel, high-quality research that contributes to the advancement of knowledge in my field.

University of Southampton

Faculty of Environmental and Life Sciences

School of Ocean and Earth Science

**Design, Fabrication and Characterisation of a Low-Cost, Acoustically Focussed
Imaging Flow Cytometer for Automated Analysis of Phytoplankton**

by

Anthony James Willis Lindley

(ORCID: 0009-0005-8447-969X)

Thesis for the degree of Doctor of Philosophy

November 2023

University of Southampton

Abstract

Faculty of Environmental and Life Sciences

School of Ocean and Earth Science

Doctor of Philosophy

Design, Fabrication and Characterisation of a Low-Cost, Acoustically Focussed Imaging
Flow Cytometer for Automated Analysis of Phytoplankton

by

Anthony James Willis Lindley

Phytoplankton are a diverse group of organisms which are globally important from perspectives of ecology, environmental health, climate and socioeconomics, yet are severely understudied. The ocean is vast and largely inaccessible, and while many recent advances have taken place in the *in situ* measurement of physiochemical variables, analysing phytoplankton abundance and diversity is still a major challenge. Bulk approaches such as fluorimetry and satellite colorimetry, which each measure the fundamental properties of an entire phytoplankton population, can provide abundance and crude taxonomic data at a large spatial scale. Despite widespread use of these techniques, there remains a need for higher taxonomic and spatiotemporal resolution data that can only be provided by light microscopy and flow cytometry, two time-consuming and expensive methods.

To address these challenges, this thesis details the development of a novel, high-throughput, acoustically focussed Imaging Flow Cytometer for low-cost imaging of phytoplankton in natural water samples, making use of off-the-shelf optical and mechanical components. Acoustic focussing is used as it is a contact-free, gentle and reliable particle positioning method which allows high-throughput imaging of cells. Analytical and experimental testing of the acoustic focussing performance is detailed using Finite Element Modelling and imaging of polystyrene beads as a proxy for phytoplankton. A protocol for measurement of imaging resolution is developed and verified before being used to characterise the optical performance of the device.

In order to rapidly and automatically analyse the images captured by the device, various image processing techniques were investigated. In the finalised system, cutting-edge convolutional neural networks were designed, implemented, and verified by way of comparison with manual counting of plankton cells within images.

Finally, to demonstrate the effectiveness of the cytometer to address real research challenges, two experiments are described. In the first, the device automatically and successfully measures the density of preserved plankton cells within a test sample with an accuracy comparable to manual microscopy, the gold standard for this analysis. In the second experiment, the unique capability of the cytometer to generate high-temporal resolution measurements of live cells within growing cultures over an extended period was demonstrated. This experiment showed a discrepancy between the automatic measurements and manual verification, which is discussed at length, in the process uncovering a potential systemic bias occurring in phytoplankton research. The implications of these findings are explored.

Table of Contents

Table of Contents	i
Table of Tables	v
Table of Figures	ix
Research Thesis: Declaration of Authorship	xxvii
Acknowledgements	xxix
Abbreviations and Definitions	xxxii
Chapter 1 Introduction	1
1.1 Motivation	1
1.2 Research Objectives	2
1.3 Chapter Outline	2
1.3.1 Background (Chapter 2)	2
1.3.2 Design and Characterisation of a Novel Low-Cost IFC using Acoustophoretic Focussing (Chapter 3).....	3
1.3.3 Image Analysis for the Quantification of Phytoplankton in IFC Images (Chapter 4).....	3
1.3.4 Use of the Acoustically-Focussed IFC in Phytoplankton Experiments (Chapter 5).....	4
Chapter 2 Background	5
2.1 Phytoplankton and Other Aquatic Microorganisms	5
2.1.1 Ecological Role.....	5
2.1.2 Diversity.....	7
2.2 Harmful Algal Blooms (HABs)	9
2.2.1 Description	9
2.2.2 Impacts	9
2.3 Observed Trends and Future Behaviours.....	12
2.3.1 Climate Change and Ocean Acidification	12
2.3.2 Land use change and Eutrophication	12
2.4 Microalgal analysis techniques	14

Table of Contents

2.4.1	Background.....	14
2.4.2	Technologies.....	15
2.5	Imaging Flow Cytometry	18
2.5.1	Introduction.....	18
2.5.2	Current IFC devices used for phytoplankton analysis.....	19
2.5.3	Camera Technologies	24
2.5.4	Illumination Techniques.....	26
2.5.5	Focussing Techniques.....	28
2.6	Acoustofluidics	32
2.6.1	Introduction.....	32
2.6.2	Current Acoustically Focussed IFC Devices	33
Chapter 3	Design and Characterisation of a Novel Low-Cost IFC using Acoustophoretic Focussing	36
3.1	Introduction.....	36
3.2	Design Brief	37
3.3	Flow-cell implementation	39
3.3.1	Approach	39
3.3.2	Analytical Investigation of Flow Cell Acoustic Properties (1-D transfer model) 42	
3.3.3	Finite Element Modelling of Flow Cell Acoustic Focussing (COMSOL)	50
3.3.4	Acoustic Focus Analysis – Experimental.....	59
3.4	Imaging Performance Characterisation	74
3.4.1	Introduction.....	74
3.4.2	Modulation Transfer Function	76
3.4.3	Development of an experimental protocol for IFC resolution characterisation 78	
3.4.4	Experimental determination of the effect of camera exposure on measured system MTF	79
3.4.5	Experimental determination of the effect of edge slant rotation on measured system MTF	84

Table of Contents

3.4.6	Experimental determination of the effect of focus accuracy on measured system MTF	88
3.5	IFC System Design	95
3.5.1	Introduction.....	95
3.5.2	System Camera Selection	95
3.5.3	First system design approach.....	96
3.5.4	Improved system hardware	98
3.5.5	Imaging Performance of IFC System	102
3.6	Conclusion	106
Chapter 4	Image Analysis for the Quantification of Phytoplankton in IFC Images	109
4.1	Introduction.....	109
4.2	Traditional Image Processing	111
4.2.1	Introduction.....	111
4.2.2	Example Algorithm	113
4.2.3	Limitations.....	115
4.2.4	Limitations in IFC Imagery	118
4.3	Machine Learning-based Image Processing.....	122
4.3.1	Introduction.....	122
4.3.2	Neural Networks	134
4.3.3	Convolutional Neural Networks.....	145
4.4	Implementation of CNN Processing of Phytoplankton Images.....	150
4.4.1	Introduction.....	150
4.4.2	Hardware.....	151
4.4.3	CNN Investigation for Phytoplankton Detection	152
4.4.4	Evaluation of CNN Performance	162
4.4.5	YOLOX for Phytoplankton Detection.....	170
4.4.6	YOLOX Training.....	176
4.4.7	Evaluation of YOLOX Performance.....	179
4.5	Conclusion	185

Chapter 5	Use of the Acoustically-Focussed IFC in Phytoplankton Experiments	187
5.1	Introduction.....	187
5.2	Experimental Validation of Complete System	189
5.2.1	Experimental Objectives	189
5.2.2	Materials and Methods	189
5.2.3	Results and Discussion	190
5.3	Use of Acoustically-Focussed IFC in a Live Phytoplankton Experiment.....	193
5.3.1	Introduction.....	193
5.3.2	Materials and Methods	197
5.3.3	Results	200
5.3.4	Discussion.....	207
Chapter 6	Conclusion & Future Work.....	215
6.1	Conclusion	215
6.2	Future work	219
6.3	Summary of Novel Contributions.....	221
Appendix A	Parameters for KLM Model of Layered Resonator Flow Cell.....	223
Appendix B	COMSOL Finite Element Model Parameters and Variables	225
Appendix C	Pseudocode for Bead Detection and Separation Measurement in Double-Exposure Images	227
Appendix D	Bill of Materials for IFC	228
Appendix E	Growth Rate Experiment: Manual and Automated Sample Times.	231
Appendix F	Calculated Specific Growth Rates from IFC Counts of R. salina During Growth Rate Experiment.....	233
Bibliography	235

Table of Tables

Table 1 – Recorded socioeconomic damages resulting from Harmful Algal Blooms, in no particular order.	11
Table 2 - Table of commercial imaging flow cytometers used in biological oceanography for the purpose of plankton observation. Reproduction of data produced in Lombard et al.. ESD: equivalent spherical diameter	20
Table 3 – KLM-modelled acoustic pressure node depth within the flow cell for various driving frequencies representing a 20 kHz sweep with centre frequencies of the two determined resonant frequencies (bold) for the flow cell geometry - 1.776 MHz and 2.056 MHz.....	49
Table 4 - Comparison of hardware features and cost of two NVIDIA devices in the Jetson family of embedded computers designed for machine learning:	152
Table 5 – Full CNN architecture used to localise and count phytoplankton cells within images from the acoustically-focussed IFC. The network consists of 5 convolutional layers, one deconvolutional layer and one max pooling layer. The input is 3-channel (full-colour RGB) and the output is a single channel probability density map where each pixel value represents the confidence of the network of the presence of a phytoplankton cell.	153
Table 6 – Table of descriptive statistics of the cell counts from 50 IFC images of <i>R. salina</i> cells, with cells counted either manually or by the CNN presented within this section.	163
Table 7 – Table of descriptive statistics of the cell counts from 50 additional IFC images of <i>R. salina</i> cells, with cells counted either manually or by the CNN presented within this section.	165
Table 8 – Pixel values in the probability density map for each of the two sets of 50 images, split by whether the pixels are contained within a manually labelled bounding box for a cell or are in the background. Confidence values for cell detections should be much higher than those in the background, which they are, at approximately 10 times higher. However, if the model was very effective, confidence values would be expected to be far higher than presented, at around 5%. Furthermore, the mean confidence of detections in Image set 1 are different than those in image	

Table of Tables

set 2, which indicates inconsistent performance of the model across different input image conditions.	169
Table 9 – Comparison of the parameter complexity of the YOLOX-s and YOLOX-nano object detection models and the framerates at which they can process 512x512 images. Adapted from Nguyen et al. (2022).....	176
Table 10 - Table of descriptive statistics of the cell counts from 50 IFC images of <i>R. salina</i> cells, with cells counted either manually or by the YOLOX-s object detection network presented within this section. For the purposes of comparison, the results of the same analysis using the CNN investigated in section 4.4.4 are also included.	179
Table 11 - Table of descriptive statistics of the cell counts from the second set of 50 IFC images of <i>R. salina</i> cells, with cells counted either manually or by the YOLOX-s object detection network presented within this section. For the purposes of comparison, the results of the same analysis using the CNN investigated in section 4.4.4 are also included.	181
Table 12 – Measured specific growth rate (μ) of cultures of <i>R. salina</i> determined using the cell count measured with manual microscopy after pipette sampling.	203
Table 13 – Table of parameters defining the transducer used within the KLM model of the acoustically-focussed flow cell.	223
Table 14 - Table of parameters defining the particles used within the KLM model of the acoustically-focussed flow cell.	224
Table 15 - Table of parameters defining the layers of the flow cell used within the KLM model of acoustic focussing.	224
Table 16 - Table of parameters defining the fluid channel inside the flow cell, used within the KLM model.	224
Table 17 - Table of parameters used in Finite Element Modelling of the acoustically-focussed flow cell.....	225
Table 18 - Table of variables used in Finite Element Modelling of the acoustically-focussed flow cell.....	226
Table 19 - Bill of Materials to construct the optical, imaging and computational hardware for the IFC.	228

Table of Tables

Table 20 - Table of sampling times of the *R. salina* cultures during the growth rate experiment (Section 5.3) which were preserved and manually counted using microscopy.231

Table 21 - Times at which *R. salina* samples were automatically pumped into the IFC for imaging. The IFC was restarted on 26th March at 15:00 due to a connectivity issue, which resulted in an additional sample being processed that day and a change in the hours of sampling on subsequent days.232

Table 22 – Specific Growth Rate of each of the 4 cultures, calculated using the cell counts measured by IFC, using only those measurements taken at the same time in order to reduce the potential influence of diurnal changes in vertical distribution. Bold=maximum growth rate.233

Table of Figures

Figure 1 – A selection of light microscopy images of phytoplankton sampled by Balzano et al. during the 2009 MALINA cruise in the Beaufort Sea.....	6
Figure 2 – Phylogenetic Diversity of Marine Microorganisms – From Collins et al. (Collins et al., 2014).....	8
Figure 3 - Principle of standard flow cytometric analysis of particles suspended in a fluid sample. Particles flow past a detector where they are interrogated by a laser. Detected scattering and wavelength changes of the laser signal can be used to infer particle size, approximate shape, fluorescence.	16
Figure 4 - (Adapted from Lombard et al) Graphical illustration of the size ranges (in ESD, equivalent spherical diameter) each of the IFC instruments discussed in their review is capable of imaging. Note that CytoSense/CytoBuoy, while here presented as having the largest range, is actually a combination standard and imaging flow cytometer, and that imaging alone is likely to have a very similar measurement window to the IFCB (the actual imaging size range is unpublished). The red bars indicate the possible analytical size range which has been cross-validated with other established techniques as opposed to merely those published by the device manufacturers (grey dotted bars).	22
Figure 5 - Raw holograms and reconstructed phytoplankton images generated by the prototype IFC developed by Göröcs et al. (2018). Each hologram must be processed into a RGB image via a computational process called deconvolution, which in this work was performed by a Deep Neural Network.....	24
Figure 6 – CCD and CMOS sensor architectures. While CCDs, which are more common, use a single ADC for the entire pixel array, CMOS sensors, which have seen rapid technological improvements, have more complex pixel topologies due to every pixel having its own ADC.....	25
Figure 7 – ImageStreamX Mk II optical path showing fluorescence laser and wideband red brightfield illuminator. A complex custom detector simultaneously detects different fluorescent channels and brightfield images, made possible by a spectral decomposition of the signal into different colour bands.	26

Table of Figures

- Figure 8 - FlowCam8000/8400 IFC optical setup. One laser and one flash lamp (dotted box) illuminate the detection region and two PMTs provide wavelength-configurable fluorescence detection while a CCD images in brightfield.27
- Figure 9 - Göröcs et al. low-cost holographic IFC optical setup, showing the miniturised 3-colour illumination COB (inset). Though a holographic imaging device, this system has a similar purpose to the IFC proposed in this thesis.28
- Figure 10 – Cross-sectional illustration of the hydrodynamic focussing principle. Injected cells are focussed into a single-file stream by laminar flowing focussing fluid.28
- Figure 11 – 2-dimensional (a) and 3-dimensional (b) hydrodynamic focussing. In these simulations a red buffer fluid controls the flow profile of a blue carrier fluid. Adapted from Hamilton et al.29
- Figure 12 – Schematic adapted from Di Carlo et al. illustrating the principle of inertial focussing of particles within a microfluidic serpentine channel.29
- Figure 13 – (Reproduced from Antfolk and Laurell 2019) The four most common applications of acoustofluidic devices.....31
- Figure 14 - Block diagram for proposed acoustically-focussed imaging flow cytometer. The device itself will consist of an acoustophoretic focussing flow cell with necessary acoustic driving and control hardware, microfluidic pumping, optical components (camera, lenses and illumination). Once images have been acquired, analysis in the form of measurement and/or identification can be performed in real-time or from storage.....37
- Figure 15 – The flow cell design. a.: schematic cross-section of the acoustic focussing region showing the thicknesses of each layer. b.: Lateral view of one of the constructed flow cells showing fluidic connectors, acrylic frame, microscope slide and transducer. c.: Axial view of the same device.40
- Figure 16 – Diagram of the mounted acoustophoretic flow cell comprising a 30x6 mm rectangular glass capillary bonded via glass pillars to a standard (75x25 mm) microscopy slide using CA adhesive. Fluid connectors are formed using heat-shrink tubing and ultrasonic excitation provided by a square PZT transducer bonded to one end of the capillary using epoxy (Epotek 301).41

Table of Figures

Figure 17 – KLM model proposed by Krimholtz et al. The transducer crystal is modelled as an electric circuit consisting a transmission line, transformer and passive components.....	43
Figure 18 – Example plots of modelled Transducer Conductance (a) and Acoustic Energy Density (b) for a transfer model with carrier fluid parameters $\rho = 1024.7 \text{ kgm}^{-3}$ and $c = 1549.7 \text{ ms}^{-1}$. Two resonance frequencies are clearly defined at 1.78 MHz and 2.06 MHz. Figure plotted using MATLAB script detailed by Hill et al.	45
Figure 19 – Modelled Acoustic Pressure Amplitude (P) plot through vertical cross-section of the device at one of the half-wave resonant frequencies. Within the fluid layer (dark blue) the pressure has a single node, with particles above or below this point being subjected to forces pushing them toward the minimum, which is at the centre-height of this fluid layer. Figure generated using MATLAB script detailed by Hill et al.	46
Figure 20 – Frequency (y-axis) and Acoustic Energy Density (colour) for the lower (a.) and upper (b.) resonant frequencies predicted by the transfer model at each combination of temperature (grouping variable) and salinity (x-axis) for the layered resonator flow cell. The graphs indicate rising resonant frequencies with both temperature and salinity at both resonant frequency A and B.	47
Figure 21 – Diagram of the Finite Element Model of the flow cell cross-section showing the progression of beads (coloured) along the device with time. At the particle inlet the beads are uniformly distributed, and they are focussed to the central pressure anti-node as they travel along the x-direction. Once particles have reached the outlet at the right side, their vertical separation distance, d, is recorded.	51
Figure 22 - Schematic of the axes of the fluid layer used for the FEM modelling within this section, showing the position objective lens relative to the fluid flow direction.	52
Figure 23 – 1D plot of the x-component (v_x) of the fluid flow velocity profile within the vertical axis of the flow channel as described by Planar Poiseuille Flow. The given flow velocities given were calculated at a volumetric flow rate of 1.5 ml/min, but the parabolic distribution of velocity with y coordinate remains the same for all flow rates. The flow velocity at the centre of the channel is the highest and the velocity toward the channel boundaries (0 and 390 μm - the top and bottom of the fluid layer) tends to zero.	54

Table of Figures

- Figure 24 – Plot of the acoustic radiation force against vertical position (y-coordinate) as determined by the ultrasonic standing wave in the flow cell. The pressure minimum (zero newtons) is located at the half-depth of the fluid layer (195 μm) and also tends to zero at the upper and lower boundaries. Since particles are never located exactly at the boundary (they are separated by at least a distance equal to their radius), even particles very close to the flow cell walls will experience some force toward the centre.55
- Figure 25 – Plot of the maximum volumetric flow rate (Y axis, ml/min) under which particles of a given radius (X axis, μm) focus within the 5cm flow cell analysed by FEM, grouped by acoustic contrast factor (colour). There is significant overlap between different data points, especially at 10 ml/min flow rate where particles of $\varphi \geq 0.06$ with large diameters all focus.57
- Figure 26 – *C60 network analyser* plot of admittance (1 over impedance) of the ultrasonic transducer within the 0.5 – 3 MHz range. Measurements were conducted over 4 different conditions; unfilled (Air - black), filled with deionized water (DI H₂O - green); filled with marine media and PolySciences Fluoresbrite fluorescent polystyrene-latex beads (blue); filled with Lugol's Iodine-fixed plankton sample (sample number N20 - red). Insets: peaks around 1.70MHz; 2.10 MHz demonstrating the slight differences in acoustic performance when the device is filled with media of different salinities (and hence sound velocities and densities).....61
- Figure 27 – Example double-exposure image of 10 μm beads taken with a 2 ms delay between illumination pulses. Each of the 6 beads present in the field of view is imaged twice in the single frame such that their velocities can be calculated from the distance they travel during the 5 ms delay.64
- Figure 28 – Example of detected beads in the same image as Figure 26. The MATLAB script detects beads using standard image processing techniques (code in Appendix A.3) and assigns a number to each detection. Detection 11 in this image is a false positive and can be ignored by the user in the next step; labelling the exposure pairs which result from the same physical bead.....65
- Figure 29 – After detection, the user enters the numbers of the exposures relating to the same physical bead (e.g. 1 and 2; 3 and 4). The MATLAB script measures the distance

Table of Figures

- between the centroid pixels of the two labelled beads and the angle (in degrees from vertical), stores this data and presents it as an overlay on the image. ...66
- Figure 30 - Results of bead-separation experiment for measuring acoustic focus quality in layered resonator flow cell with resonant frequency = 1.75 MHz, transducer voltage = 10 V peak-peak. Volumetric flow rate = 1.0 ml/min (a) and 2.0 ml/min (b).67
- Figure 31 - Results of bead-separation experiment for measuring acoustic focus quality in layered resonator flow cell with resonant frequency = 1.75 MHz, transducer voltage = 15 V peak-peak. Volumetric flow rate = 0.5 ml/min (a), 1.0 ml/min (b) 2.0 ml/min (c) and 3.0 ml/min (d).68
- Figure 32 – Qualitative assessment of the acoustic focus quality at volumetric flow rates of 1.0 ml/min and 3.0 ml/min, with the same transducer settings of $f=1.75$ MHz, $V=15$ V peak-peak.69
- Figure 33 - Results of bead-separation experiment for measuring acoustic focus quality in layered resonator flow cell with resonant frequency = 2.11 MHz, transducer voltage = 10 V peak-peak. Volumetric flow rate = 0.5 ml/min (a), 1.0 ml/min (b) and 2.0 ml/min (c)71
- Figure 34 - Results of bead-separation experiment for measuring acoustic focus quality in layered resonator flow cell with resonant frequency = 2.11 MHz, transducer voltage = 15 V peak-peak. Volumetric flow rate = 2.0 ml/min.72
- Figure 35 – Process for determining resolution using the method inspired by Vainrub (2008). 13a) a test target with high-frequency content (made by scratching an ink blot on a standard 75x25mm slide with a razor) is imaged using the system under test. 13b) a 2-dimensional Fast Fourier Transform (FFT) is applied to the image to convert from spatial to frequency domains. 13c) a predetermined number of radial samples are taken from the centre of the FFT image outward to the edges, where all pixels in the given radius are sampled (as pixels have integer addresses, bilinear interpolation is used). 13d) the values corresponding to each radius are averaged and plotted.76
- Figure 36 – Schematic of the procedure for determining the resolution of a digital imaging device by calculating the Modulation Transfer Function (MTF) using the Slanted Edge Method (ISO 12233). Values are sampled from pixels perpendicularly to an imaged slanted edge, making up the Edge Spread Function (ESF). The ESF is

Table of Figures

differentiated to produce a Line Spread Function (LSF). Finally, the MTF is yielded from the 2-dimensional Fourier transform of the LSF. Adapted from Estriebeau and Magnan (2003) and Vhengani, Griffith and Lysko (2012).77

Figure 37 – MTF curves derived from images of a slanted edge on a resolution test target (R1L3S5P, Thorlabs Inc., USA) under camera exposure conditions from 1 ms to 15 ms. All images were taken under the same Z-axis focus position and with the same edge slant angle (5 °). Here, the blue curves (MTF from images taken with 1 ms – 7 ms exposure duration), which are highly consistent, are generated from images without any oversaturated pixels. The red curves (MTF from images taken with 8 ms – 15 ms exposure duration) are from images with pixel values of 0 or 4095 (i.e. fully saturated). These are erroneously high and abnormally smooth, indicating the importance of using a nonsaturating exposure time when measuring and comparing system resolving power.80

Figure 38 - Cropped regions from images taken at a)5ms exposure and b)10ms exposure in the 3.2.1 Exposure experiment. At 5 ms exposure, pixels in the white side are unsaturated and still contain information; at 10 ms exposure, the white side is overexposed and the camera is limited by the full-well capacity of its pixels. All pixels in the white section of this image have pixel intensity values of 4095, the limit of the camera. MTF curves generated from data with such saturated pixels are erroneously high and inflate the measured system resolution to unrealistic levels.81

Figure 39 – Graph of MTF10 (blue) and MTF50 (orange) values extracted from the MTF curves for images with exposure times ranging from 1 ms to 15 ms. Both metrics are consistent for exposure times of 1 to 7 ms, where images do not have any over- or undersaturated pixels. Exposures of 8ms and longer (red dashed box) have erroneously elevated MTF10 and MTF50 metrics than at faster exposures, indicating the inflating effect of saturated pixels on measured system resolution. MTF10 values which are above the sampling limit imposed by the Nyquist theorem (a minimum of 2 pixels per cycle required to avoid aliasing) are unrealistic.83

Figure 40 – Slanted edge MTF curves generated for incrementally increasing edge slant angles, starting from 0° (completely vertical) and incrementing by 5° up to 45°.85

Table of Figures

Figure 41 - Graph of MTF10 (dark blue) and MTF50 (light blue) values extracted from the MTF curves calculated for images with edge slant angles between 0° (completely vertical; aligned with sensor pixel columns) and 45°, in 2.5° increments. With edge rotation angles of 0-10° the determined MTF10 and MTF50 values are very consistent, but with increasing slant angle these resolution metrics display angle-dependence.87

Figure 42 – Comparison of the a)-e) central 256x256 pixels and f)-j) central 50x50 pixels of the slanted edge images taken during the section 3.2.3 focus experiment. Image a) is focused below the target and image e) above the target, as can be more clearly determined from the zoomed images f) and j). Images b), c) and d) are all subjectively ‘in focus’ and it is almost impossible to manually determine which has the most accurate focus position. As can be seen in Figure 42, there are subtle differences in derived resolution metrics between the three central in-focus images. Of particular interest is the fact that image e) has a darker white section than the others, which is reflected in the lower MTF metrics on the above-target focus positions seen in Figure 42. This is caused by the interaction of the Köhler illumination focus plane and the imaging focus plane of the objective.89

Figure 43 – Plot of slanted edge MTF10 (blue) and MTF50 (orange) metrics for increasing Z-axis focus positions. A focus point well below the target was selected and Z-height incrementally increased, with an image of the slanted edge taken every 1 micron step. Images were thus taken with the edge in focus (8-13 μm from initial position) and with the focus position above (0-7 μm from initial position) and below (14-20 μm from initial position) the target. The graph shows a nonlinear coupling between MTF10 and MTF50 values, highlighting the necessity of not solely relying on one or the other for resolution characterisation. Further, the plot demonstrates that Z-position accuracy of at least ± 2 μm is required to obtain high and consistent MTF measurements.91

Figure 44 – Diagram of the first iteration IFC mechanical arrangement. Using a stepper motor, 6 mm rods and linear bearings used in 3D printers, a carriage containing a camera can be raised or lowered as required to achieve focus of the cells passing through a flow cell below a standard 10X objective lens. By using a matt black Perspex enclosure, light from outside the device is limited and no tube between the camera and objective is required. The illumination frame can be lowered and

Table of Figures

an arbitrary number of optical components inserted above to create desired lighting conditions (e.g. Kohler).....	97
Figure 45 - 3D-printable flow cell holder for the prototype IFC. The acoustofluidic flow cell slots (from the right) into the cut-outs and is pushed until it reaches the end (left). The circular hole in the centre allows an unobstructed optical path for imaging and illumination, and standoffs in each corner, through which the cage rods are inserted, prevent flexion. Inlay: close-up of the flow cell inserted into the holder (blue).....	99
Figure 46 – Current IFC prototype arrangement. From the left: illumination from a pulsed LED is focussed via means of 2 iris and 2 lenses. The acoustofluidic flow cell slots into a 3D-printed flow cell holder (blue), which is sandwiched between metal cage plates for rigidity. The objective lens is mounted on coarse and fine focus apparatus, which themselves are connected to a camera via a SM1 tube...	100
Figure 47 – MTF curve for the complete IFC system. The spatial frequency at which the MTF drops to 50% (MTF50) is 0.167 cycles per pixel, and the spatial frequency at which the MTF drops to 10% (MTF10) is 0.264.....	103
Figure 48 – A single cell of <i>Rhodomonas salina</i> (fixed with Lugol’s Iodine solution 1%) imaged by the IFC system, demonstrating the ability of the system to partially resolve internal cell structures significantly smaller than 5 μm	104
Figure 49 - Example image created for demonstrating the traditional methods of object detection. The image contains ellipses of two different colours, sizes and eccentricities overlaid on a plain black background. The image is a 400x300 .jpg file.....	112
Figure 50 – The ellipses image after being converted to grayscale.....	113
Figure 51 – The grayscale image of ellipses after applying a 5x5 pixel gaussian blur operator.	113
Figure 52 – The images of ellipses after the Canny edge detection algorithm has been applied to the blurred, grayscale image.	114
Figure 53 – The image after grayscaleing, blurring, edge detection and thresholding. All that remains now is pixel values of 0 in regions which are not included in the ellipses, and values of 1 inside the ellipses.	114

Table of Figures

- Figure 54 – The ellipses after contour detection using the active contour algorithm. The outside edges of the objects of interest are now known and represented here in green.
.....115
- Figure 55 – The centre coordinates of each detected ellipse have been determined using the minimum and maximum extents of the previously calculated contours of the shapes, and are now plotted on the image.115
- Figure 56 – A new image of ellipses is presented in Figure 55a. The image now contains 3 ellipses of colours having a low contrast with the background. Figure 55b demonstrates the problem with the previously discussed image analysis approach; the contour detection has only operated successfully on the ellipse with the highest contrast from the background as the thresholding step used after edge detection had a threshold value set too high to include the darker objects. As a result, only that ellipse has its centre coordinates calculated, as demonstrated in Figure 55c.116
- Figure 57 – in this new image example, 2 of the 3 ellipses are partially overlapping. Though to a human, the 3 ellipses are clearly distinct objects, the algorithm which has been developed so far fails to distinguish the two objects and instead treats them as one. Figure 56b shows that the contours are calculated for the combined ellipses and therefore the centroid of the combined objects is plotted on Figure 56c.117
- Figure 58 – In this final example, Figure 57a. shows 3 ellipses which are overlaid this time on a gradient background, rather than the plain black of the previous examples. Though the 3 ellipses are still easily recognised by a human, the algorithm is unable to cope with the low contrast that the gradient produces and so falsely calculates coordinates for two of the ellipses and part of the gradient as a combined object, as shown in Figure 57b. Therefore, as Figure 57c, demonstrates, only one of the 3 ellipses is correctly identified.118
- Figure 59 – Two images of the IFC flow cell while filled with sterile L1 medium but no phytoplankton cells. Figure 58a. shows the flow cell after processing several phytoplankton samples and has obvious shadows which are cast onto the field of view by detritus adhered to the sides of the flow cell. Figure 58b. shows the same flow cell after a further week of regular sampling operation. More debris has adhered to the flow cell, so the background complexity has been significantly increased. The arrow points to an area of particularly heavy new

Table of Figures

adhered debris. This variable background poses a challenge to image processing algorithms.....	119
Figure 60 – IFC image of 3 phytoplankton cells, 2 of which are in contact with each other (arrow). As previously discussed, overlapping and touching objects pose a challenge to image processing techniques.	120
Figure 61 – close-up view of a phytoplankton cell (arrow) traversing in front of adhered detritus, the shadow of which is cast from the edge of the flow cell into the field of view.	120
Figure 62 - An image of a dog in a field of bluebells is used to demonstrate three kind of image classification. In Binary Classification, a model which has been trained to recognise images of dogs outputs a 0-1 probability of the image containing a dog. In Multiclass Classification, the model is trained on multiple classes (Dog, Cat, Bus, Plant) and for each class outputs a probability of the image containing an example of that class. The class with the highest probability score, in this case Dog, is selected as the output. In Multilabel Classification, the model can assign more than one label to the image. In this example, the model outputs both Dog and Plant as both classes have probability scores which pass some threshold value. Adapted from (MathWorks, 2023).	124
Figure 63 - Two images from the acoustophoretic IFC device are presented. For each of the two input images, a binary classification model which has been trained to detect cells could output a probability score of the image containing cells. Based on this score, frames which do not contain cells, like the first input image, could be discarded without being saved, saving memory and processing time.....	125
Figure 64 – An image of a single phytoplankton cell is used as an input for hypothetical binary and multiclass classifiers. The binary classifier can predict whether a cell is healthy or not, whereas a multiclass classifier could be trained to distinguish between several genus of phytoplankton and output a probability score of the input image being an example of each, selecting the highest probability.	126
Figure 65 – An object detection algorithm based on a binary classification of each individual pixel within an input image. The algorithm generates a probability score (0-1) of the pixel containing a cell, after which a number of algorithms can be used to, for instance, count, localise or measure cells within an input image.....	127

Table of Figures

- Figure 66 – Example of a typical input and output from a machine learning regression model for object detection. An IFC image of a single phytoplankton cell is input, and the model predicts values for the coordinates of the corners which could be used to draw a box around the cell(s) in the input image.....128
- Figure 67 – This diagram, adapted from Shin et al. (2016a), demonstrates the ability of a machine learning model to output probability density maps which correspond to predicted locations of various classes of objects of interest. In this example, an abdominal MRI input image has 4 corresponding output density maps, one for each of the organs of interest.129
- Figure 68 – Example of object detection via machine learning, where coordinates defining bounding-boxes are predicted via a regression model. In this example, the model has been trained to detect 3 different classes, and can simultaneously output a list (b.) of bounding boxes for an input image (a.) with a predicted class for each box (final column of the output list b.). Finally, the boxes can be drawn onto the input image to visually show the predictions of the model (c.).130
- Figure 69 – Adapted from <https://manipulation.csail.mit.edu/segmentation.html>. An image of a dog and three sheep is used as an example of image classification, object detection, semantic segmentation and instance segmentation. In semantic segmentation, each pixel is assigned to the class with the highest probability. In the given example, red is used for the ‘dog’ class, and blue is used for the ‘sheep’ class. Pixels in the background are not separately coloured in this case, but would have been assigned the ‘background’ class. There is no distinction made between different instances of the same class, unlike in instance segmentation, which is shown below. In instance segmentation, each pixel is assigned a class as before, but also assigned an instance variable representing the number of that object. In this manner, the image pixels belonging to the three different sheep are separately labelled and can be extracted or analysed individually.....131
- Figure 70 – two images which were generated by the DALL-E 2 image generation model. In each instance, a prompt, which is displayed here above the image, was used to generate a new image using a Generative Adversarial Network. The model is trained on many examples of images of different classes and semantic information, and can then use the learned features to produce new images based on input text. Here, an image of a dog sitting in a field of bluebells is

Table of Figures

- created, as is an image of a cell which could be from a microscope image.
Neither of these images depict real, extant objects.132
- Figure 71 – Two variants of an input image generated by the generative adversarial network model DALL·E 2. Neither of the two variants are real images but instead were generated using a combination of previously learned features and features from the input image.....133
- Figure 72 – Illustration of the operation of one neuron. The neuron computes a weighted sum of its inputs, $[a_1, a_2, a_3, \dots, a_N]$, using associated weight values, $[w_1, w_2, w_3, \dots, w_N]$. A bias value, b , is added and the output, z , is passed on to the next connection.....135
- Figure 73 – An illustration of an extremely simple neural network. Two input neurons (blue) are connected to 4 neurons (grey) within a single hidden layer, which are in turn connected to a single output neuron (orange). Each connection represents one of the inputs to a matrix operation.136
- Figure 74 – The sigmoid and Rectified Linear Unit (ReLU) functions. By passing the output value of a neuron through these functions, the neural network is no longer performing a simple linear regression and is able to represent complex nonlinear relationships between input and output data. The specific activation function used influences the performance of a given model.137
- Figure 75 - A deep neural network with 3 hidden layers is presented. In this network, each layer is fully-connected; that is, each neuron is connected to every neuron in the prior layer. By introducing additional hidden layers of neurons, the deep neural network is able to model increasingly complex relationships between input and output data. (Strauß, 2018)138
- Figure 76 – Intersection over Union (IoU) is a performance metric which calculates the accuracy of a bounding box prediction by dividing the intersection (overlap) by the union (area enclosed by both boxes), for the predicted and the ground truth bounding boxes.....144
- Figure 77 – A visual representation of a convolution operation on an image represented as a 2D matrix (left). Convolution computes the value of each destination pixel in an output matrix (right) by multiplying each value in the convolution filter (centre) by the corresponding value in the input image, and then sums the results. This

Table of Figures

- process is repeated for every pixel in the input image, resulting in a new output image where each pixel is the result of applying the convolution filter to the corresponding pixel in the input image. Within a Convolutional Neural network, the inputs to a convolution, the outputs from the operation and the filter itself can be of any number of dimensions, with CNNs having many layers potentially having hundreds of dimensions within some layers.146
- Figure 78 – An illustration of a 2x2 Max Pooling operation on a 4x4 matrix. The original matrix is divided into 2x2 regions and only the maximum value from each region is retained in the output matrix. Therefore, the output is of shape 2x2, and the input has been downsampled.147
- Figure 79 – Example of an input image to the network (top): a 1920x1200 pixel, colour image of the IFC flow cell with 3 phytoplankton cells (a close up of 1 is provided on the right). On the bottom is the probability density map output by the CNN, with each pixel having a value which represents the confidence of the CNN that there is a cell there.156
- Figure 80 - Evolution of training (orange) and validation (blue) loss (Mean Squared Error, MSE) across 150 epochs for an object detection CNN, showing the prompt minimization of training loss, a transient surge in validation loss around epoch 8, and subsequent convergence to a value close to 0.157
- Figure 81 – Input image (left), probability density map representing model confidence of the presence of a cell (middle) and binarized object mask after thresholding the density map (right).160
- Figure 82 – Input IFC image (a.) and the same image overlaid with a binarized output mask of the CNN-predicted cell locations (b.), illustrating how the 3 cells visible can be easily localised based on the probability density map output by the model.161
- Figure 83 - Bland-Altman plot comparing cell counts from the machine learning (CNN) method and manual counting of cells within Image Set 1. The mean difference is -2.34, indicating a lower mean count by the CNN method. The upper and lower limits of agreement are 0.56 and -5.2, respectively, illustrating the range within which 95% of the differences between the two methods are expected to lie.164
- Figure 84 - Bland-Altman plot comparing cell counts from the machine learning (CNN) method and manual counting of cells within Image Set 2. The mean difference is 3.18,

Table of Figures

indicating a significantly higher mean count by the CNN method. The upper and lower limits of agreement are 6.3 and 0.1, respectively, illustrating the range within which 95% of the differences between the two methods are expected to lie.	166
Figure 85 – Input image after resizing (top) and masked CNN output probability density maps. The different masking threshold values selected (0.01 and 0.05) produce outputs with significantly different numbers of cell detections, with a threshold of 0.01 resulting in many false positives.....	167
Figure 86 - Top-level overview of the behaviour of a YOLO object detection model when detecting objects within an image. The image is divided into a grid, then the network predicts bounding boxes and confidence scores for potential objects within each grid cell, along with their associated class probability maps. These predictions and class probability maps are combined to create a detection result, which identifies the locations and classes of objects within the image. (Wu and Zhou, 2019).....	171
Figure 87 - The Darknet53 architecture, which consists of 53 convolutional layers and is used as the backbone for YOLOv3, a modern object detection algorithm. Diagram from (Redmon and Farhadi, 2018).	172
Figure 88 - The CSPDarknet53 neural network architecture, first introduced in YOLOv4. Used as the backbone for advanced YOLO object detection models, CSPDarknet53 employs Cross-Stage Partial (CSP) connections, facilitating superior feature extraction capabilities and enhanced object detection performance across a range of scales in the input image. The schematic on the right illustrates the structure of a CSP block, which divides the input features into two streams: one stream passes through a sequence of convolutional layers, while the other bypasses these layers. Subsequently, both streams are merged, resulting in an efficient and effective combination of extracted features. Adapted from (Xu et al., 2021).	173
Figure 89 - The complete YOLOX network architecture, featuring inlaid diagrams of its core components: the Spatial Pyramid Pooling (SPP) module, the Neck (Path Aggregation Network, or PANet), and the YOLO Prediction Head. These components enable efficient feature extraction, hierarchical feature aggregation, and accurate bounding box predictions with class probabilities,	

Table of Figures

- ultimately resulting in cutting-edge object detection for objects of a wide range of scales. Adapted from (Chou, 2022)175
- Figure 90 – Graph of AP50 (Blue) and AP95 (Orange) for YOLOX-s object detection model while being trained to detect and localise cells of *Rhodomonas salina* phytoplankton within images from the acoustically focussed IFC. After 150 epochs of training the AP50 value is close to 1.0, indicating a good overlap between predicted and ground truth bounding boxes. The lower AP95 score suggests that the model does not predict bounding boxes with a very tight alignment to those which were manually labelled to produce the training dataset, the 95% IoU threshold requires a large degree of overlap between the predicted and ground truth bounding boxes. This is unlikely to be a cause for concern as the human-labelled bounding boxes were not highly precise to a pixel level, especially in comparison to the size of the cells themselves, which each take up a very small proportion of the image at 10X magnification.....178
- Figure 91 - Bland-Altman plot comparing cell counts from YOLOX and manual counting of cells within Image Set 1. The mean difference is 0.0, indicating a lack of systematic bias. The upper and lower limits of agreement are +1.3 and -1.3, respectively, illustrating the range within which 95% of the differences between the two methods are expected to lie. These values further demonstrate the absence of systematic bias.....180
- Figure 92 - Bland-Altman plot comparing cell counts from YOLOX and manual counting of phytoplankton cells within Image Set 2. The mean difference is 0.18, indicating a slight systematic bias where YOLOX detects more cells per image. The upper and lower limits of agreement are +1.5 and -1.1, respectively, illustrating the range within which 95% of the differences between the two methods are expected to lie. These values further demonstrate the slight over-counting of cells by YOLOX compared with ground-truth data.182
- Figure 93 – An illustrative example of a misclassification of detritus as a plankton cell by YOLOX. While the three cells present in the image are correctly identified with >70% confidence, the clump of detritus has been falsely labelled as a cell by the algorithm with 63.5% confidence.....183
- Figure 94 - Box and whisker plot of the cell counts generated by: IFC & YOLOX (left/blue), the fully automated system involving imaging using acoustically focussed IFC and

Table of Figures

processing using the YOLOX object detection model; and manual cell counts (right/orange) using a microscope. The deviations from the mean do not indicate inaccuracies in the counting method, but rather highlight the non-uniform distribution of cells in each imaged volume.....	191
Figure 95 – Layout of system used in the growth rate experiment to measure cell count using acoustically focussed IFC.	199
Figure 96 - Cell counts measured through manual microscopy of pipetted and preserved <i>R. salina</i> throughout the growth rate experiment.....	200
Figure 97 - Automatically measured cell counts based on IFC images of the 4 cultures of <i>R. salina</i> during the growth rate experiment. L1 is sterile medium, and acts a control.	205
Figure 98 – Combined graph of manual (dots) and automatic IFC (lines) cell counts measured throughout the growth rate experiment. Orange bars represent the times when the incubator’s lights were on, and black vertical lines are the times at which the samples were agitated and manual samples drawn off for verification. L1 is sterile medium, and acts a control.	206
Figure 99 - A single frame taken from the IFC during imaging of <i>R. salina</i> cells during the growth rate experiment, demonstrating successful acoustic focussing of all cells. ...	208

Research Thesis: Declaration of Authorship

Print name: Anthony Lindley

Title of thesis: Design, Fabrication and Characterisation of a Low-Cost, Acoustically Focussed Imaging Flow Cytometer for Automated Analysis of Phytoplankton

I declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. None of this work has been published before submission

Signature: Date:

Acknowledgements

Every day is a journey,

And the journey itself is home.

Bashō

With heartfelt thanks to all those who have supported me throughout the journey. First and foremost, my gratitude to Dr Peter Glynn-Jones, whose guidance and belief in my ability were instrumental in completing this thesis. Thanks also to Dr Kathryn Cook, whose assistance and unwarranted patience in the lab were invaluable in overcoming the numerous challenges I faced during my research. Thank you to Professor Doug Connelly for being a supportive chair to my PhD panel, and to Dr Dan Mayor for his continued enthusiasm about my project.

To Rachael, thank you for supporting me even when I didn't deserve it. Your patience and love kept me going through the toughest moments.

Finally, thank you to my friends and family for your constant encouragement. Your belief in me has been incredibly moving.

Abbreviations and Definitions

Adam	Adaptive Moment Estimation: A stochastic optimization algorithm that adaptively adjusts learning rates for each parameter based on estimates of first and second moments of the gradients.
AP	Average Precision: A summary statistic of precision-recall curves, reflecting the average precision value for recall values over the entire range.
BN.....	Batch Normalisation: A type of normalisation method applied in CNNs which normalises the inputs of each layer for each batch.
CCD.....	Charge Coupled Device: A sensor used in digital imaging to convert light into an electronic signal.
CMOS	Complementary Metal Oxide Semiconductor: A fabrication process used for integrated circuits including image sensors.
CNN	Convolutional Neural Network: A class of machine learning model: a neural network that uses convolutional layers to perform a mathematical operation on input data, typically images.
CPR.....	Continuous Plankton Recorder: A device towed behind ships to collect plankton samples.
CSP	Cross-Stage Partial: A neural network design approach that separates the a layer's output into two stages and designed to improve the network's efficiency.
DVM	Diel Vertical Migration: A behaviour observed in many aquatic organisms involving movements up and down in the water column over a 24-hour cycle.
EOV.....	Essential Ocean Variable: Fundamental physical, chemical, and biological variables used in monitoring, modelling, and managing the ocean.
ESD	Equivalent Spherical Diameter: The diameter of a sphere that has the same volume as a given particle.
ESF.....	Edge Spread Function: The spatial response of an imaging system to an edge.

Abbreviations and Definitions

FC	Flow Cytometry: A technique used to measure the properties of particles in a fluid as they pass through a laser.
FCRN.....	Fully Convolutional Regression Network: A type of CNN which uses convolutional layers for both feature extraction and regression.
FEM	Finite Element Model: A computational method for finding approximate solutions to partial differential equations representing physical systems.
FITC.....	Fluorescein Isothiocyanate: A fluorochrome used in flow cytometry to label antibodies.
FN	False Negative: An outcome where the model incorrectly predicts the negative class.
FP.....	False Positive: An outcome where the model incorrectly predicts the positive class.
FPN	Feature Pyramid Networks: A type of CNN designed to efficiently build high-level feature maps at different scales, improving the detection of objects at varying scales in image-based tasks.
GOOS.....	Global Ocean Observing System: The international collaboration for sustained observations of the oceans.
GPU	Graphics Processing Unit: Computer hardware created for performing parallel computations on large data sets. Commonly used in graphics, simulations and deep learning.
HAB	Harmful Algal Bloom: A rapid increase in the population of algae in an aquatic system, which is harmful due to the production of toxins and/or oxygen depletion.
IFC	Imaging Flow Cytometer: An instrument combining flow cytometry and microscopy to analyse particles in flow.
IFCB	Imaging Flow CytoBot: An automated, submersible Imaging Flow Cytometer for analysing marine microorganisms.
IoU.....	Intersection Over Union: A metric for the overlap between two bounding boxes or other areas, calculated as the area of intersection divided by the area of union.

Abbreviations and Definitions

KLM	Krimholtz, Leedom and Matthaei: A model, named after its authors, that predicts the electrical properties of piezoelectric transducers.
LED	Light Emitting Diode: A semiconductor device that emits light when an electric current is passed through it.
LRN	Local Response Normalisation: A type of normalisation method applied in CNNs which normalises the activities of neurons based on the activity of adjacent neurons.
LSF	Line Spread Function: The first derivative of the ESF.
MAE.....	Mean Absolute Error: A measure of prediction error calculated as the average of the absolute differences between predicted and actual values.
MSE	Mean Squared Error: A measure of prediction error calculated as the average of the squares of the differences between predicted and actual values.
MTF	Modulation Transfer Function: The spatial frequency response of an imaging system; measures the ability of the system to reproduce the contrast of an object at different spatial frequencies.
NPP.....	Net Primary Productivity: The residual of photosynthetic carbon fixation minus respiration,
OD	Optical Density: The measure of light absorbance in a material.
PAN	Path Aggregation Network: A type of neural network architecture used in object detection which improves flow of data between different layers.
PMT	Photomultiplier Tube: A sensitive and fast photodetector.
PPFD.....	Photosynthetic Photon Flux Density: The number of photosynthetically active photons incident per unit area per unit time.
PSU	Practical Salinity Units: A unit of measurement for the salinity of water based on its electrical conductivity.
PZT.....	Lead Zirconate Titanate ($\text{Pb}[\text{Zr}(x)\text{Ti}(1-x)]\text{O}_3$): A type of ceramic that exhibits piezoelectric properties.

Abbreviations and Definitions

- ReLU Rectified Linear Unit: A function used in neural networks that outputs the input directly if positive, otherwise, outputs zero.
- SPP Spatial Pyramid Pooling: A layer used in CNNs which generates fixed-size outputs regardless of input size, by applying pooling operations at various scales.
- TN True Negative: An outcome where the model correctly predicts the negative class.
- TP True Positive: An outcome where the model correctly predicts the positive class.
- TRL..... Technology Readiness Level: A scale used to assess the maturity level of a particular technology.
- WFD..... Water Framework Directive: The European Union directive that sets a framework for protection and management of water bodies.
- YOLO..... You Only Look Once: A real-time object detection system that identifies objects in images in a single pass.

Chapter 1 Introduction

1.1 Motivation

Understanding the processes, pressures and changes acting upon the oceans requires sustained, global observation of physical, chemical and biological variables. Global quantification of ocean state has been the ultimate aim of observational oceanography from its inception and has resulted in the establishment of an intergovernmental panel on ocean observation, the Global Ocean Observing System (GOOS). In the past decade a clear picture of what measurements are required for adequate assessment of ocean state has emerged. These measurements have been described as the *Essential Ocean Variables* (EOVs) and fall into *physical, biogeochemical and biological/ecosystems* categories.

Of these, the biological/ecosystems EOVs face particular challenges with respect to the collection, processing and standardisation of measurements. Accurate quantification of phytoplankton biomass and biodiversity is of critical importance to biological oceanography yet is underdeveloped due to technical limitations which are explored in Chapter 2. The phytoplankton EOV is important due to the reliance of higher trophic levels on phytoplankton quality and abundance; the essential role of phytoplankton in carbon cycling and oxygen production; the effects of harmful algal blooms and the reflection of water quality in phytoplankton community structure.

Phytoplankton diversity has been studied for over a century, with modern observations still relying on similar procedures to the earliest investigations. At present, studies into taxonomic distribution usually rely on the collection of cells using drag nets and manual microscopic examination in order to determine species or lower order classification. This process is labour-intensive, time-consuming, unreliable and requires the decoupling of spatial and temporal information. Autonomous devices exist and are capable of automatic identification of phytoplankton in flowing samples at high accuracies but are at presently too expensive and complex to allow widespread adoption – as will be discussed in the following sections, a typical autonomous in situ phytoplankton microscope can easily cost over £100,000, which when

Chapter 1

including training, staff time and related equipment costs could be too expensive for many research programmes.

The motivation of this research, therefore, is the design, fabrication, characterisation, and use of a novel, low cost and simple device for the autonomous imaging of phytoplankton.

1.2 Research Objectives

The primary objective of this thesis are to develop and use a low-cost, automatic imaging flow cytometer by using acoustophoretic levitation to position phytoplankton cells into a sheet within the focal plane of a camera, and thus image cells at high throughputs. The device will be developed at as low a cost as reasonably possible such that it is accessible to resource-constrained laboratories, in order to enable an increased use of imaging within global ocean observations. Included are chapters on the background, motivation and requirements of the device, the use of cutting-edge image processing algorithms to automatically analyse images and an investigation of the complete system applied to real-world phytoplankton research scenarios. The following sections briefly describe the content of each of the following chapters.

1.3 Chapter Outline

1.3.1 Background (Chapter 2)

The background chapter covers a literature review investigating the ecological, climatological and socioeconomic importance of phytoplankton on global and regional scales, the history of phytoplankton observation, the state of the art and next-generation observation techniques. Imaging Flow Cytometry is identified as a candidate next-generation technology holding great promise for increasing the frequency and richness of observation while reducing cost compared to traditional methods. Following a review of current and prototype/research Imaging Flow Cytometers (IFCs) used in phytoplankton observation and monitoring, the constituent components of IFCs are investigated. These include the imaging technologies (cameras and illumination) and particle focussing techniques (hydrodynamic, inertial and acoustophoretic).

Chapter 1

Acoustophoresis is shown to hold promise for low-cost novel IFCs and a summary of the literature of acoustic focussing is presented.

1.3.2 Design and Characterisation of a Novel Low-Cost IFC using Acoustophoretic Focussing (Chapter 3)

The engineering work detailed within this chapter is broken into the rationale and process for development of each subsection of the imaging flow cytometer, the design and fabrication directions chosen and the performance characterisation of the imaging and acoustic focussing. Transfer impedance modelling (KLM) and computational (FEM) simulations of the acoustic flow cell are detailed. A frequency-response experiment is carried out to determine the resonant frequencies of the fabricated flow cells, and a methodology for determining acoustic focussing performance from double-exposed images of polyester beads developed. The latter method proved unsuitable, and the reasons for this are explored.

A method for determining the resolution of the optical system (MTF) is developed and quantified, and experiments to measure MTF using a research microscope presented. The benefits and limitations of the MTF approach are discussed.

Finally, the complete system (acoustically-focussed flow cell, illumination apparatus, imaging optics and structural components) is designed, assembled and the image quality assessed using the previously discussed technique.

1.3.3 Image Analysis for the Quantification of Phytoplankton in IFC Images (Chapter 4)

After developing the acoustic, fluidic, optical and structural design of the IFC, an investigation into automatic image analysis for autonomously detecting and counting phytoplankton cells within images from the IFC was carried out. This chapter details the process of using traditional, feature-based image process algorithms before employing machine learning in the form of convolutional neural networks (CNNs). A simple CNN used in a similar IFC is used as a baseline, and is modified for the application within colour images of phytoplankton from the acoustically-

Chapter 1

focussed IFC. A comparison of the performance of this algorithm with that of manual cell-counting by a human was carried out and it was demonstrated that the automatic approach was unreliable. Finally, a cutting-edge object detection CNN was implemented and tested, and showed performance rivalling that of a human operator.

1.3.4 Use of the Acoustically-Focussed IFC in Phytoplankton Experiments (Chapter 5)

In this chapter two experiments to quantify the performance of the IFC in addressing real-world research problems are presented, showcasing its potential as a low-cost, high-throughput imaging method to detect and count phytoplankton cells within preserved and living samples. In the first experiment, the accuracy of the instrument in determining the abundance of preserved phytoplankton cells is demonstrated, in a comparison with manual microscopy, and it is determined that the novel IFC presented previously has the potential to reliably replace the labour-intensive process of manually counting cells with a higher throughput and at a fraction of the cost of commercially available IFC instruments.

Subsequently, the IFC system is applied to the study of live phytoplankton cells during their growth within an incubator, where it was found that that the swimming velocity of these motile cells is great enough to avoid their sampling into the fluidic system for processing. This finding has implications for many other phytoplankton research studies, which do not regularly consider the potential biases introduced by under-sampling motile species.

Chapter 2 Background

2.1 Phytoplankton and Other Aquatic Microorganisms

2.1.1 Ecological Role

The ocean is a critical component of the global carbon cycle, containing as much as 50X more CO₂ than the atmosphere (Legendre et al., 2001). The oceanic carbon cycle has three components; physical, chemical and biological, which together mediate fluxes of CO₂ to and from the ocean. The biological component of this system drives a globally significant net flux of CO₂ into the ocean and is almost exclusively driven by photosynthetic net primary production (NPP) (Falkowski et al., 1998).

Phytoplankton (Figure 1) are a diverse group of prokaryotic and eukaryotic photoautotrophic microorganisms consisting thousands of described species and form the biological basis of the oceanic carbon pump. Fossilised ancestors of contemporary cyanobacteria, the most abundant group of phytoplankton, have been dated at around 3.5 billion years (Ga) old (Schopf, 1993) and it is widely believed that oceanic photoautotrophy not only gave rise to the oxygen-rich atmosphere suitable for higher forms of life but are direct predecessors of terrestrial plants (Bhattacharya et al., 1998).




Figure 1 – A selection of light microscopy images of phytoplankton sampled by Balzano et al. during the 2009 MALINA cruise in the Beaufort Sea.

Across the ocean phytoplankton ‘fix’ approximately 45 gigatons of carbon per year, of which 16 gigatons sinks to the deep ocean where it is permanently stored in sediment (Falkowski et al., 1998).

In addition to acting as the primary mechanism drawing CO₂ from the atmosphere to the ocean, these primary producers act as the basis of almost all higher forms of marine life and are adapted to environments as diverse as arctic glaciers and subtropical gyres. As these unicellular microorganisms also have fast lifecycles, high morphological heterogeneity, can be found in all natural waters and respond quickly to environmental conditions, phytoplankton can also act a biomarker of marine ecosystem health (Suikkanen et al., 2007).

Phytoplankton are not restricted to marine environments; freshwater bodies including lakes, rivers and even potable water supplies also support large and complex phytoplankton assemblages. As well as providing the primary production required to support freshwater aquatic

Chapter 2

life, freshwater phytoplankton occupy environments which are particularly influenced by anthropogenic loadings such as nutrient enrichment and pollution. The response of the plankton community to these factors can be measured and used to infer the level of impact faced by particular water bodies.

The EU's Water Framework Directive (WFD) uses *the composition, abundance and biomass of phytoplankton* as a biological marker of freshwater ecosystem health (Schmidt et al., 2008); EU countries are therefore legally required to measure the number and taxonomic composition of microalgae, at regular intervals, in their lakes and rivers (Pasztaleniec and Poniewozik, 2010). The WFD uses a series of 'indicator species' which are counted and have corresponding 'scores'; the sum of these scores, combined with indices according to algal size spectra and total biomass, is then used to infer the relative health of the water body in question (Molina-Navarro et al., 2014, Katsiapi et al., 2011).

2.1.2 Diversity

Phytoplankton are so globally ubiquitous and abundant that they contribute up to 45% of total global photosynthesis despite representing only 2% of photosynthetic biomass (Thyssen et al., 2008). Representing 12 taxonomic divisions within 3 kingdoms (Cavalier-Smith, 1993), the phytoplankton group contains both pro- and eukaryotic organisms which can have autotrophic, auxotrophic and mixotrophic nutritional strategies but are united by their possession of chloroplasts.

With sizes ranging from one micron to several millimetres, the cellular volumes of phytoplankton span at least 6 orders of magnitude (Maranon, 2015, Cloern, 2018). The relevant size fractions to which phytoplankton belong are described (by cell width) as nano- (2.0-20 μm), micro- (20-200 μm) and mesoplankton (0.2-20 mm). Within these classifications, phytoplankton are joined by the protozooplankton, a group including colourless flagellates and ciliates (Sieburth et al., 1978). Though these species are phagotrophic (Schnepf and Kühn, 2000) or osmotrophic (Richards et al., 2012) and are thus not primary producers, protozooplankton nevertheless play an important role ecologically (Burkholder and Glasgow Jr., 1997, Burkholder et al., 2008) and in carbon cycling (Levinsen et al., 1999).

Chapter 2

The taxonomic diversity of phytoplankton is extensive, with great interspecies differences in cell size, cell-wall composition, sinking rates and impacts on the biological carbon pump (Collins et al., 2014). A summary of the major phylogenetic groups of phytoplankton is presented in Figure 2. Some of the most studied subdivisions include the Rhizaria (a supergroup containing Radiolarians); the Diatoms, Dinoflagellates, Coccolithophorids and Prasinophytes (chlorophytes, containing green algae). An alternative approach to the classification of marine microorganisms, which has particularly been developed since the development of metagenomics, is division into functional groups, which links plankton with similar ecological roles (e.g. nitrifiers – predominantly cyanobacteria; or calcifiers, such as some coccolithophores) (DeLong, 2009).




Figure 2 – Phylogenetic Diversity of Marine Microorganisms – From Collins et al. (Collins et al., 2014)

2.2 Harmful Algal Blooms (HABs)

2.2.1 Description

Phytoplankton periodically experience dramatic increases in biomass of a single or few species, in localised and transient events known as *algal blooms* which can be so large as to be visible from space (Behrenfeld, 2010, Niemi, 1973). These events can provide large amounts of particulate organic carbon to higher trophic levels as well as changing physical properties of the water body such as light penetration and dissolved nutrients. Polar and subpolar regions of the ocean are characterised by strong and predictable vernal blooms which are critically important ecologically, dominating net primary production throughout the year (Moore and Abbott, 2000).

The effects of algal blooms can be negative, reducing dissolved oxygen availability in the water column when they sink to depth and are decomposed by aerobic bacteria (Anderson et al., 2002). Of the over 100,000 planktonic species described, at least 300 are capable of causing 'red tides', where coastal regions experience severe deoxygenation resulting in widespread mortality of higher trophic levels (Smayda, 1997). In addition, blooms of certain species such as the diatom *Chaetoceros convolutes* can clog fish gills and cause asphyxiation directly (Hallegraeff, 1993). At least 70 species are further capable of producing biotoxins (Smayda, 1997), such as *domoic acid* (C₁₅H₂₁NO₆) produced by *Pseudo-nitzschia australis* (Lim et al., 2012). When blooms occur which are toxic or otherwise indirectly harmful to ecosystems or human health, the events are termed *harmful algal blooms* (HABs) (Anderson et al., 2002). HABs are not limited to marine and coastal environments; freshwater species are capable of causing HABs posing particular risks due to their proximity to humans, drinking water supplies, recreational spaces etc.

2.2.2 Impacts

HABs can have severe economic effects when occurring in or around fisheries: most first-world countries mandate fishery closure when a certain level of cells from harmful species are detected in water samples (or toxic products are detected above a regulatory threshold). The biotoxins produced can undergo bioaccumulation/biomagnification in plankton-grazing shellfish, reaching levels in individuals high enough to cause human mortality if ingested (Prakash et al., 1971, Carmichael, 2001).

Chapter 2

The causative factors leading to the development of the blooms of toxin-producing species are complex and it is not yet possible to predict HAB occurrence based solely on environmental measurements. This is a problem as physiochemical variables are far easier to obtain than biological factors, as elaborated on in Section 2.4.

The human illnesses caused by ingestion of bioaccumulated products of harmful algae are paralytic-, diarrhetic-, neurotoxic-, amnesic- and azaspiracid shellfish poisoning. Except for amnesic shellfish poisoning, which is caused by toxins produced by diatoms, all HAB-related poisons are produced by dinoflagellates. A selection of the socioeconomic impacts caused by HABs are summarized in the literature in Table 1. Even in purely economic terms HABs are clearly a major source of concern – with costs ranging in the tens of millions of pounds per year it makes sense to invest considerably into monitoring and understanding harmful phytoplankton.

Table 1 – Recorded socioeconomic damages resulting from Harmful Algal Blooms, in no particular order.

Reference	Locale	Description	Period	Socioeconomic Impact
(Anderson et al., 2000)	USA	Annual estimate of all HAB costs	1987 – 1992	Losses of US\$49M annually during study period
1979 (Figley et al., 1979)	USA – East Coast	Bloom of <i>Ceratium tripos</i> (Dinoflagellate)	1976	Losses of US\$1.3B
(Ritzman et al., 2018)	USA – West Coast	Study of a bloom of <i>Pseudo-nitzschia</i> (Diatom)	2015	Losses of \$97M and multiple social impacts
(Yang and Hodgkiss, 2004)	Hong Kong & China	Bloom of <i>Karenia digitata</i> (Dinoflagellate)	1998	Losses of HK\$250M
(Kim, 2010)	South Korea	Bloom of <i>Cochlodinium polykrikoides</i> (Dinoflagellate)	1995	Losses of US\$95.5M
(Kouakou and Poder, 2019)	European Union	Analysis of annual healthcare impact of HABs	2002-2006	Increased healthcare costs of US\$11M per year

2.3 Observed Trends and Future Behaviours

2.3.1 Climate Change and Ocean Acidification

The possible effects of climate change on algal blooms (and phytoplankton more generally) have been extensively researched. Although a consensus has not yet emerged around the *magnitude* of the impacts of a warmer, acidified ocean will have on these microorganisms (Hallegraeff, 2010), some conclusions about the *types* of effects likely to occur are widely backed up by modelling and observational studies. Firstly, while the phylogenetic diversity of phytoplankton and their disparate evolutionary histories provide community-level protection against change, warming environments will favour smaller species and hence reduce overall diversity (Zohary et al., 2021, Moran et al., 2010).

Shifts in both the spatial distribution of phytoplankton and the temporal dynamics of their blooming is likely to occur: changes have been observed in historical measurement records (Richardson and Schoeman, 2004, Trombetta et al., 2019) (though are difficult to decouple from natural variability (Elsworth et al., 2020)) and models suggest an amplification of these effects through the 21st century (Henson et al., 2010, Henson et al., 2018). The shifting phenology of algal blooms will impact higher trophic levels and could adversely affect fisheries (Asch et al., 2019). Biophysical modelling has indicated that heatwaves and changing water properties will increase the occurrence of harmful cyanobacterial blooms (Paerl and Paul, 2012, Jöhnk et al., 2008, Paerl and Huisman, 2009, Hayes et al., 2015), and that harmful dinoflagellate blooms will increase in frequency, spatial extent and duration (Gobler et al., 2017, Glibert et al., 2014).

2.3.2 Land use change and Eutrophication

Changing land use has direct and indirect impacts on freshwater and marine ecosystems. Lakes and rivers are especially affected by nutrient enrichment caused by runoff from agricultural land. Freshwater phytoplankton are hence vulnerable to an additional set of anthropogenic forcings, which in combination with a warming climate can change bloom timing (Shi et al., 2019), community structure (Pomati et al., 2017) and overall abundance (Bussi et al., 2016).

Chapter 2

The importance of phytoplankton and the environmental changes that they face reinforce the need for adequate monitoring and observation. Many of the changes hypothesised to be caused by anthropogenic climate change are challenging to unpick from seasonal, annual, or decadal cycles, and there is significant variability caused by a lack of observational data. The next sections will summarise the state of the art in phytoplankton observation and lead into a discussion of possible future improvements made possible by technological advances.

2.4 Microalgal analysis techniques

2.4.1 Background

Given their critical roles in global carbon transport and marine ecology, and considering their potential socioeconomic impacts, the quantity and diversity of phytoplankton and other organisms found in the same size fractions are among the most measured features of the oceans. '*Phytoplankton biomass and diversity*' is one of the priority '*Essential Ocean Variables*' defined by the *Global Ocean Observing System*, an international programme specifying key measurement parameters and operated by the Intergovernmental Oceanographic Commission (Muller-Karger and Kudela, 2016).

Phytoplankton have been studied since before 1800 (Fogg, 1990), with routine observations becoming common during the latter half of the 20th century. The *Continuous Plankton Recorder* (CPR) (Hardy, 1939), for example, has been used since 1931 and is now used globally within five regional surveys (Batten et al., 2003). Consisting of a silk mesh continuously fed through a roller mechanism within an outer body, the CPR is towed by ships of opportunity (e.g. ferries, freight vessels etc), and collects phyto- and zooplankton large enough to be retained by the 270 µm mesh pore size. CPR data has been the subject of over 1000 peer-reviewed articles and allows the analysis of long-term trends, crucial for understanding the impact on plankton of future climate change.

The CPR has the largest spatial extent of any phytoplankton observation programme and, as the phytoplankton are collected and stored, can provide taxonomically-resolved abundance data (Lombard et al., 2019). Despite the many benefits of the CPR there are still significant limitations. Most importantly, manual microscopy is still the fundamental analysis technique for preserved CPR samples. This necessitates the retention of a team of specialist taxonomists and severely limits the number of collected cells that can be analysed: only around 1/8000th of the total number of cells collected are identified. Furthermore, the large footprint of the unit (approx. 1m long and weighing 85kg) and requirement that it be towed limits the ease of deployment, especially for short-term studies, and prevents observations on coasts/rivers, fixed moorings or on smaller vessels.

2.4.2 Technologies

Characterisation and monitoring of microalgae has mainly been conducted using optical microscopy (Karlson et al., 2010, Tsaloglou, 2016), which is costly, labour-intensive, low-throughput and relies on analysis of images by expert taxonomists, a specialism which is in steep decline (Culverhouse et al., 2006). Furthermore, microscopy remote from the point of sampling also necessitates the use of fixative agents which require additional staff time and chemicals and can, depending on species, drastically alter cellular morphology and pigmentation from that seen in the wild (Dunker et al., 2018). Species composition measurements with high temporal and spatial resolutions are currently impossible, yet remain a central aim of research into aquatic primary production (Smetacek and Cloern, 2008).

Due to these limiting factors several remote and in-situ analysis techniques have been developed, including the use of *satellite colourimetry* and *bulk chlorophyll fluorescence*. These bulk analysis techniques have been widely adopted for approximating aquatic primary production rates and phytoplankton size ranges (Brotas et al., 2013, Sun et al., 2022) but provide little or no taxonomic data and are of no use in monitoring HAB-forming species specifically (Yoder et al., 2010). Additionally, extrapolations from bulk measurements to the characteristics of single cells can be unreliable due to high heterogeneity in, for example, chlorophyll to biomass ratios of different species (Haraguchi et al., 2017).

As such, there has been a recent drive towards sensors which analyse cells individually; the first major breakthrough in this effort was the application of flow cytometry (FC) to plankton enumeration and identification (Collier, 2000) (Figure 3). FC measures, using photomultiplier tubes (PMTs), PIN photodetectors or avalanche photodiodes (APDs), the fluorescence intensity, forward- and side-scatter of each cell intersecting one or more collimated light sources, typically lasers (Collier, 2000). The value of FC monitoring is limited by the low information content of recorded data: while cell size, pigment concentration/ratio and basic morphological measurements (e.g. circularity, equivalent spherical diameter, complexity) can be derived from scatter and fluorescence signals, additional analysis and verification by microscopy is required to determine taxonomic identity, health, life stage, etc. (Marie et al., 2014, Dubelaar and Gerritzen, 2000)




Figure 3 - Principle of standard flow cytometric analysis of particles suspended in a fluid sample.

Particles flow past a detector where they are interrogated by a laser. Detected scattering and wavelength changes of the laser signal can be used to infer particle size, approximate shape, fluorescence.

While most flow cytometers are large instruments confined to laboratories, miniaturised FCs can be brought aboard research vessels in order to analyse at the point of sampling, i.e. without the use of fixatives (Álvarez et al., 2011). Recently, several automated, *in-situ* FC instruments have been developed (Dubelaar et al., 1999, Olson et al., 2003) which enable the automated continuous analysis of phytoplankton optical properties at fixed or towed locations (Dubelaar et al., 1999, Olson et al., 2003). However, current platforms can only sample semi-continuously, are expensive and lack autonomy (Tsaloglou, 2016).

The combination of the throughput of flow cytometry, the information content of microscopy and high-accuracy automated image analysis in a single package has been a significant aim of sensor development for biological oceanography over the last two decades (Sieracki et al., 2010, Blaschko et al., 2005, Benfield et al., 2007). Platforms which image individual particles in a sample moving through a flow cell are known as *Imaging Flow Cytometers* (IFCs) and have been developed simultaneously for biomedical and environmental applications. IFCs have been used in

Chapter 2

a variety of research targeting both marine and freshwater phytoplankton and will be the focus of the next section.

2.5 Imaging Flow Cytometry

2.5.1 Introduction

Possible imaging modalities of IFCs encompass bright field, fluorescence, holography, phase-contrast, and other techniques, but all collect multidimensional data via digital image sensors rather than PMTs or other 1-dimensional sensors. This potentially enables continuous taxonomic classification to the same level as optical microscopy while sampling many more cells per second than possible manually. Furthermore, automated IFC enables novel sampling methodologies previously impossible, including continuous monitoring of plankton morphological change induced by changing nutrient conditions (Lei et al., 2016) and identification of viability and cell-cycle phase (Stavrakis et al., 2019, Dashkova et al., 2017).

Successful implementation of continuous automated IFC poses a number of challenges which commercially available instruments have yet to fully overcome (Heo et al., 2017). Imaging of moving particles requires camera frame rates and exposure times sufficient to avoid motion blur along the flow axis (Wu and Chan, 2013) – to achieve high throughput, particle velocities typically exceed 1ms^{-1} (Goddard et al., 2007). Acquiring an image with enough contrast necessitates the collection of a sufficient number of photons during the camera exposure window; the light source must be capable of providing this high-intensity illumination (Miura et al., 2018, Hess et al., 2015). Particles must be confined to the focal plane of the imaging system; those flowing too close or far from the objective will be out of focus (Holzner et al., 2018). Data volumes from IFC instruments pose a significant storage and transmission challenge; image files are extremely large compared with FC data and IFC devices may generate many hundreds of gigabytes of data per day of operation (Han et al., 2016). Data processing is a significant bottleneck as classifying high volume, complex images is a nontrivial computing task. Difficulties shared by both FC and IFC design include clogging of the flow cell and shear forces on particles, which, under high flow regimes, can deform or damage cells (Zmijan, 2016). Finally, the size, robustness, cost and complexity of the instruments must be minimised in order to facilitate wide deployments.

These challenges will be further investigated, with respect to current instruments and recent developments in the fields of microfluidics, imaging devices and manufacturing techniques, in the following sections.

2.5.2 Current IFC devices used for phytoplankton analysis

Commercial IFCs

Imaging flow cytometers have been used in biological oceanography to evaluate phytoplankton species dynamics (Campbell et al., 2016, Brownlee et al., 2016, Bolanos et al., 2020, See et al., 2005) and taxon-specific responses of phytoplankton to short- (Anglès et al., 2015) medium- (Anglès et al., 2019, Kamykowski et al., 2013) and long-term environmental change. *In situ* IFC has allowed the detection of HABs (Campbell et al., 2013, Campbell et al., 2010, Kudela et al., 2008, Buskey and Hyatt, 2006) and analysis of bloom dynamics (Laney and Sosik, 2014, Zarauz et al., 2008, Cetinic et al., 2014). Analysis of ballast waters (for invasive plankton species)(Romero-Martinez et al., 2017), detection of harmful algae in watersheds for drinking water (Park et al., 2019) and algal life-cycle characterisation (Traller and Hildebrand, 2013, Dapena et al., 2015) indicate the potential of IFC for novel phytoplankton research for which traditional microscopy would be insufficient or impractical.

A recent review article which investigated the tools and technologies available for *in situ* biological oceanography highlighted *Optical Sensors for in situ Plankton Monitoring* as one of two areas needing rapid cost reductions in order to expand observational capacity and fully democratise access to instrumentation (Wang et al., 2019).

A separate article in the same special issue looked in more detail at the specific platforms and instruments available for optical detection and identification of plankton cells and provided a review of the commercially available imaging flow cytometers currently known to be used in the ocean sciences (Lombard et al., 2019). An adapted table of the instruments identified in this article is presented below, in Table 2.

Table 2 - Table of commercial imaging flow cytometers used in biological oceanography for the purpose of plankton observation. Reproduction of data produced in Lombard et al..
ESD: equivalent spherical diameter

Instrument	Size Range (ESD) [μm]	Sample Throughput [$\mu\text{L}/\text{min}$]	Use Case	Approximate Cost [UK £]	Seller
<i>Imaging FlowCytoBot (IFCB)</i>	10 – 80	250	<i>In situ / on-board / lab</i>	125,000	<i>McLane</i>
<i>CytoSense and CytoSub</i>	1 – 800	5 – 1000	<i>In situ / on-board / lab</i>	90,000	<i>CytoBuoy b. v.</i>
<i>FlowCam Nano</i>	0.3 – 30	20	On-board / lab	90,000	<i>Fluid Imaging Technologies</i>
<i>FlowCam</i> (different objectives & flow cells)	2X: 75 – 1000 4X: 20 – 3000 10X: 10 – 100 20X: 3 – 50	50 – 5000 (depending on flow cell)	On-board / lab	45-75,000	<i>Fluid Imaging Technologies</i>
<i>FlowCam Macro</i>	300 – 5000	100,000 – 900,000	On-board / lab	45,000	<i>Fluid Imaging Technologies</i>

Chapter 2

The commercial instruments presented above clearly reinforce the arguments laid out by *Wang et al.* regarding the need for lower-cost imaging devices in oceanography. With a minimum cost of £45,000 (more when considering flow cells, objectives, software, training, etc), current IFCs are too costly for widespread adoption on routine scientific cruises and in small-scale experiments. The budgetary considerations of laboratories in resource-constrained environments means that IFC is a methodology inaccessible to many researchers, and hence large swathes of global coast lack any phytoplankton imaging data. Furthermore, the limited size ranges capable of being analysed by the instruments described in Table 2, and their low throughputs, demonstrate the trade-offs that must be considered when addressing specific scientific goals.

To further illustrate the limited scales of phytoplankton for which each instrument is appropriate, Figure 4 provides a graphical representation of the size spectra each instrument is able to analyse, with the red bars giving the size ranges that have been quantitatively cross-validated with other more developed measurement techniques (Lombard et al., 2019).

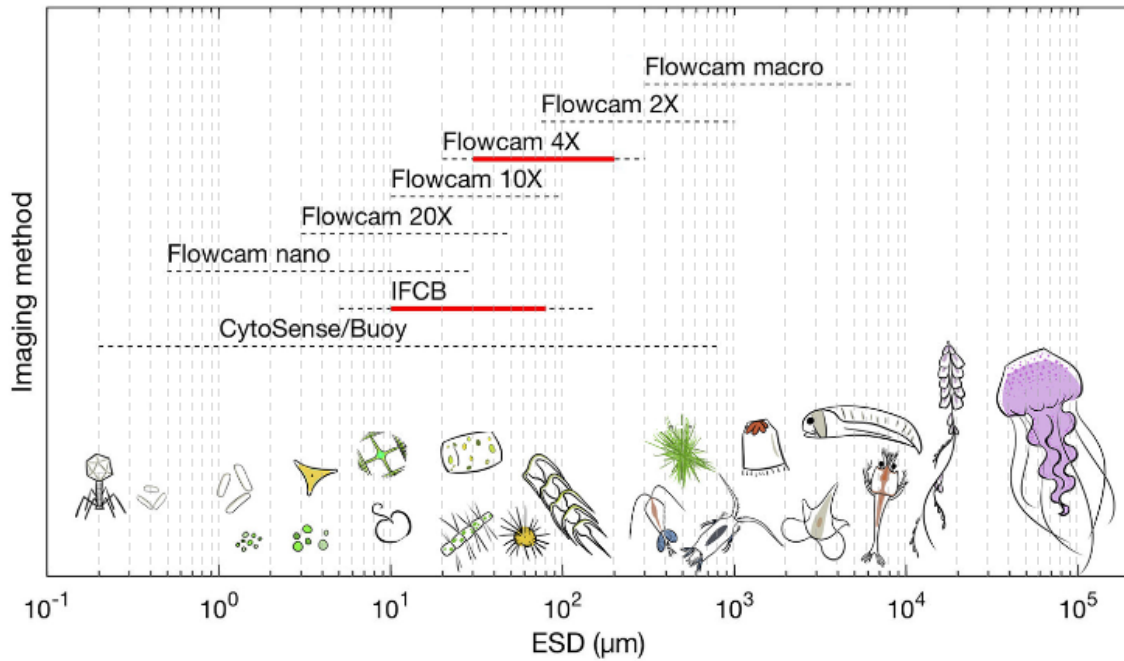


Figure 4 - (Adapted from Lombard et al) Graphical illustration of the size ranges (in ESD, equivalent spherical diameter) each of the IFC instruments discussed in their review is capable of imaging. Note that CytoSense/CytoBuoy, while here presented as having the largest range, is actually a combination standard and imaging flow cytometer, and that imaging alone is likely to have a very similar measurement window to the IFCB (the actual imaging size range is unpublished). The red bars indicate the possible analytical size range which has been cross-validated with other established techniques as opposed to merely those published by the device manufacturers (grey dotted bars).

The available IFC instruments are therefore expensive, have low sample throughput, and are only able to simultaneously image plankton of a narrow range of sizes. An analysis of the constituent elements of imaging flow cytometers follows, in which the source of these limitations and possible avenues for addressing them is discussed.

Research/Prototype IFCs

Several IFC systems in various stages of development have been demonstrated for phytoplankton imaging. A team of French researchers based at IFREMER (*L'Institut Français de Recherche pour l'Exploitation de la Mer*) and The University of Mons sought to develop a cheaper and higher-throughput adaptation of the FlowCam, which they named FastCAM (Colas et al., 2016). The

Chapter 2

FastCAM uses off-the-shelf optomechanical components from Thorlabs and a 2-megapixel, high-speed (340 FPS) camera. The device uses a FlowCam flow cell, and has a sample throughput of 10 mL/15 min (40 mL/hr). Images from the FastCAM have higher spatial resolution than the FlowCam, and overall cost is lower. The main drawback is that the device is still in a prototype stage and requires manual assembly and alignment, which can be challenging for non-specialists.

One device which was developed to meet similar goals as those discussed for this project is the PlanktonScope, a modular IFC recently developed by engineers at Stanford University.

PlanktonScope is based around a *Raspberry Pi* microcontroller and camera module, and is extremely low cost (< US \$500) (Pollina et al., 2020). This device is open-source, based around disposable plastic flow cells, and can image 1.7 mL/min. The optical setup can achieve a magnification of 0.75X and the camera has a resolution of 1.5 μm , but the minimum object size demonstrated by the authors to be successfully imaged was 35 μm . The PlanktonScope represents a major step toward lower-cost IFCs and is primarily aimed at citizen scientists, with the stated aim of the authors being to enable citizen owned small vessels (e.g. sailing yachts) to analyse plankton with very low startup and training cost.

A group at the University of California recently demonstrated a low-cost (~US \$2500) prototype IFC based around the principle of holography (Göröcs et al., 2018). Incoherent light is used to illuminate phytoplankton within a sample and 3D reconstructions of the cells can then be computationally derived using the process of deconvolution (Figure 5). This particular IFC uses deep-learning based phase-recovery for hologram deconvolution, which allows the accurate recovery of true-colour images without explicitly programming the relevant optical transfer equations, which would be very complex. An additional advantage of this approach is that deep neural networks run extremely well on graphics cards. Göröcs et al. exploit this fact by using a powerful GPU to perform image recovery in real-time (ie. Faster than the camera's frame rate). The device has a throughput of 100 mL/hour.

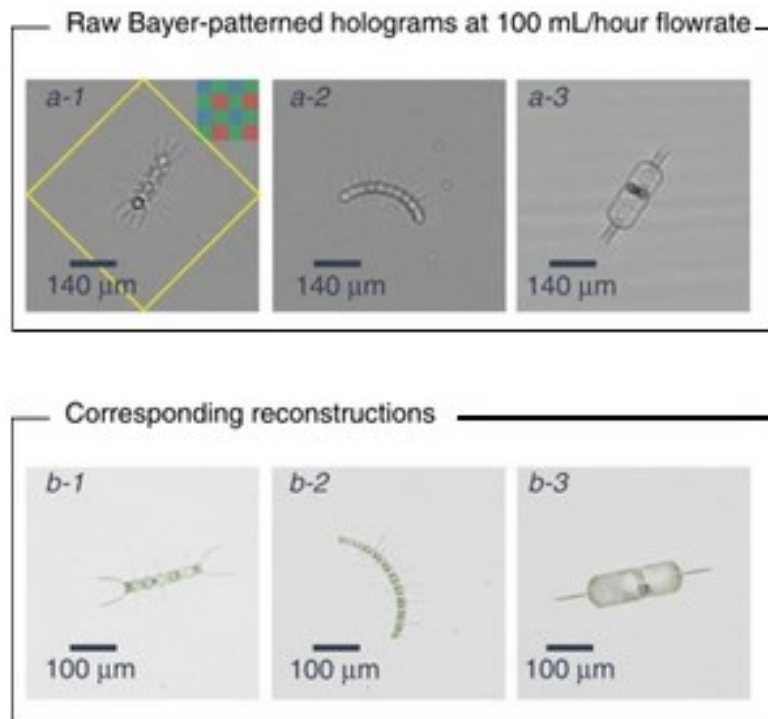


Figure 5 - Raw holograms and reconstructed phytoplankton images generated by the prototype IFC developed by Göröcs et al. (2018). Each hologram must be processed into a RGB image via a computational process called deconvolution, which in this work was performed by a Deep Neural Network.

2.5.3 Camera Technologies

The two fundamental technologies behind digital imaging are *charge-coupled devices* (CCDs) and *complementary metal-oxide-semiconductor* (CMOS) sensors, both of which have seen use in commercial IFC. These sensors both convert accumulated photons into charge on pixel arrays using the photoelectric effect. While entire CCD arrays are read by a single analogue-to-digital converter (ADC) and amplified, on a CMOS sensor, each pixel has its own amplifier and ADC. This difference is illustrated in Figure 6 which depicts the pixel layouts of each sensor architecture.




Figure 6 – CCD and CMOS sensor architectures. While CCDs, which are more common, use a single ADC for the entire pixel array, CMOS sensors, which have seen rapid technological improvements, have more complex pixel topologies due to every pixel having its own ADC.

The practical result of this difference is that CCDs usually produce higher-quality and lower-noise images but have a slower readout time due to the entire pixel array being converted to signal simultaneously. CCD manufacturing is more developed than CMOS, but recent, rapid improvements in CMOS technology have led to comparable image quality between the sensors. A major advantage of CMOS, in addition to much faster image acquisition and readout, are their far lower power requirements than CCDs (Up to 100X lower Wattage) (Hain et al., 2007).

The *Imaging Flow Cytobot* (IFCB)(Olson and Sosik, 2007a), a submersible IFC platform capable of in-situ imaging of plankton in the 10-100 μm range, uses a *Sony EXview HAD* monochrome CCD, which, when coupled with the on-board 10X objective, can image with a resolution of $\sim 1\mu\text{m}$ and a frame rate of $\sim 11\text{fps}$. The FlowCam series of instruments use both CCD and CMOS, depending on model.

A series of benchtop instruments used extensively in biomedical research and the life sciences (Headland et al., 2014, Lopez-Riquelme et al., 2013, Gautam et al., 2018)e.g. (Headland et al., 2014, Lopez-Riquelme et al., 2013, Gautam et al., 2018) but which is not routinely used in oceanography due to the cost and size (and therefore is not included in the review in Section 2.5.2), the *Amnis ImageStreamX*, use a custom *time-delay-integration* (TDI) CCD to image flowing

particles (Figure 7). This specialised readout technique allows high quality imaging despite particle motion by shifting pixel values along the sensor at the same speed as particle movement, which is equivalent to ‘tracking’ the particle by moving the entire camera (George et al., 2004).




Figure 7 – ImageStreamX Mk II optical path showing fluorescence laser and wideband red brightfield illuminator. A complex custom detector simultaneously detects different fluorescent channels and brightfield images, made possible by a spectral decomposition of the signal into different colour bands.

2.5.4 Illumination Techniques

Numerous lighting setups have been used on commercial and research IFCs. Developing upon traditional FC implementation, the *FlowCam8400* IFC (Figure 8) uses a single diode laser and a xenon flash lamp to illuminate cells in order to image simultaneously in brightfield and two fluorescence channels. While this arrangement has the benefit of providing additional morphological information via detection of autofluorescent or stained components, it requires expensive dichroic mirrors and precise optical alignment.




Figure 8 - FlowCam8000/8400 IFC optical setup. One laser and one flash lamp (dotted box) illuminate the detection region and two PMTs provide wavelength-configurable fluorescence detection while a CCD images in brightfield.

Recently there has been a trend towards using inexpensive, high-power LED illumination as part of IFC optical setups; these offer smaller footprints and comparable power to laser sources at a fraction of the price. For fluorescence imaging, however, LEDs are suboptimal compared to diode lasers due to their typically wider emission spectra and lack of collimation, both of which are required for optimal detection limits.

The holographic IFC presented by Göröcs et al. (2018) and described in the previous section uses a pulsed chip-on-board (COB) LED to provide narrowband red, blue and green illumination to the flow cell. Combined with a modern low-cost CMOS camera, this optical setup enables the IFC to maintain a small footprint while displaying very promising results at a total cost of under \$2500 USD (Figure 9). This setup provides fluorescent imaging but requires real-time image deconvolution to reconstruct the holograms it records and thus necessitates connection to a computer with a high-performance graphics card, which are often expensive.




Figure 9 - Göröcs et al. low-cost holographic IFC optical setup, showing the miniturised 3-colour illumination COB (inset). Though a holographic imaging device, this system has a similar purpose to the IFC proposed in this thesis.

2.5.5 Focussing Techniques

Particle focussing is a significant area of IFC development and is one of the major controls on IFC throughput and image quality. In traditional FC (Figure 10), coaxial hydrodynamic focussing typically aligns particles into a single-file stream through the detection region. This is achieved by injecting sheath fluid around the sample, narrowing the sample flow profile while maintaining laminar flow. Hydrodynamic focussing can take place in 2- or 3 dimensions, depending on placement and number of sheath fluid ports as illustrated by Figure 11.




Figure 10 – Cross-sectional illustration of the hydrodynamic focussing principle. Injected cells are focussed into a single-file stream by laminar flowing focussing fluid.




Figure 11 – 2-dimensional (a) and 3-dimensional (b) hydrodynamic focussing. In these simulations a red buffer fluid controls the flow profile of a blue carrier fluid. Adapted from Hamilton et al.

Other techniques include inertial focussing, which has the advantage of not requiring sheath fluid or other external forces. Inertial focussing uses wall effects and gradients in shear forces developed using specially designed flow channels (e.g. serpentine channels) to position particles in the centre of the detection region (Yang et al., 2018b). A schematic of the process of inertial focussing within a serpentine channel is presented in Figure 12, adapted from Di Carlo et al. (2007). While holding great potential for medical IFC where cells will usually be of a known, narrow size range, the large heterogeneity of microalgae cell sizes has so far prevented inertial focussing being successfully demonstrated for an instrument designed to focus natural phytoplankton communities.




Figure 12 – Schematic adapted from Di Carlo et al. illustrating the principle of inertial focussing of particles within a microfluidic serpentine channel.

Chapter 2

Some IFCs are designed to function without particle focussing. The FlowCam series of instruments, which is particularly popular in biological oceanography, being the most cited device in a recent review of IFC in the field (Dashkova et al., 2017), does not move particles in flow using any of the methods discussed. Instead, patented optics based around a Fresnel lens phase plate enhance the depth of focus of the objective lens, at the cost of image resolution. This allows the mechanical complexity of the instrument – and the likelihood of clogging or other failure – to be minimised while maintaining high sample throughput. For larger particles (e.g. large diatoms, zooplankton), this is clearly advantageous; but for smaller phytoplankton, the lower resolution will be a limiting factor which prevents analysis below a certain size threshold. For the purposes of identification (rather than just measurement and counting), sacrificing resolution will significantly lower the accuracy of taxonomic classification, including classification performed by automated methods such as neural networks, which are explored in Chapter 4.

The final focussing technique of note, and the technique upon which the remainder of this thesis will build, is acoustic focussing. Acoustic levitation, or *acoustophoresis*, makes use of acoustic pressure to manipulate particles and can be applied to microfluidic systems by generating ultrasonic waves within the carrier fluid. The study of *acoustofluidics* has been developed over the last decade due to its purely mechanical, gentle repositioning of particles and has been used in research where non-contact, label-free particle separation and manipulation is required. As well as acoustic focussing and alignment of particles within fluids (Yang et al., 2018a, Jayasinghe, 2020), acoustofluidic devices have been used for cell trapping/concentration (Leibacher et al., 2015, Evander et al., 2007, Fornell et al., 2019, Bach and Bruus, 2020), cell sorting/separation (Huang, 2019, Zhang et al., 2020, Yang et al., 2018c), cell differentiation (Wu et al., 2019a, Wu et al., 2019b, Xie et al., 2020), mechanical stimulation during cell culturing (Li et al., 2014) and many other applications within the micron size range.

Chapter 2

Concentration, buffer-exchange, separation and alignment, four of the most common acoustofluidic applications which are particularly of interest in the life sciences are illustrated in Figure 13 (Antfolk and Laurell, 2019). For the purpose of acoustic focussing within an IFC, alignment into a single layer is most common, though parallel processing of two or more layers may be desired in some instruments.




Figure 13 – (Reproduced from Antfolk and Laurell 2019) The four most common applications of acoustofluidic devices.

The principle and details of the acoustophoretic approach to particle positioning in imaging flow cytometry are expanded upon in Section 2.6.

2.6 Acoustofluidics

2.6.1 Introduction

Acoustic radiation pressure within a fluid system causes acoustic discontinuities (particles or second-phase fluids) to experience small forces. By inducing an acoustic standing wave with a half-wave resonance, suspended particles are typically pushed toward a central pressure minimum known as an acoustic pressure node. The acoustic force acting upon particles is described by Equation 1 (Bruus, 2012), which in the case of a *planar* standing wave simplifies to Equation 2 (Lenshof et al., 2012).

$$F^{rad} = -\left(\frac{4\pi a^3}{3}\right) \nabla \left(f_1 \frac{\langle p_1^2 \rangle}{2\rho_0 c_0^2} - f_2 \frac{3}{4} \rho_0 \langle v_1^2 \rangle \right),$$

$$f_1 = 1 - \frac{k_p}{k_0}, \quad f_2 = 2 \frac{(\rho_p - \rho_0)}{(2\rho_p + \rho_0)}$$

Equation 1 – Equation for Primary Acoustic Radiation Force, F^{rad} , where: a is the particle radius; K_0 is the compressibility of fluid; K_p is the compressibility of the particle; p_1 is the pressure field; v_1 is the velocity field; ρ_p and ρ_0 are density of particle and fluid, respectively; c_0 is the speed of sound in the fluid.

$$F_{ax} = 4\pi a^3 E_{ac} k \sin(2kz) \Phi$$

Equation 2 – Acoustic Radiation Force, F_{ax} in a one-dimensional standing plane wave. Where a is the particle radius; E_{ac} the acoustic energy density (Equation 4); k the wavenumber ($2\pi f/c_0$); f the frequency; z the distance from the pressure node in the wave propagation (thickness) axis; Φ the *acoustic contrast factor* (Equation 3):

$$\Phi = \frac{\rho_p + \frac{2}{3}(\rho_p - \rho_0)}{(2\rho_p + \rho_0)} - \frac{1}{3} \frac{\rho_0 c_0^2}{\rho_p c_p^2}$$

Equation 3 – Acoustic Contrast Factor, Φ , which determines the sign of the Acoustic Radiation Force acting upon a particle in a standing acoustic wave. Particles with negative Φ move away from the pressure node; particles with positive Φ are pushed toward it.

$$E_{ac} = \frac{1}{4} \frac{P^2}{\rho_0 c_0^2}$$

Equation 4 – Acoustic Energy Density, E_{ac} . P is the pressure amplitude, which can be measured experimentally in an acoustofluidic system or analytically calculated based on modelling, as in Section 3.3.2.

Two key observations to make from these equations are that the acoustic force acting upon particles scales with the cube of the particle radius (Equation 2) and that the direction of motion induced by the force is determined by the relative differences between the density and compressibility of the particle and that of the carrier fluid, as described by the acoustic contrast factor (Equation 3).

By applying a sufficient acoustic radiation force to particles in a flow cell, *acoustophoretic levitation* can be achieved. In this situation, the sinking rate of particles due to gravity is overcome by the acoustic forces, and all particles which have a positive acoustic contrast factor with their carrier media are pushed toward pressure minima whose positions depend on the excitation frequency. By exciting a half-wave resonance with one pressure node in the (axial) centre of the flow channel, all suspended particles focussed to this central position. Particles with negative acoustic contrast factors with their media undergo the opposite force, moving toward pressure anti-nodes at the top and bottom of the flow cell. As lipids have negative acoustic contrast with water, algae with extremely high lipid content, such as those used in biotechnology, e.g. biofuel production, may never focus to the pressure node of a half-wave resonator.

2.6.2 Current Acoustically Focussed IFC Devices

Acoustic focussing has been developed due to a number of benefits compared with other techniques: the forces acting upon cells are gentler than in hydrodynamic focussing, with less tendency to damage or deform morphological features; the focussing is achievable across a wide range of cell sizes; there is no need for consumables and the microfluidic complexity is greatly reduced. Acoustic focussing has been demonstrated in research IFC instruments (Zmijan et al., 2015) and as add-ons to the hydrodynamic focussing equipment of the IFCB (Olson et al., 2017)

Chapter 2

and Thermo-Fischer Attune® NxT IFC (Bautista-Chamizo et al., 2018); as yet, however, there is no commercial IFC which solely uses acoustic levitation for focussing. The next section will build upon the principle of acoustophoretic focussing and investigate the application of acoustic-only particle focussing within a prototype IFC for phytoplankton imaging and analysis

Chapter 3 Design and Characterisation of a Novel Low-Cost IFC using Acoustophoretic Focussing

3.1 Introduction

As previously discussed, manual examination of microalgae using optical microscopy comes at a high cost and low reliability and is unsustainable due to the limited and declining number of taxonomists. Commercial IFCs, which, as was noted in Chapter 2, hold potential to address these issues, have so far faced barriers to implementation as part of routine ocean observation due to current instrument cost and complexity. As such, this chapter details work to address this challenge, in the development an IFC with the lowest complexity and the cheapest possible off-the-shelf components, and with as few custom-made components as reasonably possible.

A significant additional design goal is to increase cell throughput compared to existing (hydrodynamically-focussed) instruments such as the IFCB. As was discussed in the previous chapter, hydrodynamic focussing requires cells to be imaged one at a time. By utilising acoustophoretic focussing, the new instrument developed in this chapter will be capable of simultaneous imaging of multiple phytoplankton cells in a 2-dimensional plane, and so will have a throughput significantly higher than would be possible with hydrodynamic or inertial focussing. Since phytoplankton will be presented to the imaging optics in a plane of a constant distance from the objective, no depth-of-focus extending optics are required, which ensures that cells will be imaged at a higher resolution compared to unfocussed IFCs such as the FlowCam.

Furthermore, the complexity of this new system can be significantly lower than hydrodynamically-focussed IFCs. Acoustic focussing works well in a simple straight flow cell which is far simpler to fabricate than a serpentine or other channel suitable for inertial focussing, a technique which is not currently used by major commercially available IFCs. Overall, the method holds great promise for improving throughput while reducing the cost, size and complexity versus current commercial instruments, while also allowing higher image quality than recent low-cost prototype IFCs such as those described in Section 2.5.2.

3.2 Design Brief

A block diagram of the proposed IFC system architecture is presented in Figure 14. In summary, the suspended microalgae cells will be injected into a layered resonator flow cell using a pump, which could be syringe pump, peristaltic pump or even simple gravity-fed, depending on configuration and volume/temporal resolution requirements. Once analysed, sample waste can be disposed of, or, for applications requiring repeated measurements of the same source over time (e.g. monitoring a single live culture over time), could be recirculated.

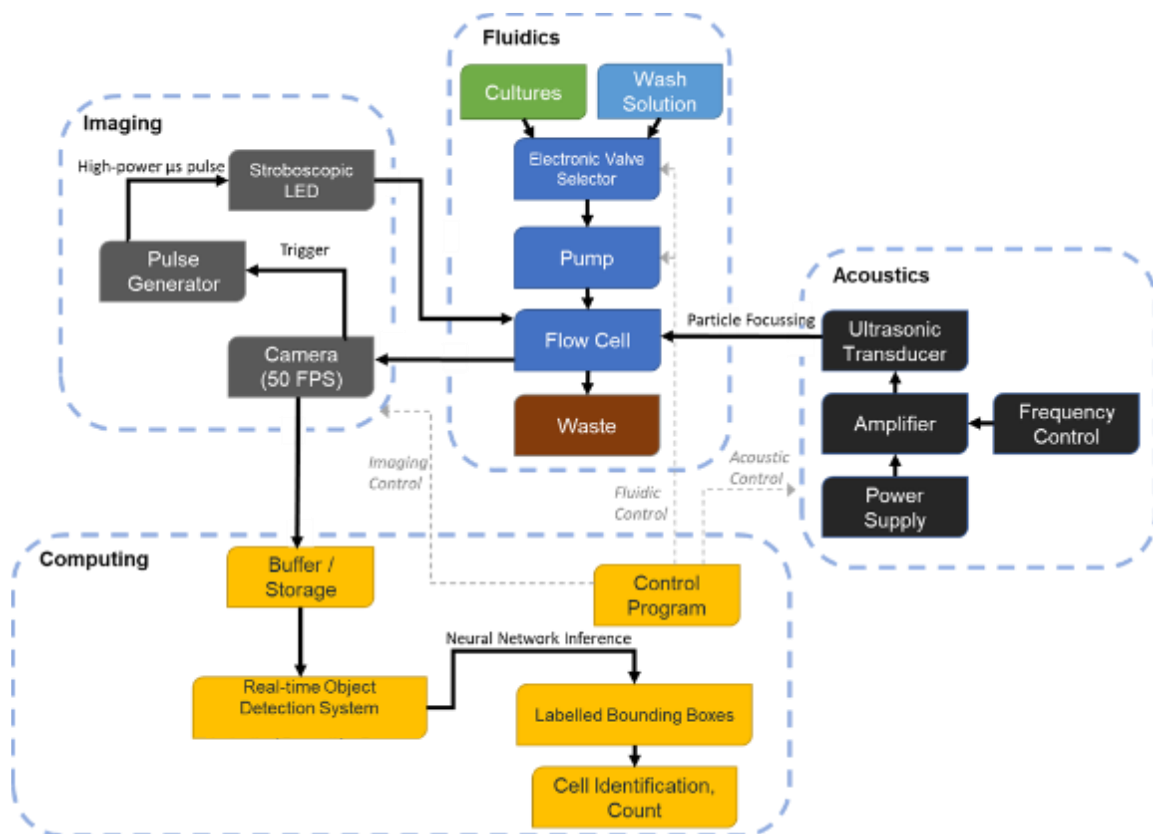


Figure 14 - Block diagram for proposed acoustically-focussed imaging flow cytometer. The device itself will consist of an acoustophoretic focussing flow cell with necessary acoustic driving and control hardware, microfluidic pumping, optical components (camera, lenses and illumination). Once images have been acquired, analysis in the form of measurement and/or identification can be performed in real-time or from storage.

The flow cell will be acoustically excited by means of a piezoelectric transducer driven at a resonance frequency such that an acoustic standing wave forms in the flow cell with a pressure

Chapter 3

minimum (node) around the channel half-height. Cells in the flow stream will thus be pushed towards this central node, remaining at approximately the same position in the flow cell Z-axis (along the optical path) after passing the transducer, as their sinking rate due to gravity will be negligible compared to their linear velocity over the distance from focussing to imaging region. A high-frequency power amplifier will be needed to boost the signal generator output to the voltages required to drive the transducer. Acoustic resonance will be maintained by setting the signal generation system to perform a frequency sweep around the predetermined flow cell resonant frequency such that minor drifts in exact resonance frequency (i.e. due to temperature/salinity changes in the flow medium) will not significantly impact the total energy absorbed by the system, and no active feedback control will be necessary. Previous research has demonstrated that when the sweep period is fast enough (e.g. 50ms), this technique results in an averaging of the force profiles of each frequency activated with the result approximating resonance (Glynne-Jones et al., 2010a). An analytical validation of this frequency sweep approach will be presented in Section 3.3.2.

A camera will be positioned downstream of the transducer and will image at high speed the focussed cells as they flow past, the algae being illuminated by a high-power stroboscopic LED in a transillumination arrangement (with illumination pulses synchronised with the camera by trigger electronics). By using pulse durations of 1 – 20 μ s and low duty cycles of less than 1:100, it will be possible to drive the LED at over 10X its rated current for corresponding luminosity increases of over 5X without causing damage (Willert et al., 2010). The requirement for synchronisation of illumination pulses with the camera's aperture will necessitate the use of cameras with external trigger inputs or outputs; these are common in life science and industrial applications.

Images will be transferred from the camera to a computer system either directly or using frame-grabber electronics (depending on camera choice and required framerate) and will then either be stored for later analysis or be analysed in real-time. This analysis could take the form of, for example, image processing to determine particle size spectra, or conducting taxonomic identification using a pre-trained classification system such as a convolutional neural network, (CNN). This analysis capability is not within the scope of this chapter, and instead fully explored in Chapter 4.

3.3 Flow-cell implementation

3.3.1 Approach

Design 1

The first flow-cell design made use of a *matched-layer half-wave planar resonator* device following the implementation of Zmijan et al. (2015). In this type of device, an ultrasonic transducer is bonded to a *matching/carrier layer*, which with a *reflector layer* sandwiches a fluid layer. In this arrangement an acoustic standing wave is excited in the device with a half-wavelength pressure anti-node located in the fluid channel. This design has the advantage of being simple to produce, consisting only of layers of substrate bonded together, and fits well the simple geometry required for an IFC. The IFC developed here builds upon previous work by redesigning the layered flow cell for integration into a complete, portable and low-cost system as will be described below.

In the work of Zmijan's group, the acoustically focussed flow cell was designed around a fluid layer sandwiched between two standard 1 mm-thick microscope slides. As in the new IFC presented here, a major design criteria is the development of an imaging platform capable of resolving cells which are smaller than those resolvable by existing cheap instruments such as the *PlanktonScope* (Pollina et al., 2020) discussed in Section 2.5.2, optical quality is a key concern. In testing of the flow cells reproduced according to their designs, the thickness of the microscope slide through which the cells were imaged resulted in a degradation of the imaging performance to an unacceptable extent. While this could potentially be mitigated by use of a high working-distance objective, these lenses are generally far more expensive than would be appropriate given our design goals. Therefore, the Zmijan's flow cell design was improved upon by designing a new flow cell making use of a 0.17 mm thick cover slip for the matching layer, instead of a microscope slide. This novel flow cell design has the advantage of significantly greater optical transmission, allowing greater imaging quality of the objects in flow and the use of low- working distance objectives. A result of changing the layers within which the acoustic standing wave is established, however, is that it is not possible to use their assessments of acoustic properties and

Chapter 3

instead there is a need to comprehensively model and test the new design, work which is detailed in the subsequent sections.

Fabrication of the novel flow cell entails a process similar to that described by (Zmijan et al., 2015). Briefly, three layers of double-sided tape (9629PC, 3M, USA) were laser-cut to form a straight channel of dimensions 6 x 60 x 0.390 mm, which was then sandwiched between a standard glass microscope slide (25 x 75 x 1 mm) with drilled fluid inlet ports, and a 75 x 25 x 0.17 mm coverslip. A PZT transducer (PZ26, Ferroperm, Denmark) was bonded to the coverslip using epoxy (301, Epotek, USA) and the same epoxy used to bond a Perspex frame around the flow-cell as shown in Figure 15. The tubing used for introducing samples to the flow cell was LDPE, of 0.58 mm Inner Diameter, 0.96mm Outer Diameter (Smiths Medical Portex, Fisher Scientific, USA).

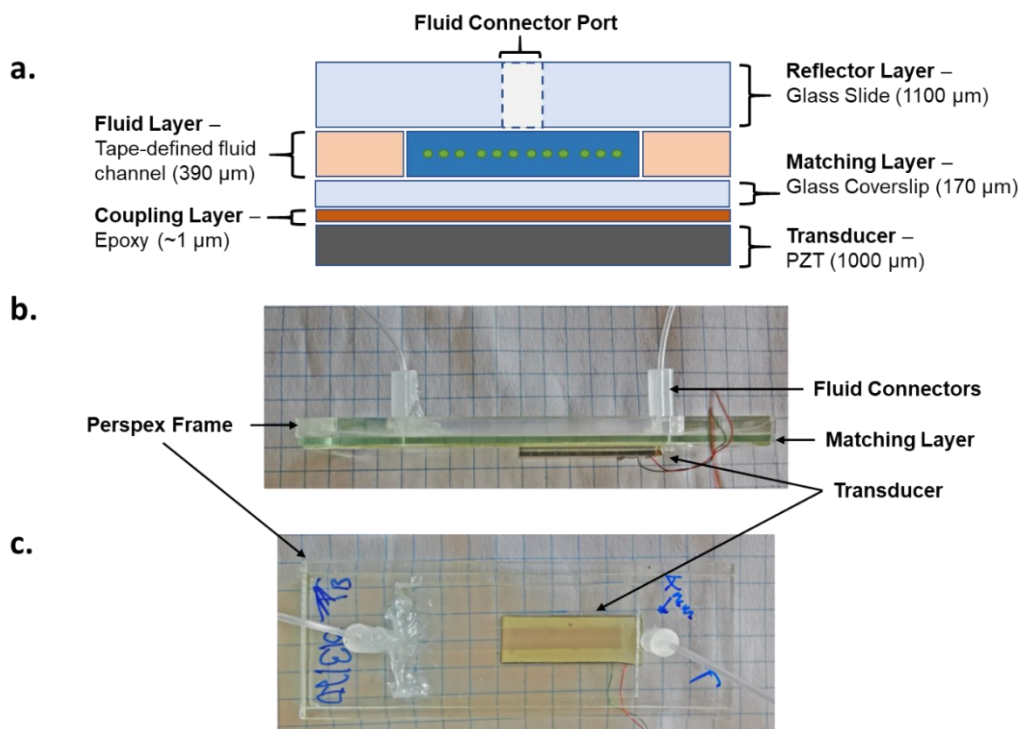


Figure 15 – The flow cell design. a.: schematic cross-section of the acoustic focussing region showing the thicknesses of each layer. b.: Lateral view of one of the constructed flow cells showing fluidic connectors, acrylic frame, microscope slide and transducer. c.: Axial view of the same device.

Chapter 3

Initial qualitative testing of the fabricated flow cell under a FITC-filter cube on an inverted microscope (*IX71, Olympus*) using fluorescent 10 μm polyester beads (*YG Fluoresbrite microspheres, Polysciences Inc., USA*) demonstrated good acoustic focusing. Further characterisation of the acoustic performance of the flow cell follows in Sections 3.3.2 and 3.3.3.

Design 2

A second flow cell implementation was explored but ultimately proved unacceptable, the reasons for which will be detailed below. The flow cell design was based around a borosilicate glass capillary of dimensions 30 x 6 x 0.3 mm, (*5356-050, VitroCom, USA*). An identical PZT transducer to that used in Design 1 was bonded to the wide face of the capillary using the same epoxy. Fluidic connectors were formed by means of heat-shrink tubing around the ends of the capillary with microfluidic tubes (0.2 mm diameter) pressed against the capillary openings. A complete schematic of this design is shown in Figure 16. Initial testing of this flow cell format initially demonstrated good acoustic focussing when compared with the custom approach detailed above due to the omission of acoustically absorbing materials such as tape and Perspex.




Figure 16 – Diagram of the mounted acoustophoretic flow cell comprising a 30x6 mm rectangular glass capillary bonded via glass pillars to a standard (75x25 mm) microscopy slide using CA adhesive. Fluid connectors are formed using heat-shrink tubing and ultrasonic excitation provided by a square PZT transducer bonded to one end of the capillary using epoxy (Epotek 301).

However, two issues with the flow cell prevented further integration into the overall system design. Firstly, the acoustic focussing, while initially proving successful, on further investigation had a strong lateral component causing unwanted movement of particles within the focus plane. In practice, this resulted in particles agglomerating into large balls, parts of which extended out

Chapter 3

of the depth of focus of the imaging system, and hence was severely deleterious to imaging quality and accuracy.

Secondly, the glass pillars were very challenging to consistently epoxy in place such that the capillary would be flat within the imaging plane. This resulted in decreasing accuracy of focus from one side of the image to the other, which again reduced image quality and, by making it impossible to adequately image both side of the flow cell simultaneously, reduced possible throughput. Due to these two issues the second flow cell design was not further investigated.

3.3.2 Analytical Investigation of Flow Cell Acoustic Properties (1-D transfer model)

To determine the appropriate frequency at which to drive the piezoelectric transducer (in order to excite half-wave resonance within the flow cell), it is possible to model the system using an equivalent circuit transducer model coupled to an acoustical impedance transfer model. The transducer is modelled as an equivalent circuit as described by Krimholtz, Leedom and Matthei ('KLM') (Krimholtz et al., 1970). The acoustical interaction between the transducer and the other layers of the system can be approximated as 1-D acoustical propagation problems (Hill et al., 2002). This approach avoids the significant complexity of deriving analytical solutions to the wave equation for a given acoustic resonator system (Sherrit et al., 1999). The circuit diagram for the transducer part of the model is presented in Figure 17.




Figure 17 – KLM model proposed by Krimholtz et al. The transducer crystal is modelled as an electric circuit consisting a transmission line, transformer and passive components.

Full details of the concept, formulation and analytical operation of the transducer equivalent circuit model can be found in the original paper by Krimholtz, Leedom and Matthei (Krimholtz et al., 1970), and the acoustical propagation equations are fully detailed by Hill, Shen and Hawkes (Kinsler et al., 1999). The MATLAB script used to carry out numerical simulation of the pressure amplitude within the flow cell within a range of frequencies is described by Glynn-Jones et al. (Glynn-Jones et al., 2010b).

The transfer model can be used to analyse the frequency response of a layered resonator and determine its acoustic focussing performance, providing estimates of the primary radiation force on particles of a given radius and acoustic contrast factor with the specified fluid medium. For the purposes of further modelling of the IFC's acoustofluidic flow cell, and understanding the different acoustophoretic performances across different sample characteristics, a study was carried using a parameter sweep across differing temperatures, T , and salinities, S , representing freshwater, brackish/estuary conditions and a range of open ocean values up to highly saline. The full results of this study are presented in Figure 20.

Chapter 3

In this experiment, values for T ranged from 5°C to 30°C, while S ranged from 0 *Practical Salinity Units* (PSU) to 39 PSU. Particles in the model were set to a uniform radius of 5 μm , density ρ of 1056 kgm^{-3} and sound velocity c of 1962 ms^{-1} . These values are average measured properties (Hill et al., 2010) of the PolySciences *Fluoresbrite* fluorescent polystyrene-latex beads which are used in later experimental device characterisation.

The density of the carrier fluid at each T & S combination were calculated according to the formulae of Fofonoff and Millard (*Fofonoff and Millard Jr, 1983*) and the sound velocities calculated according to Medwin (*Medwin, 1975*).

The calculated minimum and maximum densities and speed of sound of the carrier fluid is given below, and a full list of used model parameters can be found in Appendix A (Table 13, Table 14, Table 15 and Table 16).

$$\rho_{\min} = 995.7 \text{ kgm}^{-3} \text{ at } T = 30^\circ\text{C}, S = 0 \text{ PSU};$$

$$\rho_{\max} = 1030.9 \text{ kgm}^{-3} \text{ at } T = 5^\circ\text{C}, S = 39 \text{ PSU};$$

$$c_{\min} = 1425.7 \text{ ms}^{-1} \text{ at } T = 5^\circ\text{C}, S = 0 \text{ PSU};$$

$$c_{\max} = 1549.7 \text{ ms}^{-1} \text{ at } T = 30^\circ\text{C}, S = 39 \text{ PSU}.$$

The Acoustic Contrast Factor, Φ , between the modelled particles and carrier fluid, determined using Equation 3, was: $\Phi_{\min} = 0.276$ at $T = 5^\circ\text{C}, S = 39 \text{ PSU}$ and $\Phi_{\max} = 0.295$ at $T = 5^\circ\text{C}, S = 0 \text{ PSU}$.

Under every input condition of T and S , the model predicted two distinct frequencies between 1 and 3 MHz which result in half-wave resonances within the fluid layer; frequencies above 3 MHz resulted in more than one pressure anti-node within the fluid channel and so were ignored.

Figure 18 shows an output from the KLM model for the proposed flow cell with parameters $T = 30^\circ\text{C}, S = 39 \text{ PSU}$, (highly saline, hot seawater) which result in $\rho = 1025.1 \text{ kgm}^{-3}$ and $c = 1522.1 \text{ ms}^{-1}$ for the medium. This parameter combination results in resonances at 1.78 MHz and 2.06 MHz, with expected Acoustic Energy Densities (E_{ac}) of 103.31 Pa and 83.39 Pa respectively.

At each of these determined resonant frequencies, it is possible to model the acoustic energy distribution throughout the device as illustrated in Figure 18. Such modelling for this device, at each of the resonant frequencies under each *T&S* parameter combination, shows that none of the pressure amplitude plots have ‘turning points’ within the fluid channel. This is important as different layer arrangements can create areas close to the channel edges where the acoustic force actually pushes toward the edges rather than the pressure node as seen in Figure 19. The fact that these turning points are not observed in our modelling gives us confidence in the like performance of our novel flow cell design.




Figure 18 – Example plots of modelled Transducer Conductance (a) and Acoustic Energy Density (b) for a transfer model with carrier fluid parameters $\rho = 1024.7 \text{ kgm}^{-3}$ and $c = 1549.7 \text{ ms}^{-1}$. Two resonance frequencies are clearly defined at 1.78 MHz and 2.06 MHz. Figure plotted using MATLAB script detailed by Hill et al.

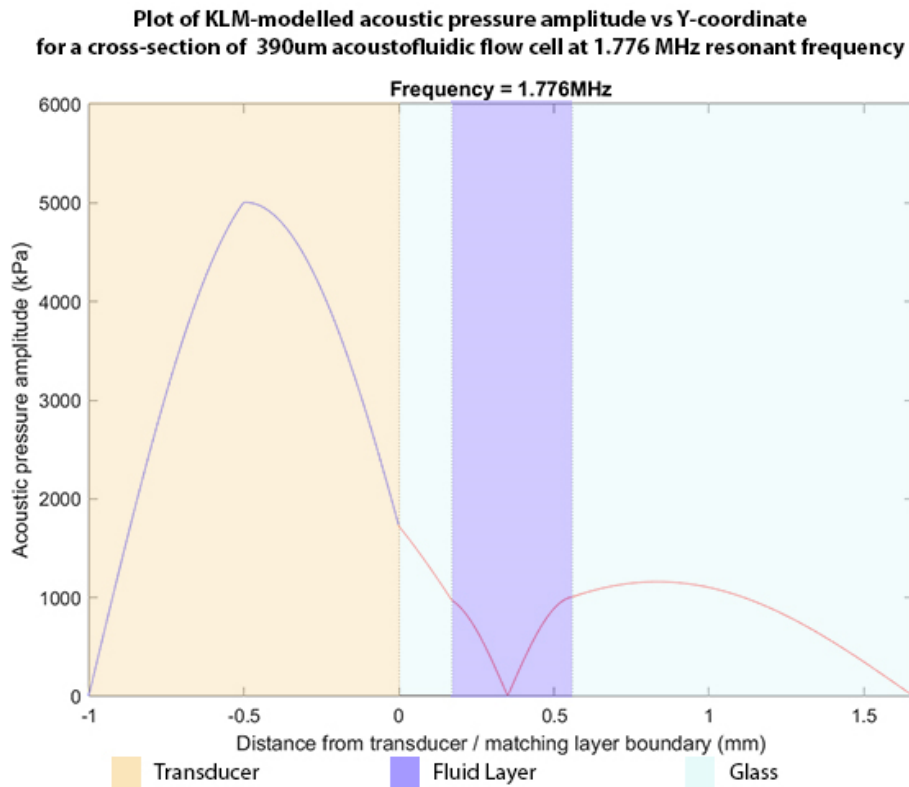


Figure 19 – Modelled Acoustic Pressure Amplitude (P) plot through vertical cross-section of the device at one of the half-wave resonant frequencies. Within the fluid layer (dark blue) the pressure has a single node, with particles above or below this point being subjected to forces pushing them toward the minimum, which is at the centre-height of this fluid layer. Figure generated using MATLAB script detailed by Hill et al.

The modelled plots of acoustic pressure amplitude for the two respective resonant frequencies have slight variations in the exact location in the fluid layer at which the pressure minimum forms. Such variations in the anti-node axial position necessitate the ability to finely adjust the imaging focus plane as part of system design. Therefore, in our IFC design there will not be a fixed distance between the flow cell and the objective, instead, apparatus enabling a manually adjustable focus distance will be integrated.

To explore the potential effects of varying carrier fluid properties on the frequency and strength of the resonances in the flow cell, the transfer model was analysed at each combination of fluid density and sound velocity. The two peaks in modelled acoustic energy density (i.e. Figure 18) were recorded (resonant frequency and acoustic energy density at the peak) and plotted for each temperature and salinity pair as seen in Figure 20.

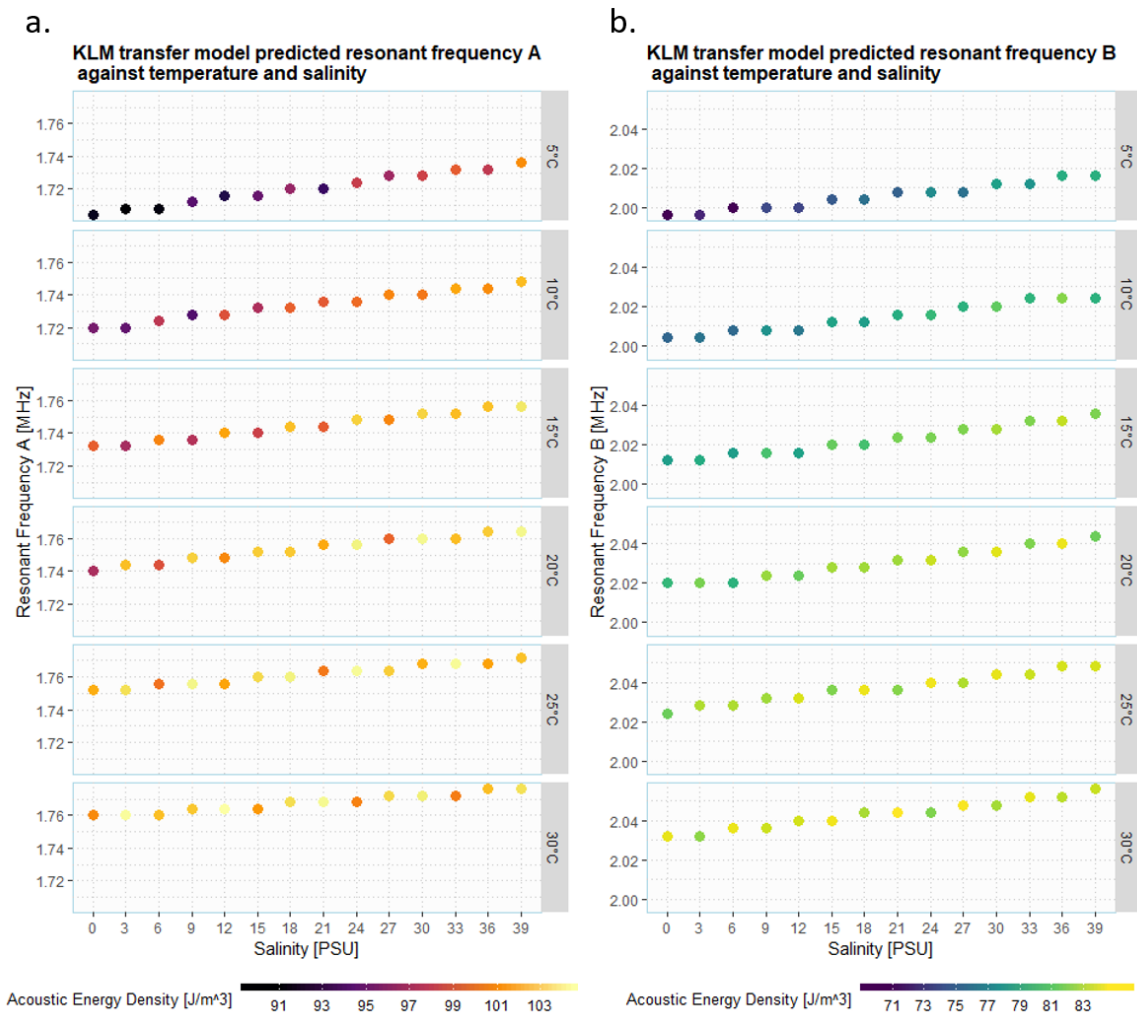


Figure 20 – Frequency (y-axis) and Acoustic Energy Density (colour) for the lower (a.) and upper (b.) resonant frequencies predicted by the transfer model at each combination of temperature (grouping variable) and salinity (x-axis) for the layered resonator flow cell. The graphs indicate rising resonant frequencies with both temperature and salinity at both resonant frequency A and B.

The data presented in Figure 20 show that there is a linear relationship between the output variables and the input parameters. Between the lowest temperature and salinity pair and the highest, the frequency of the lower resonance increased by approx. 80kHz and the upper by approx. 60kHz. The gradient of this increase in both frequency and energy density with increasing salinity appears to reduce as temperature increases.

As previous mentioned in Section 3.2, one approach previously validated within the literature to address the ‘drifting’ of the precise resonant frequency of the flow cell due to changing input

Chapter 3

conditions is to set the ultrasound generator to carry out a 'frequency sweep' around the target frequency. Previous work by Glynne-Jones et al. (2010a) and others has discussed this time-averaging of acoustic fields by using frequency modulation and shown it to be effective. The results of the exploration of different salinity and temperature on acoustic energy density and resonant frequency suggest that our IFC will require a wide enough sweep in the driving frequency to ensure resonance is achieved even when analysing significantly different water samples, especially *in situ*, where temperatures can be more variable.

In order to validate that a frequency sweep would appropriately focus cells in later validation work, the KLM model was run again with parameters set to explore the effect of a 20 kHz frequency sweep on the position of the acoustic pressure node within the fluid layer. 20 kHz was chosen for sweep width as the results presented in Figure 20 demonstrate an approximately 20kHz shift in resonant frequency with a temperature change of 5 °C, which is a reasonable maximum temperature change expected in a controlled laboratory environment. The effect of salinity on resonant frequency is smaller, and salinity is not expected to change when taking samples from a monoculture grown in an incubator, without additional inputs of dissolved salts.

This frequency sweep experiment was carried out using constant KLM parameters for the carrier fluid of $\rho = 1024.7 \text{ kgm}^{-3}$ and $c = 1549.7 \text{ ms}^{-1}$, the same as for Figures 18 and 19. The depth of the acoustic pressure node within the fluid layer was determined by KLM modelling using $\pm 10 \text{ kHz}$ around the two previously determined resonant frequencies of 1.776 and 2.056 MHz, and the results are presented below in Table 3. This experiment shows that at frequencies $\pm 10 \text{ kHz}$ around the resonant frequency, the variation in depth of the acoustic pressure node within the fluid layer is only 2.3 μm , suggesting that a 20 kHz frequency sweep around the chosen frequency will be experimentally acceptable when using the real device in a lab environment and will not cause cells to leave the plane of optical focus. For use in the field, the range of temperature and salinities likely to be experienced by the device would need to be carefully considered such that a frequency sweep would cover all likely resonant frequencies without unacceptably moving the pressure node within the device and causing a degradation in imaging quality.

Table 3 – KLM-modelled acoustic pressure node depth within the flow cell for various driving frequencies representing a 20 kHz sweep with centre frequencies of the two determined resonant frequencies (bold) for the flow cell geometry - 1.776 MHz and 2.056 MHz.

Frequency / MHz	Pressure Node Depth / μm
1.766	210.6
1.776	209.4
1.786	208.3
2.046	214.9
2.056	216.1
2.066	217.2

The transfer model has several limitations which necessitate experimental characterisation of layered resonators. Firstly, the model takes several parameters which are difficult if not impossible to accurately determine. The loss factor, Q , of the different layers is not known, as are the exact thicknesses of the glue and electrode layers. Without accurate inputs for these values the frequency response predicted is likely to significantly vary from the response of the physical system.

Secondly, the transfer model drastically oversimplifies the behaviour of the resonator and reduces the acoustic wave to a 1-dimensional, uniform field. In reality, transducers do not produce spatially homogenous acoustic pressure across their surfaces and have 'acoustic hot spots' that are particularly strong compared with the mean acoustic energy density. This is not an issue for particularly large transducers as the acoustic energy density of the transducer varies sinusoidally and will average out over time.

Chapter 3

The model also does not represent the lateral modes of the standing wave which occur in real systems. Structural aspects of layered resonators generate complex acoustic pressure fields which do not have a uniform minimum at the nodal plane. Instead, variations in the pressure at this plane will move suspended particles laterally and can lead to effects such as acoustic streaming and particle agglomeration. It is highly desirable to minimise these effects in the IFC system as agglomeration of phytoplankton cells would impede accurate imaging and analysis.

Finally, the Acoustic Contrast factors used in the transfer model varied between 0.276 and 0.295. While these values correspond to the ACF of the 10 μm beads used for the acoustic characterisation, they are significantly higher than the values of 0.01 to 0.07 measured in 4 phytoplankton strains by Hincapié Gómez et al. (2018) in a study which sought to experimentally characterise the acoustic properties of microalgae. The radius of the particle being focussed is the highest determinant of the acoustic force acting upon it, being the cubed term in the acoustic radiation force Equation 2, but this force also scales linearly with ACF. As such, the beads should experience between 4 and 30 times more acoustic radiation force than phytoplankton in the same system and it will be important to understand this difference when characterising focus quality with beads.

3.3.3 Finite Element Modelling of Flow Cell Acoustic Focussing (COMSOL)

Before moving to an experimental characterisation of the fabricated device, the system was modelled in *COMSOL Multiphysics* in order to address the uncertainty of focussing performance for different phytoplankton properties. By simulating the flow channel and acoustic forces on a range of particles of different radii and acoustic contrast factors, at a range of flow rates, sensible bounds on the likely performance of the device can be determined more easily than would be possible experimentally. A full list of variables and parameters used in the FEM model are given in Appendix B, Table 17 and Table 18.

The geometry of the COMSOL model consisted of a cross-section of the part of the flow cell to which the transducer is bonded (5 cm length x 390 μm height x 4 mm width). Only the first 5 cm of the length was modelled as that is the area in which focussing occurs. The mesh used for finite element modelling (FEM) was a triangular mesh of minimum dimension 1 μm . To determine the trajectories of particles at a range of starting positions as they enter the flow cell, 1,000

(independent) particles were uniformly distributed along the leftmost boundary. The model predicts their path under the action of the acoustic force, flow, and viscous drag. Once particles reach the outlet, their vertical separation distance is recorded. A diagram of the behaviour of this model is given in Figure 21, with a schematic illustrating the axes given in Figure 22.

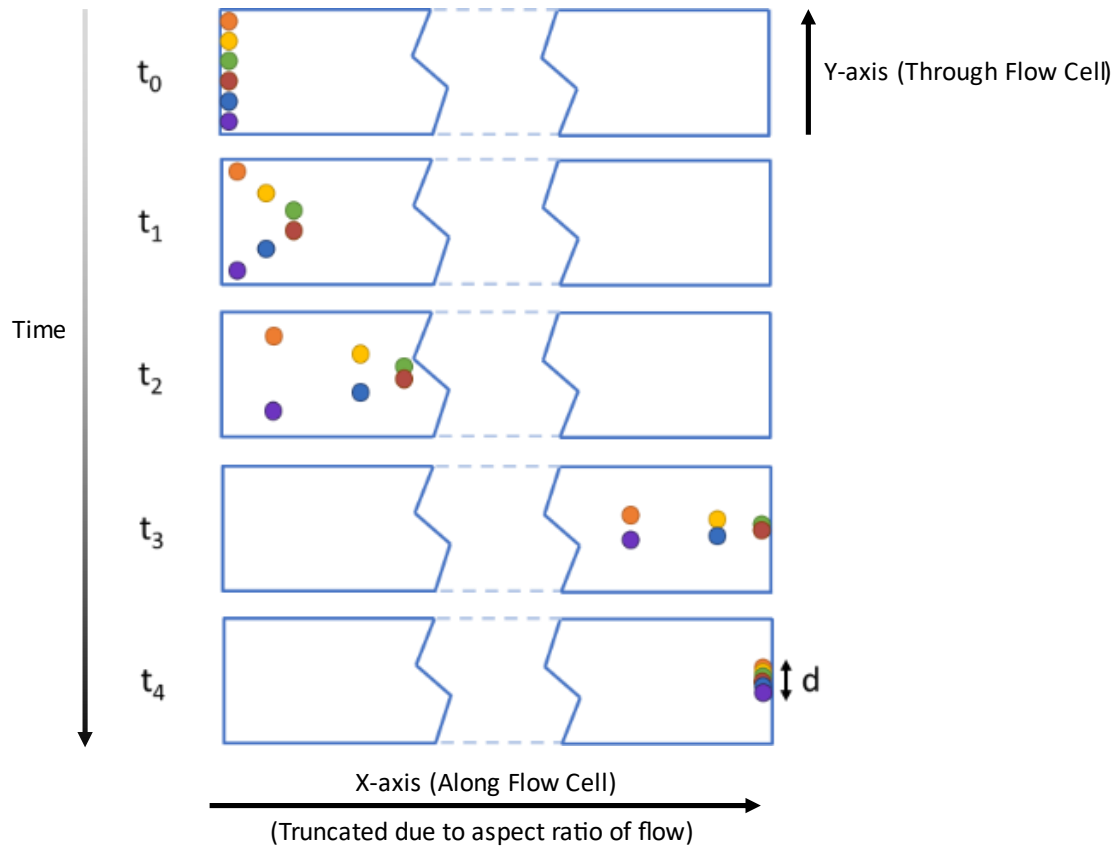


Figure 21 – Diagram of the Finite Element Model of the flow cell cross-section showing the progression of beads (coloured) along the device with time. At the particle inlet the beads are uniformly distributed, and they are focussed to the central pressure anti-node as they travel along the x-direction. Once particles have reached the outlet at the right side, their vertical separation distance, d , is recorded.




Figure 22 - Schematic of the axes of the fluid layer used for the FEM modelling within this section, showing the position objective lens relative to the fluid flow direction.

The range of particle starting positions was variable such that particles would never be closer to the top and bottom boundaries of the flow cell than their radius, in order to prevent particles intersecting the glass sides. Thus, where $r = \text{particle radius}$, the particle inlet was set to span from $0 \mu\text{m} + r$ to $390 \mu\text{m} - r$ on the Y axis.

Fluid flow within the flow cell behaves according to *Planar Poiseuille Flow*. This flow condition exists where an incompressible, Newtonian fluid is in steady laminar flow between two stationary plates forming a channel of constant cross-section which is substantially longer than its cross-sectional area. The equation for the velocity of the flow with respect to y , the vertical position within the flow channel, is given by Equation 5, the *Hagen-Poiseuille Equation for Planar Poiseuille flow*.

$$u(y) [ms^{-1}] = \left(\frac{H^2}{2\eta}\right)\left(-\frac{dP}{dx}\right)\left(\frac{y}{H}\left(1 - \frac{y}{H}\right)\right)$$

Equation 5 – Formula for linear (along-flow) velocity of fluid under Planar Poiseuille Flow with respect to vertical position (y). Where: H is the separation of the two planes/walls

Chapter 3

[m] (height of the flow channel); η is the dynamic fluid viscosity [Pa·s] and $\frac{dP}{dx}$ is the pressure gradient along the direction of flow [Nm^{-3}].

The pressure gradient in the flow cell is given by Equation 6.

$$-\frac{dP}{dx} [Nm^{-3}] = \left(\frac{Q \cdot 12\eta}{w}\right) \left(\frac{1}{H^3}\right)$$

Equation 6 – Formula for pressure gradient within the flow channel, where Q = volumetric flow rate [m^3/s], η is the fluid's dynamic viscosity [Pa·s], w is the width of the flow cell [m] and H is the height of the flow cell [m].

The flow velocity is highest at the centre of the channel and tends to zero at the flow cell boundaries. The 1-dimensional profile of the resulting fluid velocity profile across the height of the modelled flow channel is given in Figure 23.

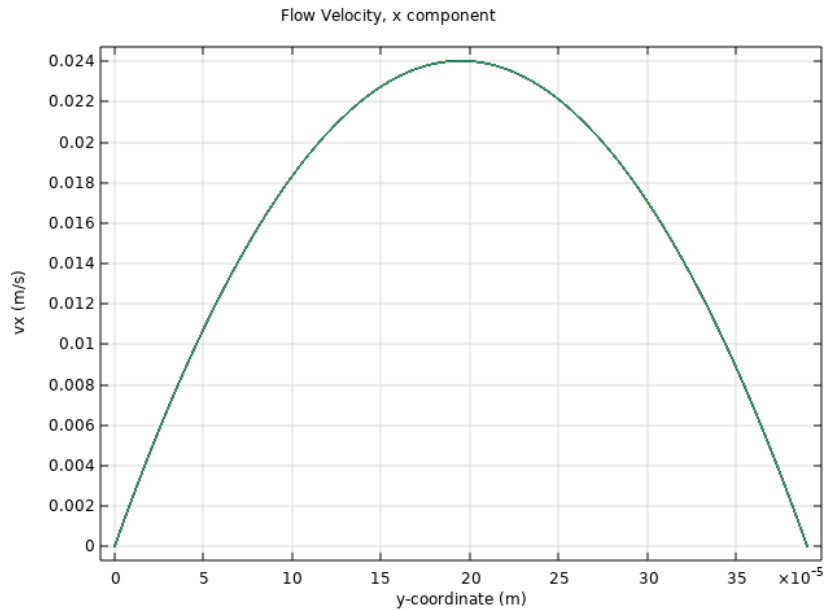


Figure 23 – 1D plot of the x-component (v_x) of the fluid flow velocity profile within the vertical axis of the flow channel as described by Planar Poiseuille Flow. The given flow velocities given were calculated at a volumetric flow rate of 1.5 ml/min, but the parabolic distribution of velocity with y coordinate remains the same for all flow rates. The flow velocity at the centre of the channel is the highest and the velocity toward the channel boundaries (0 and 390 μm - the top and bottom of the fluid layer) tends to zero.

The acoustic radiation force (Equation 2) across the height of the flow cell is similarly implemented in the model. The resonant frequency was given as $\frac{c_0}{2H}$, where c_0 = sound velocity of the fluid and H = height of the flow cell. While the transfer model used in the previous section models the behaviour of the entire layered resonator system, including the air/glass boundary and epoxy bonding transducer to the flow cell, the FEM analysis considers only the fluid layer. Therefore, by specifying the frequency as $\frac{c_0}{2H}$, a pressure anti-node is always formed halfway through the fluid layer's vertical profile and the more complex transducer and structural acoustical behaviours included in the KLM model can be ignored. This is appropriate for the purpose of the FEM analysis as it is the ability of the resonant frequency to focus particles rather than the particular frequency to use (in the real system) that is of interest.

The force profile (in Newtons) given by the idealised acoustic radiation force (assuming a half-wave bounded by perfectly reflecting walls) is illustrated in Figure 24 (1D vertical profile). As

Chapter 3

expected, the acoustic radiation force at the centre of the channel ($y = 1.95 \times 10^{-6}$ m) equals zero and has two maxima at the quarter and three-quarter heights in the fluid layer. The force also tends to zero at the top and bottom boundaries, which suggests that particles situated there will not experience any force toward the central plane and therefore never be focussed. In practice, however, the particles will always be separated from the flow cell walls by at least their radius, which will ensure they experience an increasing acoustic force, however small initially.

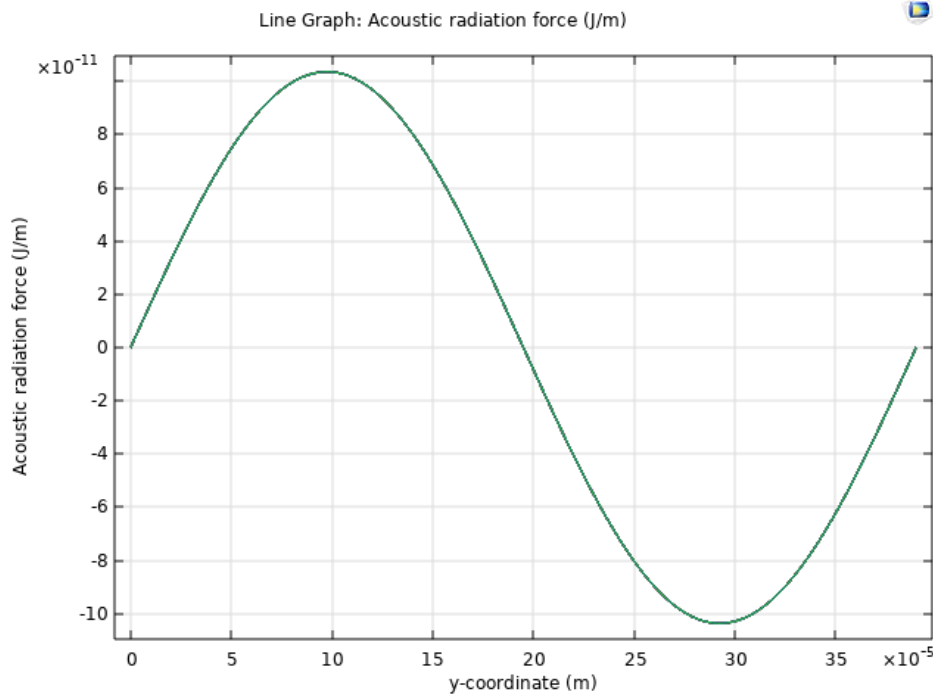


Figure 24 – Plot of the acoustic radiation force against vertical position (y-coordinate) as determined by the ultrasonic standing wave in the flow cell. The pressure minimum (zero newtons) is located at the half-depth of the fluid layer (195 μm) and also tends to zero at the upper and lower boundaries. Since particles are never located exactly at the boundary (they are separated by at least a distance equal to their radius), even particles very close to the flow cell walls will experience some force toward the centre.

The model was set to simulate 150 seconds or until a stop condition was met. This stop condition was set such that the outlet (rightmost boundary) of the flow cell acted as a particle counter and stopped the simulation when all 1000 particles reached the other end of the fluid layer. As such, at higher flow rates, COMSOL only needed to simulate the system until the flow velocity had

Chapter 3

carried all particles through, and particles would not continue to be focussed by the acoustic radiation force after they had left the 5cm focussing region.

A parameter sweep was established such that COMSOL would simulate the focussing of particles of between 0.5 and 10 μm radius (r), of acoustic contrast factors (φ) of between 0.01 and 0.1, and at volumetric flow rates (Q) between 50 $\mu\text{L}/\text{min}$ and 10 mL/min . Each combination of r , φ and Q had a full 1000-particle simulation run: for 39 radii, 10 acoustic contrast factors and 24 flow rates. Thus, a total of 9360 simulations were completed.

For each parameter combination, the vertical position of particles at the outlet was saved and the range (maximum minus minimum y coordinate) calculated. For each combination of φ and r , the maximum volumetric flow rate that resulted in a vertical separation range of $\leq 5 \mu\text{m}$ was recorded. 5 μm was selected as a conservative measure of focus – if the centres of the particles are within 5 microns of each other, they will be within the depth of focus of the standard microscope objectives planned for the physical system.

The results of this analysis are presented in Figure 25, which can be described (following Equation 2) as graphs of the shape $y = x^3$. From the graph it is apparent that size is the main determinant of whether particles adequately focus within the 5cm focussing region. For objects of a contrast factor of 0.05, having a 10 μm radius will allow focussing even at very high flow rates of 9 mL/min , whereas at 0.5 μm radius, particles will not focus within 5 cm even at 50 $\mu\text{L}/\text{min}$. Similarly, the ability to focus strongly depends the acoustic contrast factor, with a 5 μm particle being focussed under volumetric flow rates of between 0.05 and 1 mL/min depending on its φ value.

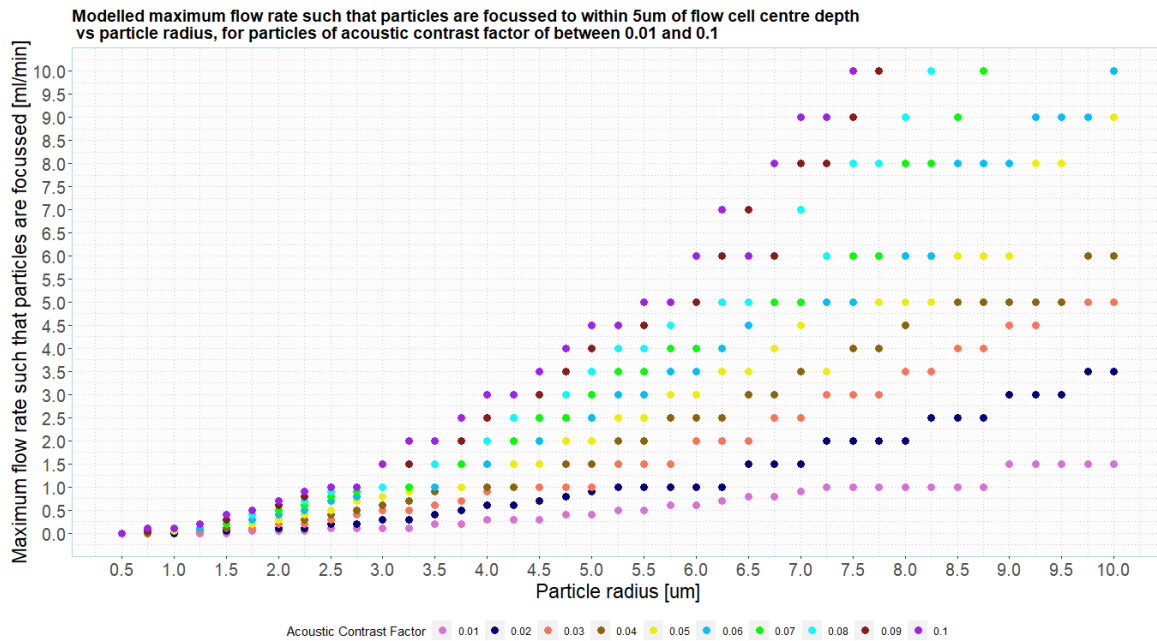


Figure 25 – Plot of the maximum volumetric flow rate (Y axis, ml/min) under which particles of a given radius (X axis, μ m) focus within the 5cm flow cell analysed by FEM, grouped by acoustic contrast factor (colour). There is significant overlap between different data points, especially at 10 ml/min flow rate where particles of $\phi \geq 0.06$ with large diameters all focus.

The results of this analysis suggest that it is important to either know the approximate expected acoustic contrast factor and radii of the phytoplankton within a sample, or, to limit flow rate such that even small plankton with low contrast factors can be imaged within a sample. For samples of unknown contrast and radius distributions, there is therefore a tradeoff in experimental design between throughput and the minimum size and contrast that will be focussed. When compared with the commercial instruments already used by biological oceanographers (Table 2), even the minimum flow rate investigated (0.5 mL/min), at which even the 1 μ m radius, 0.01 ϕ cells are adequately focussed, is competitive.

The sampling rates given for the commercial instruments describe the speed with which all cells in the fluid can be imaged. In the IFC presented here, not all cells in the sample fluid will be imaged as the objective lens will not have a field of view covering the entire width of the flow

Chapter 3

cell. Therefore, to comment on the sampling rate of our device it will be important to consider the imaged volume as well as the total sample volume, as these will not be the same.

The above modelling does however have significant limitations which reduce its predictive power for the performance of physical flow cells. In real devices, losses arise from the acoustical interactions between the layered components of the flow cell, all of which have different acoustical impedances and internal sound velocities. These losses could significantly lower the radiation force experienced by particles within the fluid layer and increase the time taken to focus them to the imaging plane. Therefore, modelling alone will not be suitable to assess device performance and an experimental approach will also be required which will be discussed in the following section.

As previously mentioned, due to their internal crystal structures, real transducers have complex 3-dimensional behaviours resulting in 'hot spots' on the transducer/flow cell interface. These can cause result in lateral sections of the fluid layer having weaker acoustic forces. At higher flow rates these effects average out to a constant force on particles but at lower horizontal velocities particles experience non-homogenous forces in the focussing region. Due to the velocity profile of particles in flow described by Equation 6 (and illustrated in Figure 23), even at higher volumetric flow rates the particles closest to the flow cell boundaries may have a horizontal velocity low enough to be affected by the transducer heterogeneity described. Viscous losses in the acoustic boundary layers can cause acoustical streaming, unwanted axial movements of particles which can result in agglomeration.

The exact axial location of the pressure node toward which particles experience a force may not be the half-width of the flow cell where the flow velocity is highest. While in the FEM simulation the frequency is set such that there is always a node corresponding to the axial height of maximum horizontal velocity, in a real device, the different resonant frequencies excitable often produce nodes above or below the centre of the flow channel. The KLM modelling work in Section 3.2.2 illustrates this point, with the 1.776 and 2.056 MHz resonant frequencies predicted in Figure 18 producing pressure minimums slightly above and slightly below the fluid layer half-height.

Chapter 3

The resonant frequency of a flow cell depends on the temperature of the fluid and of the device itself, but heat is generated by driving the transducer, which can lead to a shift in resonant frequency over time. This will be more pronounced at low flow rates, where the cooling effect of water movement through the device will be lower.

The above shortcomings mean that the experimental characterisation of the layered flow cells is likely to have significantly poorer focussing performance than suggested by modelling. The next section will detail an attempt to quantitatively investigate the performance of the flow cells experimentally, using multiple methods to assess acoustic characteristics.

3.3.4 Acoustic Focus Analysis – Experimental

Two methods were employed to experimentally determine and characterise the resonant frequencies of the fabricated flow cells. The first made use of the *C60 Impedance/Frequency Response Network Analyser* hardware produced by Cypher Instruments. The secondary approach involved using the flow cell in the IFC device, imaging fluorescent microspheres in flow, and then analysing the images to determine the velocity field within the fluid medium.

The first method allows the determination of the resonant frequencies of any given acoustical device by plotting the frequency response of the transducer under realistic use conditions. This is important to understand as it allows the appropriate setting-up of the device to successfully acoustically focus cells. The second method investigates the particle focussing performance of the flow cell under various flow conditions while driven at resonance. Combined, these methods attempt to provide both information on what frequencies should be used to drive the acoustic standing wave and how effective the respective frequencies are for the purpose of focussing phytoplankton cells to the imaging plane of the flow cell.

Method 1 – Admittance Analysis

The C60 is a network analyser which can measure the frequency response of electro-acoustic systems within the 10 Hz to 4 MHz range. By connecting the flow cell transducer to the C60 output and running a frequency sweep within the range suggested by the above modelling (1 to 3 MHz), the device can be set to measure and plot the transducer admittance ($Y = 1/Z$).

Chapter 3

In order to use the C60 to determine resonant frequencies of a given transducer, a plot of admittance was generated in the 1-3 MHz range under 4 different flow cell conditions: filled with only air (no fluid introduced); filled with deionized water; filled with Lugol's Iodine- preserved *Isochrysis galbana*; filled with F/2 marine media with *Fluoresbrite* fluorescent polystyrene-latex beads (representing the presence of phytoplankton cells in a marine sample). A plot for one flow cell (Device A) is given in Figure 26.

Chapter 3

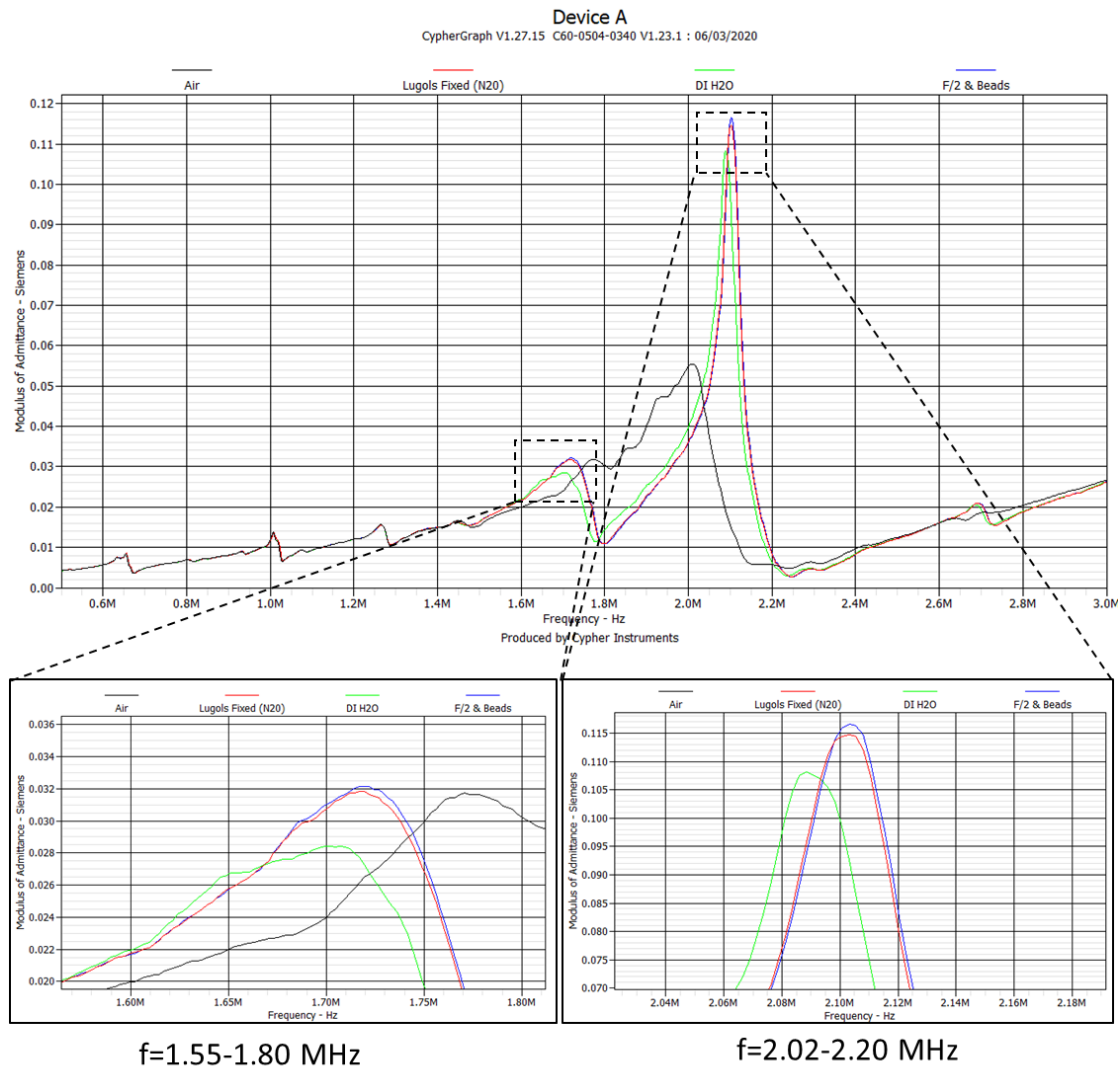


Figure 26 – C60 network analyser plot of admittance (1 over impedance) of the ultrasonic transducer within the 0.5 – 3 MHz range. Measurements were conducted over 4 different conditions; unfilled (Air - black), filled with deionized water (DI H2O - green); filled with marine media and PolySciences Fluoresbrite fluorescent polystyrene-latex beads (blue); filled with Lugol’s Iodine-fixed plankton sample (sample number N20 - red). Insets: peaks around 1.70MHz; 2.10 MHz demonstrating the slight differences in acoustic performance when the device is filled with media of different salinities (and hence sound velocities and densities).

As mentioned above, Figure 26 also shows the flow cell’s frequency response when filled with a Lugol’s Iodine-fixed phytoplankton sample (sample number N20). This sample was of *Isochrysis galbana*, a unicellular haptophyte phytoplankton with a cellular diameter of 5-10 μm , and was analysed in order to determine any frequency shift when using a preserved plankton sample. This is because the ‘fixing’ of phytoplankton cells using Lugol’s solution is common in biological

Chapter 3

oceanography and it is important to understand any effect the presence of such fixing agents on the acoustical properties of the flow cell. The very close agreement between the measured acoustic properties of the preserved sample (red) and the sample with polystyrene beads in growth medium provides good confidence that the devices' performance will be similar in both preserved and live cell samples.

Resonant frequencies can be determined from the C60 frequency sweep output by looking for admittance spikes, which occur when the flow cell is filled with fluid but not when it is filled with air. Eight flow cells were analysed in this manner and the resulting resonances, as seen in Figure 26, were around 1.7 MHz and 2.1 MHz. When filled with deionized water the flow cell admittance spikes were frequency-shifted by around 50 kHz and were slightly weakened; admittance profiles for F/2 media- and Lugol's preserved sample-filled flow cells were highly similar in frequency and magnitude of peaks. Both measured peaks were of maximum widths below 50 kHz. Each liquid used in characterisation was allowed to come to room temperature before injection.

Method 2 – Image Analysis

A second method for characterising the acoustical performance of the flow cell was formulated: this involves taking images of particles within the flow cell under various flow conditions while the transducer is driven at a resonant frequency. However as seen below, this method ultimately proved unsatisfactory.

If the ultrasonic half-wave resonance generated within the flow cell has its pressure node in a plane which is vertically halfway through the fluid layer, particles in flow which are focussed to this pressure node cell should have a normal distribution of vertical positions, with the standard deviation of this distribution giving a measure of the focussing performance (similar to the separation distances modelled in Section 3.3.3). This is because the carrier fluid within the flow cell acts according to Poiseuille flow with a normal distribution of linear (flow) velocity through the vertical profile of the flow cell, so suspended particles will also move with normally distributed linear velocities.

Chapter 3

A mismatch between the centre of the flow cell (which has maximum linear flow velocity), and the acoustic pressure node, would cause a skewness of the distribution of particle velocities, with the magnitude and direction of skew giving some indication of the position of the pressure node relative to the point of maximum flow velocity.

By imaging the same particles twice with a short, known time delay between captures, their velocities can be determined. Using this approach, work in this section was carried out with the aim of determining the performance of the acoustic focussing by determining the distribution of particles' travel distances within the time delay.

An objective of this characterisation was to measure focussing performance under realistic flow-rate conditions. At a volumetric flow rate of 2 mL/min, the COMSOL model predicts particle velocities exceeding 0.03 ms^{-1} , which at 10X magnification, where the field of view of the used camera is $1125 \times 703 \text{ }\mu\text{m}$, would result in a particle crossing the imaging region in under 30 ms. For faster flow rates and higher magnification, the framerate of the camera will not be sufficient to capture multiple exposures of the same particle; as such, a different approach is required.

In order to image the same particles multiple times, and hence calculate separation distances in a given time period, a LED pulsed illumination system capable of microsecond pulses was used, as specified in Section 3.2. By setting the LEDs to activate twice within a single exposure of the camera with a short delay between lighting pulses, each object will be imaged twice in each frame. An unprocessed example frame taken using this technique is presented in Figure 27.




Figure 27 – Example double-exposure image of 10 μm beads taken with a 2 ms delay between illumination pulses. Each of the 6 beads present in the field of view is imaged twice in the single frame such that their velocities can be calculated from the distance they travel during the 5 ms delay.

This method was carried out on 100 frames at each flow rate, with two illumination pulses of 15 μs separated by a 2 ms delay. In order to generate a representative linear velocity distribution with sufficient statistical significance, the separation distances for a large number of beads must be measured. To conduct this analysis and avoid the bias a manual estimate of bead centre (e.g. using *ImageJ*) would incur, a MATLAB script was written which, for a given IFC image, detects and numbers beads and localises their centroid pixel. The user can then input which bead numbers are in fact the same bead, and the script automatically records the distance (and angle) of separation between the two exposures, labelling this distance on the image.

Chapter 3

Figure 28 and Figure 29 show the post-detection and post-labelling steps of this process, respectively. Pseudocode for the MATLAB script is presented in Appendix C (entire script published online at

https://github.com/ajwl27/IFC_BeadImageAnalysis/blob/main/analyse_bead_images.m).




Figure 28 – Example of detected beads in the same image as Figure 26. The MATLAB script detects beads using standard image processing techniques (code in Appendix A.3) and assigns a number to each detection. Detection 11 in this image is a false positive and can be ignored by the user in the next step; labelling the exposure pairs which result from the same physical bead.


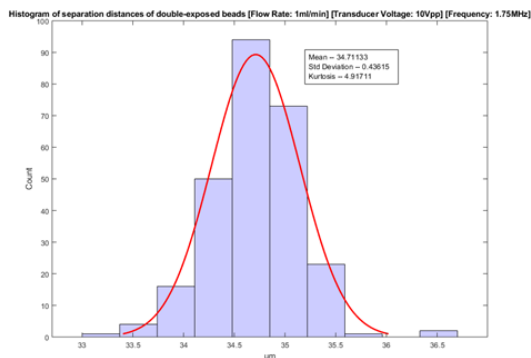


Figure 29 – After detection, the user enters the numbers of the exposures relating to the same physical bead (e.g. 1 and 2; 3 and 4). The MATLAB script measures the distance between the centroid pixels of the two labelled beads and the angle (in degrees from vertical), stores this data and presents it as an overlay on the image.

After the experimental work was completed and each set of 100 frames analysed in the described manner, histograms of the separation distances (converted to micrometres based on the scaling factor for a 20X objective - 5.7971 pixels : 1 μ m) were plotted. The mean and standard deviation of the data were calculated, and the results are shown in Figures 29, 30 and 32. Fewer experimental setups than would be optimal were tested due to the coronavirus situation limiting lab use.


Exposure-separation distributions for $f=1.75\text{MHz}$; $V=10\text{Vpp}$

a. 1.0 ml/min



Mean = 34.71 μm
SD = 0.44 μm
SD as % of Mean = 1.27%

b. 2.0 ml/min



Mean = 66.13 μm
SD = 2.35 μm
SD as % of Mean = 3.55%

Figure 30 - Results of bead-separation experiment for measuring acoustic focus quality in layered resonator flow cell with resonant frequency = 1.75 MHz, transducer voltage = 10 V peak-peak. Volumetric flow rate = 1.0 ml/min (a) and 2.0 ml/min (b).

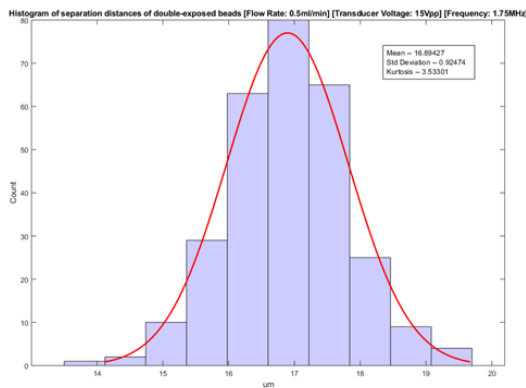
Figure 30 shows the clearest difference in distribution between an experimental condition under which the beads were (qualitatively) focussed well (a., 1.0 mL/min) vs a condition where the focus was poor (b., 2.0 mL/min). The standard deviation of separation distances increased from 0.44 μm to 2.35 μm , which divided by the mean separation distances for each distribution corresponds to a 1.27% and 3.55% SD respectively. This widening of the distribution aligns well with the expected behaviour when moving from in-focus bead images to out-of-focus bead images as described above.

Chapter 3

It is possible to conclude from Figure 31 and Figure 33, however, that the method used is not sufficiently sensitive. With the volumetric flow rate set to 3.0 mL/min and the transducer driven at 1.75 MHz, 15 V peak-peak, the acoustic focus quality was qualitatively significantly inferior to that at the slower flow rate of 1 mL/min, but this is not reflected in the plotted normal distributions. The qualitative difference in focus quality is illustrated in Figure 32 which shows a side-by-side comparison of frames taken at the two flow rates used in the below graphs.


Exposure-separation distributions for $f=1.75\text{MHz}$; $V=15\text{Vpp}$

a. 0.5 ml/min




Mean = 16.89 μm
SD = 0.92 μm
SD as % of Mean = 5.45%

b. 1.0 ml/min




Mean = 34.55 μm
SD = 0.44 μm
SD as % of Mean = 1.27%

c. 2.0 ml/min



Mean = 63.61 μm
SD = 2.46 μm
SD as % of Mean = 3.87%

d. 3.0 ml/min



Mean = 98.86 μm
SD = 1.91 μm
SD as % of Mean = 1.93%

Figure 31 - Results of bead-separation experiment for measuring acoustic focus quality in layered resonator flow cell with resonant frequency = 1.75 MHz, transducer voltage = 15 V peak-peak. Volumetric flow rate = 0.5 ml/min (a), 1.0 ml/min (b) 2.0 ml/min (c) and 3.0 ml/min (d).




Figure 32 – Qualitative assessment of the acoustic focus quality at volumetric flow rates of 1.0 ml/min and 3.0 ml/min, with the same transducer settings of $f=1.75$ MHz, $V=15$ V peak-peak.

The source of this apparent error is most likely to be methodological; the technique of determining focus distance relies on an accurate measurement of the exact central pixel of each bead, but this is extremely difficult to consistently locate either manually or using the MATLAB script. Neither the image processing algorithm used or any attempt to manually localise the centre of each bead will be accurate enough to draw statistically reliable distributions of measurements between the two exposures. It is possible, however, that more sophisticated image processing algorithms would more accurately locate the centre of the beads and thus more accurately measure the inter-exposure travel distance.

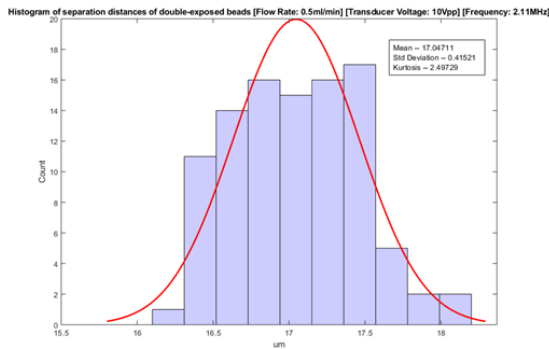
As previously mentioned, one other reason that the described method is unable to accurately determine acoustic focus performance may be that the acoustic pressure node is not coincident with the half-depth of the fluid layer, and thus the maximum of the parabolic flow profile. The result of this mismatch would be a larger variation in flow velocities for a small acoustic focussing error than if the two were coincident. This phenomenon would also cause a reduction in the closeness of the beads' speed distribution the normal distribution. Therefore, the standard deviations reported, which are those of the gaussian curves fit to the data, may be an inappropriate metric to use as a comparison of performance under different conditions.

Chapter 3

The methodological failure discussed above repeats at $f=2.11$ MHz, $V=10$ Vpp, as demonstrated in Figure 33. The reduction in standard deviation as a percentage of the mean between the distance distributions at 0.5 mL/min and 1.0 mL/min (33a. and 33b.) would not be expected if using an accurate metric of focus quality, and is not backed up by qualitative analysis of frames under the two different flow rates. The increased skewness to the right of the mean of each distribution at 2.11 MHz compared with the results from the earlier 1.75 MHz experiment suggests that the pressure node excited at 2.11 MHz is further from the fluid layer half-depth than that excited at 1.75 MHz.


Exposure-separation distributions for $f=2.11\text{MHz}$; $V=10\text{Vpp}$

a. 0.5 ml/min




Mean = 17.05 μm
SD = 0.42 μm
SD as % of Mean = 2.46%

b. 1.0 ml/min



Mean = 33.21 μm
SD = 0.44 μm
SD as % of Mean = 1.32%

c. 2.0 ml/min



Mean = 65.32 μm
SD = 2.54 μm
SD as % of Mean = 3.89%

Figure 33 - Results of bead-separation experiment for measuring acoustic focus quality in layered resonator flow cell with resonant frequency = 2.11 MHz, transducer voltage = 10 V peak-peak. Volumetric flow rate = 0.5 ml/min (a), 1.0 ml/min (b) and 2.0 ml/min (c)

Figure 34 presents the inter-exposure separation distance for the same driving frequency as in Figure 33 but the higher voltage of 15 volts peak to peak (Vpp). At this higher voltage, *acoustic streaming* was observed within the flow cell, with beads forming two single-file lines within the field of view. This is clearly undesirable as beads can agglomerate and there is a reduction of throughput and imaging quality. Other volumetric flow rates were not investigated as slower horizontal velocity would only increase these streaming effects; and higher velocity would result in insufficient axial (desirable) focussing.

Exposure-separation distributions for $f=2.11\text{MHz}$; $V=15\text{Vpp}$

a. 2.0ml/min




Figure 34 - Results of bead-separation experiment for measuring acoustic focus quality in layered resonator flow cell with resonant frequency = 2.11 MHz, transducer voltage = 15 V peak-peak. Volumetric flow rate = 2.0 ml/min.

As discussed, the experimental approach taken here has unfortunately proved incapable of accurately measuring the acoustic focussing performance of beads within the flow cell. For the purposes of the remainder of the work within thesis, however, it will be sufficient to know whether or not the acoustic focussing is successfully bringing cells into the optical focus plane for sharp imaging. A quick subjective test of the sharpness of objects within the IFC's images provides a high level of confidence that those objects have been successfully acoustically focussed, as shown in Figure 32. This is because of the relationship between the objective lens and the depth of the flow cell; typically, objectives have a depth of focus of less than $10\ \mu\text{m}$, and the flow cell is $390\ \mu\text{m}$ deep. Since objects are randomly distributed throughout the depth of the flow cell when acoustic focussing is not operational, objects regularly appear out of focus, as demonstrated in Figure 32b. If none of the objects appear out of focus, it can reliably be inferred that the acoustic focussing is operating successfully.

Although this subjective test does not help to qualify the level of performance of the focussing, it can consistently be used to check that it is working sufficiently for imaging. If further flow cell development was to be carried out associated with more precise determination of acoustic

Chapter 3

focussing efficiency, however, future experimental protocols could be used to test the acoustic focus performance by calculating the sharpness of the edge of a high-contrast object like the polystyrene beads used here. This technique would not work for organic cells, which are highly translucent, but could be used to quantify the flow cell performance in greater detail. The issue with parabolic flow however remains; even with a new imaging technique it may not be possible to use the distance travelled to during the inter-exposure period as a means to measure acoustic focus without precisely determining the height of the acoustic pressure node within the flow cell and calculate the effect of varying flow velocities.

3.4 Imaging Performance Characterisation

3.4.1 Introduction

The image resolution of an IFC is a key parameter as it determines the information content of the acquired images. With higher resolution images, both human taxonomists and automated image processing systems will be able to more accurately perform taxonomic classification, cell measurement, etc. Other metrics of image quality such as distortion and contrast, while harder to measure objectively, also contribute to the human interpretability of images and can furthermore affect automated techniques such as neural networks. To measure the suitability of the IFC developed here for various research applications, it is therefore necessary to have a *quantitative* measure of imaging quality.

The goal for this work will be to state the minimum phytoplankton cell size which can reliably be imaged by the device. This property, i.e. the smallest size of object that can still be seen to be separate from other objects in the image, is typically referred to as the *resolution* (Wu et al., 2008, Dobbins, 2000).

Therefore, in this section, a critical analysis of a quantitative approach to testing the image resolution of the acoustically focussed IFC will be presented, using a standard inverted microscope to develop a testing protocol which can later be applied to the IFC (In Section 3.5.5). By first quantifying the effects of the variables within the system on the resolution, including exposure, rotation of the flow cell and the accuracy with which the objective lens is focussed, it will be possible to apply the measurement procedure with confidence to the completed IFC. These effects of various manually controlled parameters on the measurement of resolution have not received much attention within the literature but are important to understand in our IFC.

A formal definition of the resolution of a microscope system is the smallest spatial period of a still-resolved object (Born and Wolf, 1999). As mentioned in the review of current and future trends in phytoplankton ecology (Section 2.3), cells are expected to decrease in size over time due primarily to warming of their habitats. This highlights the necessity for IFCs to have the highest possible resolution, such that the smallest cells can be imaged. As demonstrated by the

Chapter 3

modelling work in Section 3.3, the acoustophoretic focussing within the flow cell is expected to successfully acoustically levitate even cells under 5 μm diameter, therefore, the optical performance of the IFC becomes a crucial factor in the minimum size of cells possible to be analysed.

Typical characterisation of the resolution of custom microscope devices often involves manually imaging a test target (e.g. 1951 USAF test target) and subjectively determining the smallest features resolved (Greenbaum et al., 2013, Kheireddine et al., 2019, Vainrub, 2008). However, this method is inherently subjective and makes it difficult to compare competing IFC systems. A more quantitative measure can be derived from the frequency spectrums of recorded digital images by calculating the highest spatial frequency present (avoiding random noise) with non-zero amplitude (Pospíšil et al., 2017).

This can practically be achieved by taking the 2D Fourier transform of the image, plotting circularly-averaged amplitude values and from these determining the frequency at which the signal is indistinguishable from background noise (*cut-off frequency*), as demonstrated by Vainrub (2008) and depicted in Figure 35. The main issue arising during this process is the necessity of semi-subjective determination of the cut-off frequency, as in most images there is no clear boundary where the Fourier amplitudes of the image and noise are distinct. Furthermore, the frequency response of the image depends in part on the shape of the object being imaged (Dobbins, 2000). In the following section, therefore, an exploration of a quantitative approach to the measurement of imaging resolution will be presented.

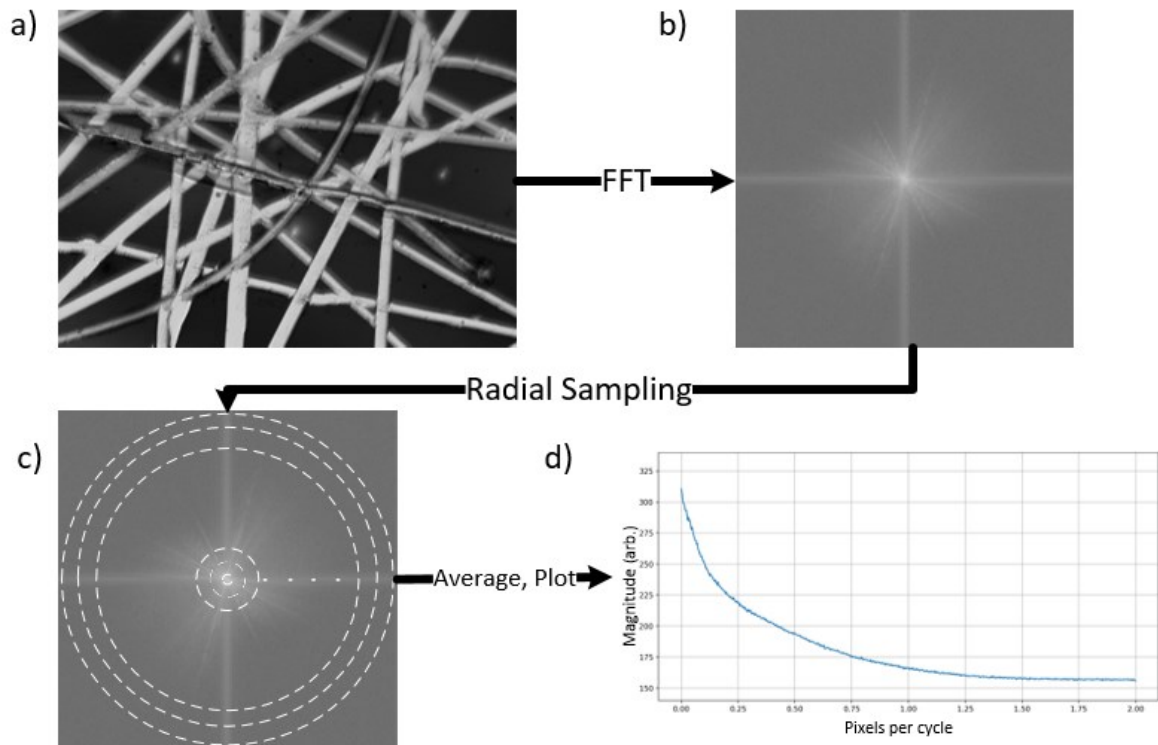


Figure 35 – Process for determining resolution using the method inspired by Vainrub (2008). 13a) a test target with high-frequency content (made by scratching an ink blot on a standard 75x25mm slide with a razor) is imaged using the system under test. 13b) a 2-dimensional Fast Fourier Transform (FFT) is applied to the image to convert from spatial to frequency domains. 13c) a predetermined number of radial samples are taken from the centre of the FFT image outward to the edges, where all pixels in the given radius are sampled (as pixels have integer addresses, bilinear interpolation is used). 13d) the values corresponding to each radius are averaged and plotted.

3.4.2 Modulation Transfer Function

The quantitative measurement of a device’s spatial frequency response, the *Modulation Transfer Function* (MTF), can be performed to assess the imaging resolution of a complete optical system (Xie et al., 2018b). To calculate the MTF of a digital imaging device, an edge (a very sharp transition from black to white) is imaged onto the sensor and the pixel values perpendicular to this edge are recorded, these values being referred to as the Edge Spread Function (ESF). The first derivative of these values are the Line Spread Function (LSF); MTF is simply the 2-dimensional Fourier transform of the LSF (Estribeau and Magnan, 2004), as depicted in Figure 36. MTF is

Chapter 3

typically measured in either Hertz, cycles or lines per pixel, where a cycle is one light and one dark line.




Figure 36 – Schematic of the procedure for determining the resolution of a digital imaging device by calculating the Modulation Transfer Function (MTF) using the Slanted Edge Method (ISO 12233). Values are sampled from pixels perpendicularly to an imaged slanted edge, making up the Edge Spread Function (ESF). The ESF is differentiated to produce a Line Spread Function (LSF). Finally, the MTF is yielded from the 2-dimensional Fourier transform of the LSF. Adapted from Estriebeau and Magnan (2003) and Vhengani, Griffith and Lysko (2012).

To avoid undersampling of the edge due to pixel size and to account for phase shift dependence, standard MTF measurement procedure involves imaging an edge which is slightly slanted with respect to the rows and columns of the imaging sensor (Burns et al., 2000, Vhengani et al., 2012). Simulated (Estriebeau and Magnan, 2004) and experimental (Xie et al., 2018c) results show that MTF curves produced using the slanted edge technique are repeatable and close to the theoretical MTF of a system when slant angles are small (5-10° from vertical) but that increased angles reduce the accuracy and reproducibility of measurements. Furthermore, the slanted edge protocol has been shown to be a reliable measure of resolution even for CMOS sensors with complex pixel topologies (Estriebeau and Magnan, 2004) despite being developed for square CCD pixels.

With the aim of developing a robust way to measure MTF of the completed IFC system including optics, camera and processing software, the following sections will detail the exploration of the

effect of various variables on measured MTF using a commercial microscope. They will present novel experimental work to assess the influence of various parameters on the measured MTF in work which will not only allow us to confidently measure MTF in our system but contribute to the field of microscopy as a whole.

3.4.3 Development of an experimental protocol for IFC resolution characterisation

The typical process used for practically determining the MTF of a microscope and our IFC system is based on the standard *slanted edge technique* described in ISO 2233 (ISO, 2017). Briefly, the edge of a black square on a transmission microscopy target slide (*R1L355P, Thorlabs Inc., USA*) is imaged with the camera rotated with respect to the slide holder in order to ensure adequate slant angle for analysis. Next the *ImageJ* (Schneider et al., 2012) plugin *Slanted Edge MTF* (Mitja et al., 2011) used to generate MTF against spatial frequency curves from the data. Values of spatial frequency where MTF drops to 50% (*MTF50*) and 10% (*MTF10*) of the low-frequency maximum MTF are extracted to allow straightforward intercomparison and to be used to estimate the minimum object size that can be imaged with sufficient contrast.

In order to quantify the effects of varying illumination techniques, exposure times, focus accuracies and edge slants, all of which can be varied in the IFC, a range of experiments were conducted using a digital camera (*ORCA-ER, Hamamatsu*) on an inverted microscope (*IX71, Olympus*) with a 10X objective (*UPlan FL N, NA=0.30, Olympus*) in brightfield configuration. These experiments are detailed in Sections 3.4.4, 3.4.5 & 3.4.6.

3.4.4 Experimental determination of the effect of camera exposure on measured system MTF

An important variable within an IFC system is the amount of light received by the camera during its exposure. This can be controlled either by changing the illumination intensity or by increasing the exposure time. Although changing the illumination intensity is typically more complex as it requires control hardware (e.g. Pulse-width modulation for LEDs), it will be easy to change the illumination pulse length in our planned stroboscopic illumination system. Since varying the exposure time with a fixed illumination brightness has the same effect as varying the illumination intensity with a fixed exposure time, as long as the imaging target is not moving, the former was chosen in this experiment due to the simplicity of implementation on a standard microscope.

Thus, a slanted edge MTF measurement was carried out with fixed focus position, lighting intensity, diaphragm apertures, edge rotation (to 5° as recommended in ISO2233) and X/Y stage positions, while varying camera exposure time from 1 to 15 ms in 1 ms increments. The MTF curves resulting from this procedure are shown in Figure 37.




Figure 37 – MTF curves derived from images of a slanted edge on a resolution test target (R1L3S5P, Thorlabs Inc., USA) under camera exposure conditions from 1 ms to 15 ms. All images were taken under the same Z-axis focus position and with the same edge slant angle (5 °). Here, the blue curves (MTF from images taken with 1 ms – 7 ms exposure duration), which are highly consistent, are generated from images without any oversaturated pixels. The red curves (MTF from images taken with 8 ms – 15 ms exposure duration) are from images with pixel values of 0 or 4095 (i.e. fully saturated). These are erroneously high and abnormally smooth, indicating the importance of using a nonsaturating exposure time when measuring and comparing system resolving power.

With the faster exposures from 1 to 7 ms, none of the pixels in the recorded images were saturated. Conversely, in all the exposures longer than 8 ms, the white half of the image had many oversaturated pixels, where these pixels had reached their *full-well capacity* (Hasinoff, 2014) and were unable to store any further irradiance data. Figure 38 illustrates this concept of saturation by comparing cropped regions from images taken with a) 5 ms exposure and b) 10 ms exposure and displaying pixel intensity values from the white sections of each image.




Figure 38 - Cropped regions from images taken at a)5ms exposure and b)10ms exposure in the 3.2.1 Exposure experiment. At 5 ms exposure, pixels in the white side are unsaturated and still contain information; at 10 ms exposure, the white side is overexposed and the camera is limited by the full-well capacity of its pixels. All pixels in the white section of this image have pixel intensity values of 4095, the limit of the camera. MTF curves generated from data with such saturated pixels are erroneously high and inflate the measured system resolution to unrealistic levels.

The resulting MTF curves calculated from the recorded images in Figure 37 show that when the camera was not fully saturated (blue lines), the MTFs showed great consistency, but that when images from the system were saturated (red lines), MTF results were erroneously high, leading to unrealistic resolution measurements.

Complimenting this finding, Figure 39 presents the MTF10 and MTF50 metrics across the tested exposure range, clearly showing significant elevation in these resolution metrics in the saturated images. With MTF10 values above the Nyquist sampling limit of 0.5 cycles per pixel (a minimum of 2 pixels per cycle is required to avoid aliasing), these figures are patently unrealistic.

Chapter 3

When applying the MTF measurement technique to our custom IFC, therefore, it will be ensured that while carrying out the slanted edge experiment, pixels within the image are not fully saturated. This will practically be achieved by selecting an appropriate combination of illumination intensity and exposure time, which may be different to those tested above on a microscope due to the sensitivity of the attached camera, and optical properties of the IFC.




Figure 39 – Graph of MTF10 (blue) and MTF50 (orange) values extracted from the MTF curves for images with exposure times ranging from 1 ms to 15 ms.

Both metrics are consistent for exposure times of 1 to 7 ms, where images do not have any over- or undersaturated pixels. Exposures of 8ms and longer (red dashed box) have erroneously elevated MTF10 and MTF50 metrics than at faster exposures, indicating the inflating effect of saturated pixels on measured system resolution. MTF10 values which are above the sampling limit imposed by the Nyquist theorem (a minimum of 2 pixels per cycle required to avoid aliasing) are unrealistic.

3.4.5 Experimental determination of the effect of edge slant rotation on measured system MTF

The next variable of interest while measuring the MTF of the IFC will be the rotation of the slanted edge target. Therefore, an experiment is carried out to determine the effect of the rotation angle on the measured MTF.

For this experiment the effect of edge slant angle on derived MTF curves was determined by fixing microscope X,Y and Z positions, illumination intensity and camera exposure (5 ms, which was previously determined to avoid saturated pixels and hence result in reliable MTF curves which do not exceed the theoretical Nyquist resolution limit).

Starting with a measured edge rotation angle of 0° with respect to the columns of the sensor's pixels (the edge completely vertical), the camera was iteratively rotated about its central axis by 2.5° , the edge image captured and MTF curve calculated. Previous studies have measured MTF using only a small number of rotation angles (e.g Xie et al. (2018c) tested 6 different rotation angles, of which half were rotated over 25° from the rotation angle which resulted in the optimal MTF curve). The aim of this experiment is to examine with finer detail the effect of rotations which are closer to the optimal theoretical rotations of $0-10^\circ$. Figure 40 displays the MTF curves measured in this experiment and demonstrates the angle-dependence of MTF measurements when rotation angle is increased past 10° .

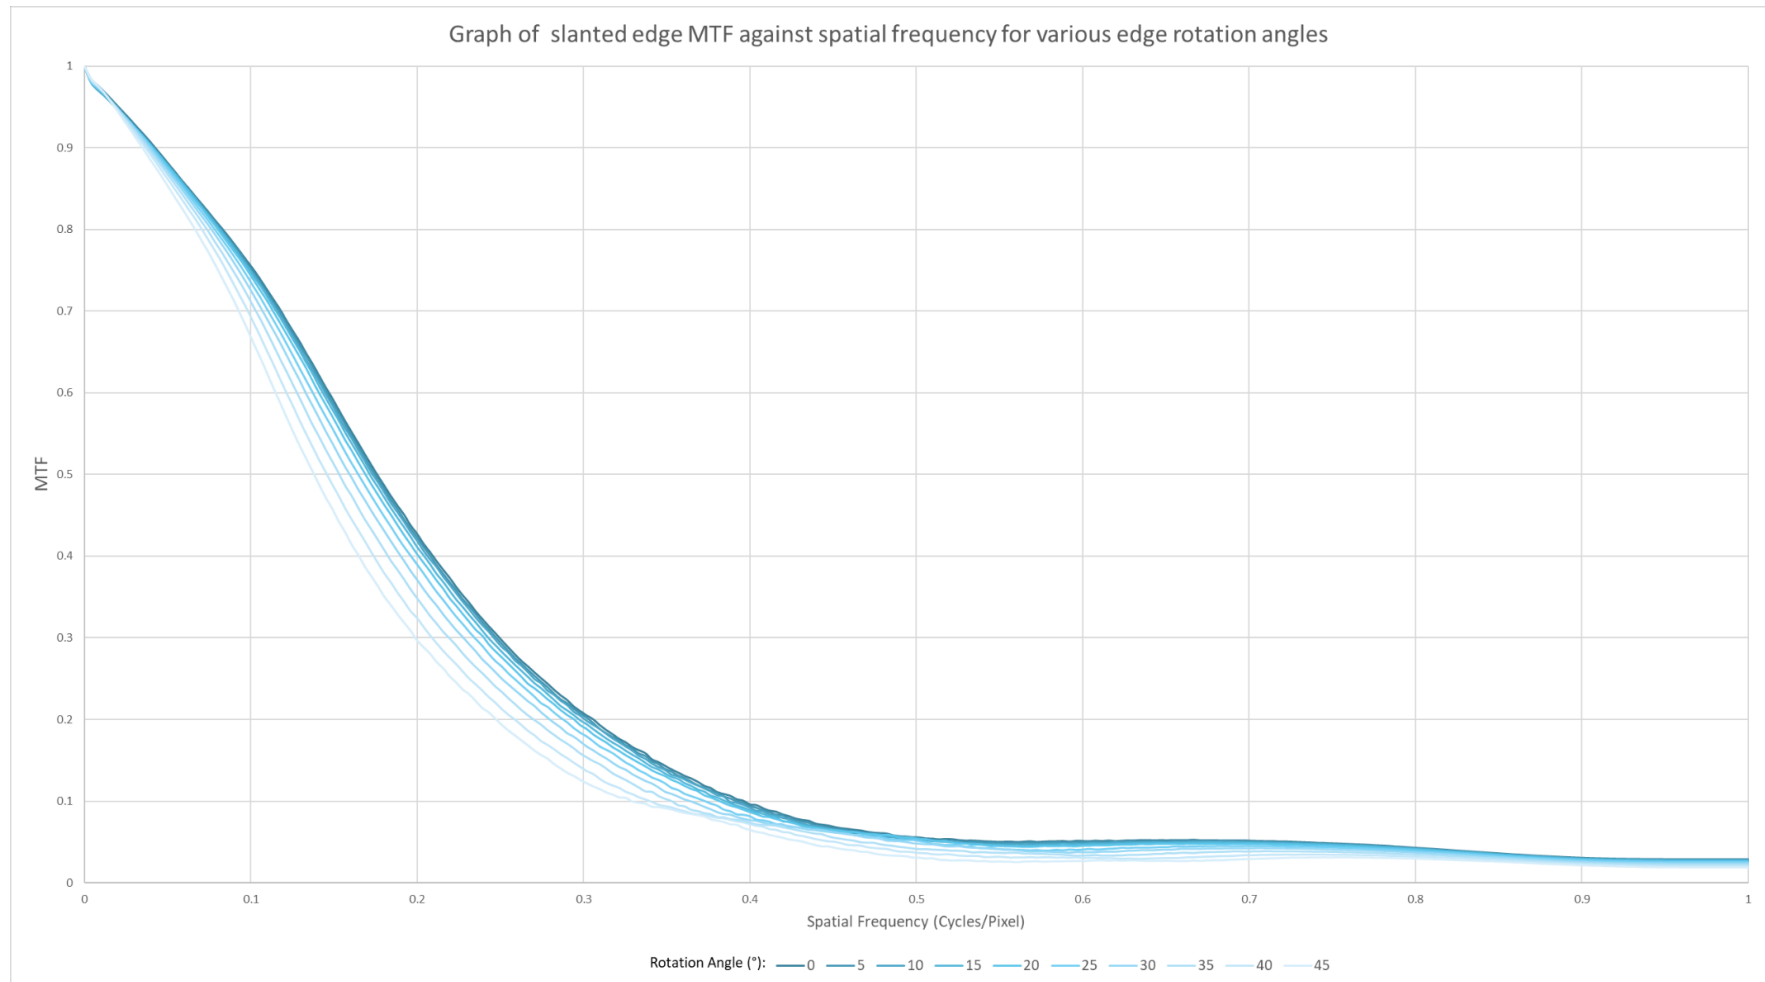


Figure 40 – Slanted edge MTF curves generated for incrementally increasing edge slant angles, starting from 0° (completely vertical) and incrementing by 5° up to 45°.

Chapter 3

Figure 40 presents the MTF curves measured at each edge 5° increment in rotation angle, overlaid onto the same axes. 2.5° increments were not included in this graph in order to reduce complexity. Good agreement in measured MTF is demonstrated for 0° to 10° rotation, but increasing past this point results in degraded areas under the curve with an angle-dependent worsening of MTF. This result confirms the finding of Estribeau and Magnan (2003) and Xie et al (2008) that MTF measurements using the slanted edge technique require a small rotation angle of not more than 10° in order to accurately assess system resolution.

Figure 41 shows the extracted values of MTF10 and MTF50 values across all the tested edge slant angles in all 2.5° increments, again showing angle-dependence of these resolution metrics when the slant is more than 10°. This finding confirms the work of Estribeau and Magnan (2003), and Xie et al. (2018), who find that only edge slant angles in the 0-10° range yield consistent MTF curves. Our experiment additionally provides an experimental description of the rate at which MTF is expected to fall off as edge rotation increases; Figure 41 shows that rotations of up to 20° result in measured MTF10 and MTF50 values that are within 95% of the best measured values and that angles of up to 32.5° result in measured MTF10 and MTF50 values that are within 90% of the best measured values.

When applying the MTF measurement protocol on our custom IFC, care will be taken to ensure that edge slant angles are in the range of 0-10°. This will allow the most reliable measurement of MTF possible and result in the best estimate of the minimum phytoplankton cell size resolvable by the completed system.




Figure 41 - Graph of MTF10 (dark blue) and MTF50 (light blue) values extracted from the MTF curves calculated for images with edge slant angles between 0° (completely vertical; aligned with sensor pixel columns) and 45°, in 2.5° increments. With edge rotation angles of 0-10° the determined MTF10 and MTF50 values are very consistent, but with increasing slant angle these resolution metrics display angle-dependence.

3.4.6 Experimental determination of the effect of focus accuracy on measured system MTF

The next variable of interest when measuring the MTF of our IFC will be the quality of focus of the microscope objective onto the slanted edge imaging target. Changing the focus of a microscope objective moves the imaging plane vertically through the target, and to obtain the best possible image quality of a given object, the objective should be focussed so that the object plane coincides with the imaging plane. This ensures that the light rays from the object are optimally collected and focused by the microscope objective, resulting in a sharp and detailed image. When the object plane and imaging plane are properly aligned, it maximizes the resolution and contrast. As previously discussed, our IFC will have a variable focus objective, so it is important to be able to assess the impact of focus accuracy on image resolution in order to accurately assess the minimum resolvable phytoplankton size.

In this third experiment, therefore, the effect of focus accuracy (in the Z-axis) on measured system MTF is investigated. Illumination intensity, X/Y stage position, sensor exposure time (5ms, as above) and slant angle (5°) will all be fixed, and the variable will be the Z-axis focus position, which will be increased in 1 μm increments from a focus position well below the test target, to one well above. In this manner, the sensitivity of the MTF measurement procedure to focus accuracy is tested. An experiment which describes the effect of focus accuracy on MTF has not yet been presented in the literature; authors simply state that their devices were appropriately focussed. Our experiments help to quantify the sensitivity of the measurement protocol to various focus conditions, a determination which is especially important in instruments without precise and repeatable, electronically-controlled focus of their imaging lenses, as will be the case for our IFC.

As illustrated by Figure 42, it is extremely challenging to subjectively determine the quality of focus position beyond an accuracy of $\pm 2 \mu\text{m}$. This reduces the repeatability of focussing of an objective onto an object and creates the risk of inaccurately measuring optimal MTF values.

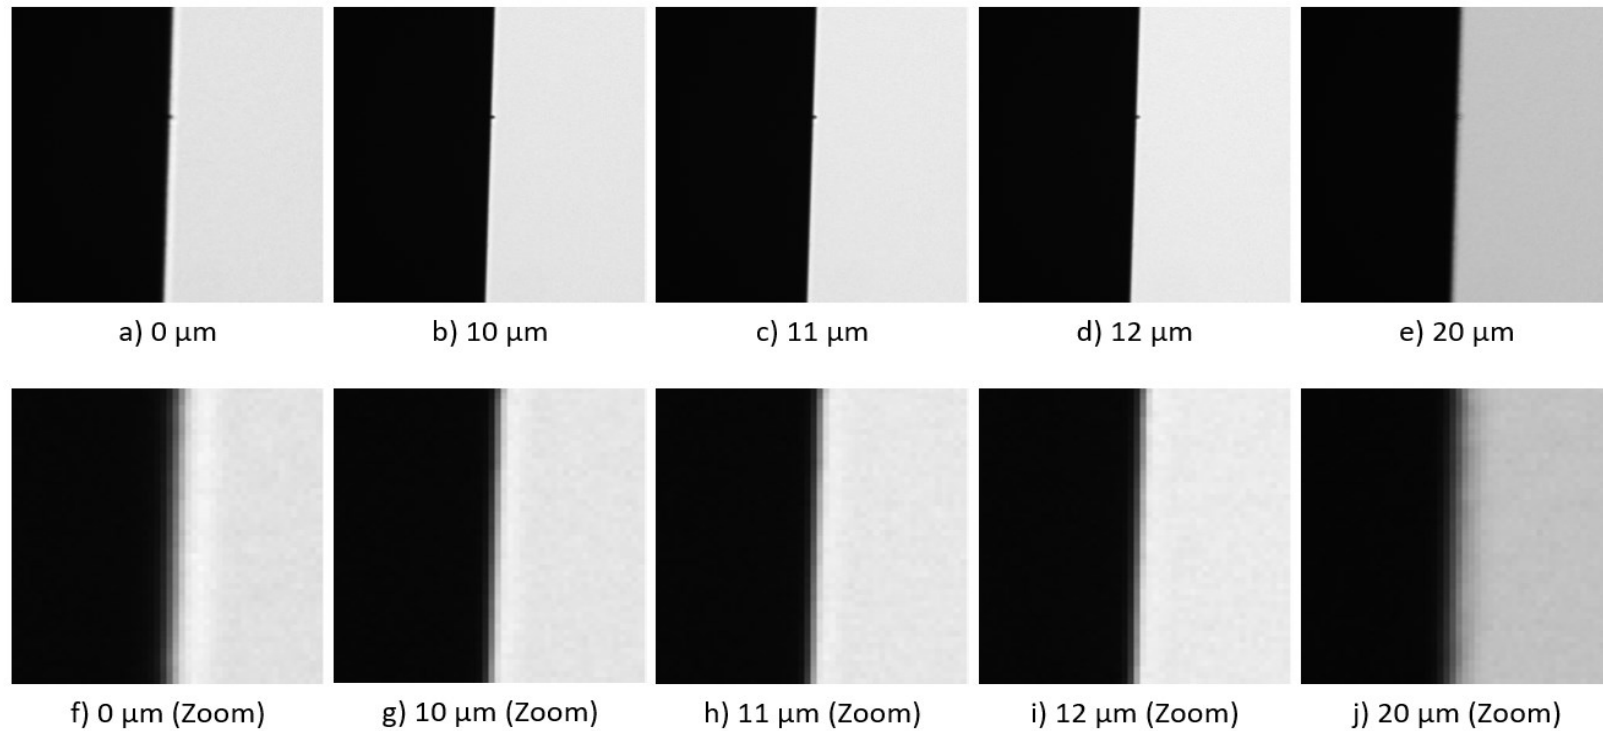


Figure 42 – Comparison of the a)-e) central 256x256 pixels and f)-j) central 50x50 pixels of the slanted edge images taken during the section 3.2.3 focus experiment. Image a) is focused below the target and image e) above the target, as can be more clearly determined from the zoomed images f) and j). Images b), c) and d) are all subjectively ‘in focus’ and it is almost impossible to manually determine which has the most accurate focus position. As can be seen in Figure 42, there are subtle differences in derived resolution metrics between the three central in-focus images. Of particular interest is the fact that image e) has a darker white section than the others, which is reflected in the lower MTF metrics on the above-target focus positions seen in Figure 42. This is caused by the interaction of the Köhler illumination focus plane and the imaging focus plane of the objective.

Chapter 3

Figure 43 demonstrates, by comparing MTF10 and MTF50 across all of the experimentally tested focus depths, that there are subtle differences in measured resolution between images that seem identical to a human user. This is a finding which has significant implications for the use of this method in instruments like our IFC which do not have discrete focus settings, e.g. traditional microscopes where the focus of the objective is controlled by manually tuning a continuously rotating knob. The same graph also shows that reliance on either MTF10 or MTF50 alone may be unsatisfactory due to nonlinear relationships between the two metrics. This finding also does not appear to have been noted within the literature. While MTF10 more reliably determines the maximum spatial resolution resolvable, MTF50 may be useful for determining overall image quality. Therefore, when measuring the resolution of our IFC system, both MTF10 and MTF50 values will be used to estimate minimum cell sizes that could be imaged.




Figure 43 – Plot of slanted edge MTF10 (blue) and MTF50 (orange) metrics for increasing Z-axis focus positions. A focus point well below the target was selected and Z-height incrementally increased, with an image of the slanted edge taken every 1 micron step. Images were thus taken with the edge in focus (8-13 µm from initial position) and with the focus position above (0-7 µm from initial position) and below (14-20 µm from initial position) the target. The graph shows a nonlinear coupling between MTF10 and MTF50 values, highlighting the necessity of not solely relying on one or the other for resolution characterisation. Further, the plot demonstrates that Z-position accuracy of at least ± 2 µm is required to obtain high and consistent MTF measurements.

As previously discussed, the custom IFC will not make use of the same precise, discrete stage focussing mechanism found in the microscope tested within these experiments. This is because these components would significantly increase the cost of the system, which as previously stated is intended to be accessible by resource-constrained laboratories. Lower-cost microscopes often make use of continuous focus mechanisms, such as rotating a focus knob or focus ring, to adjust the position of the objective lens relative to the specimen. While these mechanisms may lack the precision of their high-end electronic counterparts, they can still achieve good focus quality when used properly.

In our custom IFC, a *focus ring* will be employed to enable users to manually adjust the focus of the imaging system. This cost-effective solution is sufficient for achieving the desired focus quality, as the user can carefully and iteratively adjust the focus until the image appears sharp and well-defined. This approach allows for acceptable focus accuracy within the practical limitations of human perception and the system's intended use.

Despite the potential for slightly reduced precision and repeatability in the focus adjustment, the focus ring system provides a balance between affordability and performance, making it a suitable choice for resource-constrained laboratories. In accordance with our design goals stated in Section 3.2, this design choice will ensure that a wider range of users can access and benefit from the custom IFC, while still achieving satisfactory imaging quality.

The implication of this design choice is that MTF measurement of the IFC will not be as precise as for the research-grade microscope used in this section, and that when measuring the imaging performance it is possible that the optimal (maximal area under the curve) MTF curve for the IFC will not be found. This complicates the direct comparison of the measured MTF of our IFC with those of more sophisticated microscopy hardware, as the results may be influenced by the less precise focus adjustment mechanism. For this reason, it is likely that a direct comparison of the measured MTF₁₀ and MTF₅₀ values between the IFC and the microscope used for these experiments will demonstrate a lower imaging performance of the IFC.

Chapter 3

In order to attempt to mitigate these limitations, when measuring the MTF of the IFC, several images of the slanted edge target will be captured while varying focus conditions to account for inconsistencies in focus quality. By selecting the image resulting the highest measurements of MTF10 and MTF50, a more accurate representation of the IFC's optimal imaging performance can be obtained.

However, it is important to consider that the primary goal of the custom IFC is to provide a cost-effective and accessible platform for phytoplankton analysis, rather than competing directly with high-end microscopy systems in terms of optical performance. The primary design objective is to achieve an adequate level of image quality such that the smallest possible phytoplankton can be accurately identified.

It is therefore planned to focus on using the MTF10 and MTF50 values to determine the minimum phytoplankton size which will be resolvable by the IFC, rather than directly comparing the performance with existing instruments. These MTF values can be used to find the limiting resolution of the IFC, which is critical for determining its ability to accurately identify and differentiate various phytoplankton species based on their size and morphological features. By establishing a minimum resolvable size, users can better understand the IFC's suitability for their research applications and ensure that the acquired images provide sufficient information for accurate identification and analysis of phytoplankton populations. This approach will ultimately help to assess the overall effectiveness and utility of the IFC in the context of its intended use, rather than focusing solely on its performance relative to high-end microscopy systems. It will also allow us to state a minimum cell size for the IFC, similarly to those presented for the commercial instruments introduced in the previous chapter.

Our work to use MTF measurements to determine the minimal cell size resolvable will be presented in Section 3.5.5, after detailing the arrangement of the complete IFC system. To find the limiting resolutions using MTF10 and MTF50 values, MTF values will be converted from units of cycles/pixel to cycles/ μm by multiplying them by the reciprocal of the pixel size of the camera, and then by the magnification of the objective.

3.5 IFC System Design

3.5.1 Introduction

This section describes the development process followed to incorporate the acoustic flow cell, pulsed illumination and camera into a complete system. The following subsection (3.5.2) details the initial approach taken to this integration, which involved laser-cutting an enclosure from acrylic. This design proved insufficient for several reasons which are explored. The second approach (Section 3.5.3) revolved around using a commercial cage/rail system and a custom 3D-printed flow cell holder, which overcomes the disadvantages of the first design while maintaining acceptably low cost and complexity.

After finalising the hardware components and layout of the IFC, the imaging performance of the assembled complete system is measured using the procedure described in Section 3.4. The smallest resolvable phytoplankton cell is determined using the measured MTF values of the instrument.

3.5.2 System Camera Selection

The camera selected for the final system design was a 2.3-megapixel colour USB3 camera (*Grasshopper GS3-U3-23S6C-C, Teledyne FLIR, USA*) using an IMX174 (*Sony, Japan*) CMOS Sensor. The camera was chosen due to it satisfying the below selection criteria:

- Pixel Density: The camera has 2.3 Megapixels (1920x1200 pixels), which is sufficient to capture a wide field of view.
- Pixel pitch: A pixel pitch (size) of 5.86 μm x 5.86 μm is large compared to many alternative cameras investigated. This large pixel pitch allows the pixels to capture more light and therefore makes the camera more appropriate for the high light sensitivities required for microscopy applications.
- Framerate: The camera is capable of operating at up to 163 FPS at a reduced resolution, and over 50 FPS at full resolution. By achieving imaging speeds greater than 30 FPS, the camera allows *real-time* imaging of the flow cell.

Chapter 3

- **Global Shutter:** Of the two types of shutter used in digital cameras, i.e. rolling and global shutters, only cameras with global shutters are appropriate for the IFC system. This is because all the pixels of the camera must be exposed during the short stroboscopic illumination pulses generated by our LED system. If a rolling shutter were to be used instead of a global shutter, the pulse could be faster than the time it takes the camera to exposure the entire sensor.
- **Hardware Trigger I/O:** The camera chosen has I/O ports which can directly trigger or be triggered by illumination hardware. This is critical to allow the synchronisation of the illumination with the exposure.
- **C-mount:** The Thorlabs components selected for the IFC platform have a C-mount adapter for the camera, allowing the camera to easily integrate with the rest of the system.
- **Low Cost:** In order to address the stated requirement of producing the complete IFC at as low a cost as possible in order to allow its adoption by resource-limited researchers, it was crucial that the camera cost less than the specialised cameras designed for microscopy, which can often cost several or tens of thousands of pounds. The FLIR camera cost £842, which satisfied this requirement.

3.5.3 First system design approach

With the aims of minimising complexity and cost while maintaining the fastest possible implementation time for adoption by non-experts, initial mechanical design of the IFC heavily utilised laser-cut Perspex, combined with components (rods, bearings, lead screw, motor) used widely in the consumer-grade 3D printer industry. The complete Version 1 design of the IFC mechanical system can be seen in Figure 44. The motivation behind this initial mechanical setup was to use the automated vertical movement of a camera above a fixed objective lens to change the depth of the focal plane as required to focus on the cells in flow. The objective lens and flow cell were mounted on lockable frames, with cut-outs in the outer enclosure to allow access to the locking screws without disassembling the enclosure.

Chapter 3

A stepper motor (*NEMA 17*) controlled by an *Arduino Nano* microcontroller was used to drive the camera carriage, and limit switches used to prevent overextending the carriage or collision with the enclosure. Optical planarity throughout the system was achieved by means of using a set of 4x 6mm rods vertically secured to the top and bottom of the enclosure.




Figure 44 – Diagram of the first iteration IFC mechanical arrangement. Using a stepper motor, 6 mm rods and linear bearings used in 3D printers, a carriage containing a camera can be raised or lowered as required to achieve focus of the cells passing through a flow cell below a standard 10X objective lens. By using a matt black Perspex enclosure, light from outside the device is limited and no tube between the camera and objective is required. The illumination frame can be lowered and an arbitrary number of optical components inserted above to create desired lighting conditions (e.g. Kohler).

Chapter 3

While this system is quick and straightforward to assemble due to the speed of laser cutting and ability to incorporate design features to allow easy construction, the range of focal depths proved insufficient for adjusting focus from the top to the bottom of the flow cell without having to also move the objective lens, which is challenging and can reduce planarity through the optical path.

An additional problem which arose during the initial testing of this version of the mechanical design was due to the necessity of opening the device front (cut away to show inner components in Figure 44) in order to adjust the vertical height of the frames which hold the objective lens and flow cell. An 'open' design, without any tube between the camera and objective, while simple and cheap to implement, was found to a nuisance to focus as the device needed to be in a darkened room to operate without the light-blocking exterior enclosure in place. These issues were considered unacceptable and motivated a complete redesign, which is described in the following Section.

3.5.4 Improved system hardware

For the second iteration of the mechanical IFC setup it was decided that the use of additive manufacturing (3D printing), as opposed to laser-cutting, offered a reasonable trade-off between construction speed and design freedom. Also, by using off the shelf commercial optical components, particularly *Thorlabs SM1* lens tubing between the objective and camera, the prior problems with focussing range could be addressed without increasing device cost significantly. The fixed-length tubes would be required even in the absence of a coarse focus adjustment mechanism as the objective lens is of the *Deutsche Industrie Norm*, or *DIN* standard. *DIN* objectives, an internationally adopted standard, require a focal tube length of 160mm. This is as opposed to 'infinity-corrected' objectives, which does not focus the light on a point but emits parallel rays which must separately be focussed onto the camera's imaging plane. The design for the 3D printable flow cell holder is presented below in Figure 45.

The decision was also made to mount the flow cell vertically rather than horizontally like in the initial design. This modification means that gravity does not act to pull cells out of the acoustic pressure node, increasing the chance that they will still be within the appropriate imaging plane after leaving the transducer region and entering the imaging region.




Figure 45 - 3D-printable flow cell holder for the prototype IFC. The acoustofluidic flow cell slots (from the right) into the cut-outs and is pushed until it reaches the end (left). The circular hole in the centre allows an unobstructed optical path for imaging and illumination, and standoffs in each corner, through which the cage rods are inserted, prevent flexion. Inlay: close-up of the flow cell inserted into the holder (blue).

In order to maximise the ease with which the device can be adjusted, particularly with respect to changing the distances between optical components (e.g. flow-cell and illumination source) as required to achieve focus quality, a system of metal rods and *cage plates* used. As Figure 46 demonstrates, metal lens tubes (*SM1L range, Thorlabs Inc., USA*) couple the camera via an adjustable-length tube (*SM1V15, Thorlabs Inc., USA*) to the objective, providing planarity, coarse focus adjustment and external light exclusion. Fine focus adjustment is by means of a high-precision zoom housing (*SM1ZM, Thorlabs Inc., USA*). Condenser lens (*ACL2520U-DG6, Thorlabs Inc., USA*), collimating lenses (*47-637, Edmund Optics, USA*), and adjustable irises (*CP20S, Thorlabs Inc., USA*) were used to provide Köhler illumination (Köhler, 1893), following the design of Madrid-Wolff and Forero (2019). A Köhler setup provides uniform, collimated illumination intensity across the sample independent of light source shape and size (Voelkel et al., 2010) and is the same configuration typically used in research brightfield microscopes.




Figure 46 – Current IFC prototype arrangement. From the left: illumination from a pulsed LED is focussed via means of 2 iris and 2 lenses. The acoustofluidic flow cell slots into a 3D-printed flow cell holder (blue), which is sandwiched between metal cage plates for rigidity. The objective lens is mounted on course and fine focus apparatus, which themselves are connected to a camera via a SM1 tube.

The total cost of the optical, imaging and computational components of the system after integrating the camera was £2240.47 (A bill of materials for the IFC is presented in Appendix D). This is less than 1/20th of the cost of the FlowCam, the cheapest commercial alternative (see Table 4 in Section 2.5.2 for a cost comparison with other instruments).

This cost does not include the flow cell, transducer or electronics needed to generate the acoustic standing wave as these components are highly customisable and can be as expensive or cheap as required depending on the trade-off between complexity of assembly, ease of use, and cost. The IFC was tested with a flow cell driven by a commercial, lab-grade benchtop function generator and amplifier, as well as with a custom prototype-board-based Direct Digital Synthesiser and amplifier obtained on eBay for <£20. Both approaches demonstrated success in acoustic focussing of beads and phytoplankton cells, but the latter required in-depth knowledge of coding and electronics. The parts required to assemble the flow cell are under £20, but require access to a laser printer and glass drill, so the cost to a potential user is difficult to state conclusively.

Overall, this system represents a major step toward improving the accessibility of high-quality imaging analyses by achieving the design goals of having a higher throughput than existing commercial devices while being fully modifiable and at a significantly lower cost.

3.5.5 Imaging Performance of IFC System

After integrating the camera with the rest of the IFC hardware, the MTF of the complete system was determined using the procedure detailed in section 3.4. A slanted edge analysis experiment was carried out using the same imaging target as used when developing the procedure (*R1L3S5P*, *Thorlabs Inc., USA*) but now mounted in the flow-cell holder presented above in Section 3.5.3, and imaged using the IFC camera described in Section 3.5.4. After appropriately adjusting the Köhler illumination, the target was imaged using a DIN 10X objective (*MSB50100*, *Nikon, Japan*) and the edge slant was 10°. Several images were captured while adjusting the focus of the IFC objective lens such that the maximum measured MTF would not be dependent on the subjective quality of focus – a key point raised during the experiment of section 3.4.6. Each exposure was illuminated by a 5 μ s LED pulse, which were verified to ensure that pixels in the captured images were not fully saturated.

The MTF curve was calculated for each image captured and the curve with the highest MTF10 and MTF50 values selected as a representation of the best possible performance of the IFC. This MTF curve is presented below in Figure 47. The measured values for MTF10 and MTF50 were 0.167 and 0.264 cycles/pixel, respectively. These MTF characteristics are significantly worse than those measured for the inverted microscope system as detailed in section 3.4, which saw MTF10 values of over 0.2 cycles/pixel and MTF50 values of over 0.35 cycles/pixel, but this is to be expected given the difference in sophistication of the objective lenses used, the cameras used by the two systems, and the issue of continuous focus adjustment discussed in detail in Section 3.4.6.




Figure 47 – MTF curve for the complete IFC system. The spatial frequency at which the MTF drops to 50% (MTF50) is 0.167 cycles per pixel, and the spatial frequency at which the MTF drops to 10% (MTF10) is 0.264.

As discussed previously in Section 3.4.6, of greater importance for the IFC system than a direct comparison of MTF values with commercial equipment, is the realistic smallest particle which can be imaged, as that will determine the lower bound on the size spectra of phytoplankton which can be analysed using the device. Therefore, as previously stated, the measured MTF values will be used to estimate a realistic minimum feature size.

As stated in Section 3.4.2, MTF10 is the spatial frequency at which the contrast falls to 10%, while MTF50 is the spatial frequency at which the contrast falls to 50%. MTF50 can be used as a more conservative measure of resolution than MTF10, so MTF50 will be used in the estimation. Since MTF50 = 0.167 cycles/pixel, the spatial frequency in cycles/ μm can be determined as follows:

$$\begin{aligned}
 \text{Spatial frequency (cycles}/\mu\text{m}) &= \text{MTF50 (cycles/pixel)} * (1 / \text{pixel size in } \mu\text{m}) \\
 &= 0.167 \text{ cycles/pixel} * (1 / 5.86 \mu\text{m}) \\
 &\approx 0.0285 \text{ cycles}/\mu\text{m} \text{ (3sf)}
 \end{aligned}$$

Chapter 3

Next, it is possible to find the corresponding resolution (minimum resolvable distance) for the camera system:

$$\begin{aligned}\text{Camera resolution} &= 1 / \text{spatial frequency} \\ &\approx 1 / 0.0285 \text{ cycles}/\mu\text{m} \\ &\approx 35.1 \mu\text{m} \text{ (3sf)}\end{aligned}$$

Finally, camera resolution is divided by the magnification of the objective in order to find the actual resolvable size in the sample plane:

$$\begin{aligned}\text{Resolvable size} &= \text{camera resolution} / \text{magnification} \\ &= 35.1 \mu\text{m} / 10 \\ &\approx 3.51 \mu\text{m} \text{ (3sf)}\end{aligned}$$

A minimum resolvable cell size of 3.51 μm seems to align well when compared to a qualitative visual analysis of a cell imaged by the IFC. Figure 48 presents a cropped image of a single phytoplankton cell (*Rhodomonas salina*) taken by the IFC; while the cell is slightly more than 5 μm across its smaller axis, some internal cell structures of less than 5 μm are made visible which would be expected if the resolution limit was below 5 μm .



Figure 48 – A single cell of *Rhodomonas salina* (fixed with Lugol's Iodine solution 1%) imaged by the IFC system, demonstrating the ability of the system to partially resolve internal cell structures significantly smaller than 5 μm .

Repeating the above calculations to estimate the smallest resolvable size using the value of MTF10 rather than MTF50 leads to a value of 2.21 μm (3sf), which is a plausible lower bound to the minimal resolvable cell size within IFC images. Based on these values for MTF10 and MTF50, it is proposed that an approximate minimum cell size for the device is 3 μm diameter, which

Chapter 3

compares very favourably with the Imaging Flow Cytobot and FlowCam, both of which have posted minimum sizes of 10 μm , as described in Table 2.

3.6 Conclusion

In this chapter, work to design and systematically explore the modelled and real-world performance of a high-throughput, low-cost, disposable, acoustically-focussed flow cell was presented. Different flow cell strategies were investigated, resulting in the selection of a novel *matched-layer half-wave planar resonator* for the IFC. This design builds upon the work of previous studies, implementing a new matching layer to allow superior imaging performance within our device.

Next, finite element modelling of the final flow cell to assess its acoustic focussing performance was presented, using model parameters designed to cover the likely range of phytoplankton properties the final IFC will encounter. It was then attempted to develop a novel imaging method for assessing acoustic focus using the distanced travelled by particles during a fixed time, but this proved unsatisfactory, the reasons for which were explored. It was found that the simple test of assessing the number of objects appearing in focus due to the visual sharpness of their edges is sufficient to reliably detect whether the acoustic focussing is operating successfully. This approach will be used throughout the further experimentation with the flow cell detailed in subsequent chapters.

Overall, the analyses conducted in this chapter provide confidence that the flow cell will allow the successful acoustic focussing of a wide range of phytoplankton into a two-dimensional sheet for high-throughput imaging by the optical system of the IFC.

Next, to ensure that those optical components are of sufficient performance to image small plankton cells and to accurately characterise the lower limit of cell sizes that could reasonably be imaged, it was necessary to develop a robust protocol to measure the Modulation Transfer Function of a microscope. The MTF gives a complete overview of the optical performance of a complete imaging device, including the camera, illumination and lenses in the system. Therefore, testing of the sensitivity of MTF measurements to small variations in rotation of the target used for MTF testing, exposure of the camera, and focus depth was carried out. After carrying out these experiments, the range of conditions under which MTF could reasonably be measured for the completed IFC could be determined.

Chapter 3

Finally, the work to assemble the full IFC device was reported. The hardware design of the IFC was significantly revised after initial testing showed the many inadequacies of the first design. The finalised design successfully achieves the design goals of quality, ease of assembly and modification, using only easily accessible components, and having a low total cost. The total cost of consumables is also low; no chemical flushing agents (or similar) are required for the operation of the IFC. The flow cells have been described as 'disposable' as they are cheap to fabricate, as other than the transducer they consist only of standard microscope slides and cover slips, epoxy and acrylic. The cost to actually fabricate these flow cells may in reality be significantly increased by the relative complexity and time-consuming assembly process, but this could in principle be alleviated using batch fabrication techniques. The current design also requires the use of a laser cutter, which while not a specialist piece of equipment, may be difficult to access by less resource-rich researchers.

After detailing the design and assembly of the complete system, the optical performance was measured using the previously detailed MTF protocol. It was found that the device has a lower limit of phytoplankton size of approximately 3 μm , which is promising as it suggests the IFC can be used to analyse a large section of the plankton size-spectra.

Given the design of our IFC was optimised to balance performance characteristics while minimising cost, our design uses a moderately-priced objective lens and camera. Since the IFC is designed to be open-source and easily modifiable, it would be trivial to install a higher-quality camera and/or objective lens. Therefore, the optical performance of the device presented in this chapter is not a hard limit, and instead reflects a design approach designed to make the instrument easily accessible to resource-constrained researchers.

In conclusion, this work has detailed the design and construction of a novel acoustically-focused IFC which will enable greater access to imaging flow cytometry for resource-constrained oceanographic laboratories. A robust characterisation of the device in terms of acoustic focussing performance and image quality was presented. The device described in this chapter has the potential to not only expand the use of IFC for phytoplankton research across diverse research scenarios but to offer a new instrument with increased throughput and image quality compared to far more expensive existing instrumentation such as the *Imaging Flow CytoBot* and *FlowCam*. It

Chapter 3

is reasonable to be excited about the possibilities of our device to make a significant positive impact in real-world oceanographic, aquaculture, industrial and algal biotechnology applications.

The following chapters will detail the development of software to automatically analyse the images produced by our IFC (Chapter 4), and the use of the IFC for real ocean microbiological research (Chapter 5).

Chapter 4 Image Analysis for the Quantification of Phytoplankton in IFC Images

4.1 Introduction

Having used the IFC presented to capture an image of a water sample, the image must be analysed to process the raw pixel data into useful information (e.g., presence/absence of organisms; cell count; segmentation of cell images from the image frame; cell species and/or life-stage classification). As previously discussed, this analysis has traditionally been performed by a human familiar with the plankton species likely to be present within a sample. However, with the development of computer graphics and digital photography, algorithms have been created to automate these tasks. Automated and semi-automated digital processing methods have now been employed for a diverse range of phytoplankton research.

The FlowCam IFC, for instance, can automatically segment and measure plankton cells based on proprietary image processing software, as detailed by Poulton (2016). Modern computer vision techniques, such as machine learning, are increasingly being utilised with IFC imagery. Campbell et al. (2010) detail the use of a machine learning approach based on Support Vector Machines to detect harmful dinophytes in an estuary and Kraft et al. (2022) used a Convolutional Neural Network approach to automatically classify the species of phytoplankton sampled by an *Imaging Flow Cytobot* in the Baltic Sea.

By processing imaging data computationally, rather than manually, the volume of data that can be analysed is drastically increased, and the reliance on human experts is reduced (though not eliminated). Efficient automated analysis also presents the capability of processing samples in real-time, at the point of sampling. However, the accuracy of automated systems, particularly for those with sensitive applications such as HAB early-warning platforms, needs to be carefully considered.

This chapter presents a cutting-edge solution for fully automated phytoplankton analysis. It begins with a discussion of the shortcomings of traditional image processing approaches to automated image analysis, which rely on human experts to hand-craft suitable algorithms. This is

Chapter 4

followed by a description of work to develop a novel Convolutional Neural Network (CNN) for detecting and counting plankton cells within images, which allows for the processing of large volumes of data with greater accuracy than traditional methods. Finally, an application of the state-of-the-art object detection model 'YOLOX' to IFC images is presented, which represents the first time this high-performance model has been employed for brightfield imaging flow cytometry. By leveraging these innovative methods, the work presented here aims to significantly advance the field of phytoplankton analysis and enhance our ability to monitor harmful algal blooms *in situ* and in real-time.

4.2 Traditional Image Processing

4.2.1 Introduction

This section will discuss traditional digital image processing techniques, which have been widely used for object detection in various fields, including medical imaging, remote sensing, and robotics. However, these techniques have limitations when it comes to analysing images of phytoplankton generated by the IFC, which will be explored.

The first set of traditional techniques mentioned is *edge detection*-based. Edge detection is a fundamental technique for image processing that involves finding boundaries between objects in an image. In edge detection-based object detection, the algorithm identifies edges and then segments the image based on these edges. However, this technique often fails to detect objects with low contrast or those with borders which are not well defined, both of which are potentially problematic with respect to images of organic cells.

The second set of traditional techniques mentioned is *contrast*-based. Contrast-based methods use differences in intensity or colour between the object and the background to identify the object. This approach can be useful when the object has a high contrast with its surroundings, but it may struggle when the background is complex.

To address these limitations, traditional image processing techniques often require the use of several refinement techniques, such as colour/brightness thresholding, background removal, blurring, and morphological thresholding. Colour and brightness thresholding can be used to detect objects based on their colour or intensity, while background removal can help to eliminate unwanted background artefacts and noise. Blurring can be useful for removing small details that reduce the performance of edge detection algorithms. Morphological thresholding can limit detections to objects within a certain size range as determined by pixelwise area.

Despite their effectiveness with images having clear visual distinction between objects of interest and the background, traditional image processing techniques still have limitations when it comes to analysing the complex phytoplankton images generated by the IFC. The main challenge to

Chapter 4

traditional image processing techniques, even with extensive refinement, is the highly variable and often high-contrast background objects that can make it difficult for the algorithm to distinguish the phytoplankton from cells and detritus adhered to the flow cell. Additionally, the traditional techniques require manual fine-tuning by the user to deal with specific image conditions, which is time-consuming and limits their effectiveness.

As an example of a pipeline for object detection via traditional methods, consider the example of locating the (x,y) centre coordinates of the various ellipses in a simple image such as that presented in Figure 49:




Figure 49 - Example image created for demonstrating the traditional methods of object detection.

The image contains ellipses of two different colours, sizes and eccentricities overlaid on a plain black background. The image is a 400x300 .jpg file.

To more closely align with the detection of plankton in IFC images, the problem here is formulated such that regardless of the colour, size, eccentricity, and number of ellipses within the image, the system should be able to determine a centre coordinate for each ellipse, thus allowing the shapes to be counted, tracked over time, etc. This aligns closely with the requirements of the image processing from the IFC, but with drastically reduced complexity.

4.2.2 Example Algorithm

A simple typical pipeline for detecting the ellipses' centre coordinates using traditional image analysis techniques involves the following steps.

- 1) Conversion to grayscale (Figure 50). The image is converted from 3-channels (RGB) to 1. This pre-processing step is required as later functions operate only on single-channel images.




Figure 50 – The ellipses image after being converted to grayscale.

- 2) Blurring the image (Figure 51). A blur is applied to the image to eliminate any noise. The type and strength of the blur must be specified. For this example, a simple Gaussian smoothing is applied to the image, which involves multiplying pixel values by a sliding window filter, a weighted matrix with the highest value at the centre and gradually decreasing values as the distance from the centre increases, according to the normal curve. In this manner, high-frequency details are removed while low-frequency features are preserved. This improves the performance of subsequent processing.




Figure 51 – The grayscale image of ellipses after applying a 5x5 pixel gaussian blur operator.

- 3) Edge detection (Figure 52): The next step is to detect the edges in the image. This is done by applying a filter or gradient operator to the image that highlights the regions with the most rapid changes in luminosity or colour. Multiple algorithms have been introduced for edge detection, including *Sobel* (Duda and Hart, 1974) and *Canny* (Canny, 1986) edge detectors.




Figure 52 – The images of ellipses after the Canny edge detection algorithm has been applied to the blurred, grayscale image.

- 4) Threshold Masking (Figure 53). The results of the edge detection algorithm are often overly detailed and noisy, so a thresholding operation is applied to the image. To perform this step, pixel values below a specified lower threshold are set to zero and pixel values above a specified upper threshold are set to 1, creating a binary image output.




Figure 53 – The image after grayscale, blurring, edge detection and thresholding. All that remains now is pixel values of 0 in regions which are not included in the ellipses, and values of 1 inside the ellipses.

- 5) Contour detection (Figure 54). The contours of the objects are extracted by tracing the boundary between the foreground and background regions of the image. Several algorithms are available for contour extraction, including the Douglas-Peucker algorithm (Douglas and Peucker, 1973) and the active contour model (Kass et al., 1988). In all cases, the output is a series of connected point coordinates that define the boundary of each detected object. This step is required if for the extraction of the outline of each object. If one were only interested in the number of objects, it would be possible to just count the connected regions of pixel values of 1. If one only wanted the maximum extent of the objects – to draw a bounding box around the objects, for instance, one could find the maximum and minimum coordinates in x- and y-directions for each connected region.




Figure 54 – The ellipses after contour detection using the active contour algorithm. The outside edges of the objects of interest are now known and represented here in green.

- 6) Localisation (Figure 55). Once the contours of the shapes have been determined, the objects can be localised as required. For instance, a rectangular bounding box that encloses each entire contour can be easily determined from the output of the previous step. Once a bounding box has been determined, the centrepoint coordinates of the shapes can be easily calculated. Alternative approaches to centrepoint extraction include moment-based methods, which involve computing the first-order moments of the contour, e.g. as discussed by Yuan and Hui (2008).




Figure 55 – The centre coordinates of each detected ellipse have been determined using the minimum and maximum extents of the previously calculated contours of the shapes, and are now plotted on the image.

4.2.3 Limitations

The above algorithm works well for the extremely simplistic example image shown in Figure 49, but is easily shown to be inadequate by considering various conditions possible within the problem as specified but not present within the image previously investigated.

For instance, if an ellipse appears which has a brightness of colour close to the background (black), that ellipse may not be above the threshold value previously selected for binary masking,

so may not pass on to the contour detection step and thus would not have any associated coordinates determined. In the context of the real IFC images, this could occur as a result of cell transparency or occlusion by background debris. Figure 56 illustrates this condition; only the lightest of the 3 dark ellipses is successfully detected. In order to allow the algorithm to obtain the correct result it would be necessary to go back and manually adjust the binarization threshold.




Figure 56 – A new image of ellipses is presented in Figure 55a. The image now contains 3 ellipses of colours having a low contrast with the background. Figure 55b demonstrates the problem with the previously discussed image analysis approach; the contour detection has only operated successfully on the ellipse with the highest contrast from the background as the thresholding step used after edge detection had a threshold value set too high to include the darker objects. As a result, only that ellipse has its centre coordinates calculated, as demonstrated in Figure 55c.

If ellipses are introduced to the image with borders either very close to another ellipse, or even partly occluded, the contour detection algorithm fails, as shown in Figure 56c. In real IFC images, cells have a high likelihood of presenting close to or in contact with other cells.

To successfully account for this scenario, the software would need to be made far more complex by the additional processing step of a watershed algorithm (Roerdink and Meijster, 2003), which itself requires manually tuned parameters. Morphological transforms, transformations of the threshold mask to 'erode' the size of ellipses and allow their separate detection, may be required. These too need carefully selected parameters to ensure adequate performance.



Figure 57 – in this new image example, 2 of the 3 ellipses are partially overlapping. Though to a human, the 3 ellipses are clearly distinct objects, the algorithm which has been developed so far fails to distinguish the two objects and instead treats them as one. Figure 56b shows that the contours are calculated for the combined ellipses and therefore the centroid of the combined objects is plotted on Figure 56c.

If the background is not pure black but instead a gradient, while the ellipses are still easily distinguished by a human viewer the performance of the processing algorithm explored thus far will be severely hindered. In the real IFC, the lighting is often not of a perfectly consistent intensity and has a gradient with a maxima to the centre of the image. Figure 58 shows an example of a gradient background, with 3 ellipses overlaid on a blue gradient background. The algorithm is incapable of distinguishing between the background and foreground and so fails as shown in Figure 58b. and c., only successfully detecting the foreground object with the highest contrast from the background.




Figure 58 – In this final example, Figure 57a. shows 3 ellipses which are overlaid this time on a gradient background, rather than the plain black of the previous examples. Though the 3 ellipses are still easily recognised by a human, the algorithm is unable to cope with the low contrast that the gradient produces and so falsely calculates coordinates for two of the ellipses and part of the gradient as a combined object, as shown in Figure 57b. Therefore, as Figure 57c, demonstrates, only one of the 3 ellipses is correctly identified.

As a solution for this problem, the system could be extended to include an initial background subtraction step, whereby the mean of several images is subtracted from the image to be analysed before proceeding. This would allow the unchanging (background) elements to be removed effectively but would necessarily cause the algorithm to be unable to process some images which would be used for calculating the background.

4.2.4 Limitations in IFC Imagery

Though each of these complicating factors can be manually controlled for by increasing the algorithmic complexity and/or fine-tuning the parameters selected for thresholding, etc., the system will still be highly susceptible to error if the image content changes beyond the starting conditions that the algorithm was developed for. In the case of the IFC presented in this thesis, these unpredictable changes are highly likely and can occur due to several factors, which are briefly explored below.

The background can change significantly over the lifetime of the flow cell as various debris adheres to the inside of the glass and causes uneven shadowing, similar to the background gradient problem explored. Cells and detritus can be pushed onto the flow-cell walls by the effect of negative acoustic contrast factor discussed in the previous chapter. Figure 59 demonstrates

Chapter 4

two background images of the same flow cell, separated by one week's operation with regular flushing of phytoplankton-containing seawater through the device. Many additional cell debris have accumulated on the edge of the glass flow cell, producing background 'shadows' which obscure any cells moving in front across the foreground. Additionally, some debris which was present in the earlier image has moved or detached, demonstrating the variable nature of adhered detritus and rendering background subtraction techniques impossible.




Figure 59 – Two images of the IFC flow cell while filled with sterile L1 medium but no phytoplankton cells. Figure 58a. shows the flow cell after processing several phytoplankton samples and has obvious shadows which are cast onto the field of view by detritus adhered to the sides of the flow cell. Figure 58b. shows the same flow cell after a further week of regular sampling operation. More debris has adhered to the flow cell, so the background complexity has been significantly increased. The arrow points to an area of particularly heavy new adhered debris. This variable background poses a challenge to image processing algorithms.

Cells are not guaranteed to present singly, not overlapped and with low proximity to other cells; it is entirely possible for cells to touch or even partially occlude each other, though the probability of these occurring is dependent on the cell density, acoustic contrast factor, cell size and other factors. Figure 60 illustrates this challenge, showing two plankton cells touching each other.



Figure 60 – IFC image of 3 phytoplankton cells, 2 of which are in contact with each other (arrow).
As previously discussed, overlapping and touching objects pose a challenge to image processing techniques.

Organic material is often highly transparent, and it is possible that some plankton species will have a low contrast with their background, especially on areas where the background is darkened by adhered debris. Though these cells can be easily recognisable to a human observer, traditional processing methods are typically not robust enough to deal with such cases. Figure 61 illustrates one example of this problem, with 2 cells visible, one of which is passing in front of a background with which it has very little contrast.



Figure 61 – close-up view of a phytoplankton cell (arrow) traversing in front of adhered detritus, the shadow of which is cast from the edge of the flow cell into the field of view.

Chapter 4

Even if all of these shortcomings are fully addressed with an improved algorithm, the traditional processing techniques discussed are tailored to a specific set of initial conditions, and need to be carefully readjusted for each new use case (e.g., different phytoplankton species, each of which may have different shapes, colours, and opacities).

More powerful image processing methods based on machine learning have been shown to overcome these limitations by automatically learning the most relevant features and patterns from the input data, without the need for explicit feature engineering or manual adjustments. These approaches, particularly deep learning techniques such as convolutional neural networks (CNNs), have been successful in various image recognition and classification tasks, including the identification and analysis of phytoplankton species (Zhang et al., 2023), and will be fully explored in the following sections.

4.3 Machine Learning-based Image Processing

4.3.1 Introduction

Supervised Machine Learning refers to the use of computational optimization algorithms to determine appropriate values of parameters to complete a task, rather than those parameters being set by a human. These tasks are typically categorized as either *regression-based* (predicting numerical values), or *classification-based* (assigning a class). In the field of image processing, there is a wide variety of methods and applications of machine learning. For instance, the algorithm presented in section 4.2.2 used manually determined parameters such as the kernel filter size of a Gaussian blur and a threshold value for binarizing an edge image, but one could instead use regression-based machine learning to find optimal values for these values such as to demonstrate the highest possible accuracy of the complete algorithm.

Supervised machine learning algorithms are usually trained on large datasets of labelled examples and determine optimal values for parameters such as to model the underlying patterns in the data. Once these parameters have been optimised, the algorithm can be used to make predictions about new data it has not seen. Algorithms used to optimise parameters for image analysis include Linear and Polynomial Regression, Decision Tree and Random Forest models, Neural Networks, Support Vector Regression and Bayesian Linear Regression (Gareth James et al., 2013).

Machine learning allows image analysis algorithms to take on layers of sophistication that would be impractical or impossible using traditional approaches. The algorithm presented in 4.2.2 is extremely simple and, as demonstrated, is not very effective for the problem of cell detection in the images produced by the IFC presented in this thesis. A better but more complex approach would be to manually determine a set of features common to each cell; these may include colour and luminosity features, morphological features such as size and roundness, etc. Hand-crafting these features would be very time consuming and, as discussed in section 4.2.4, could increase the specificity of the algorithm so that small changes in image content, such as the background changing due to the movement of cells adhered to the flow cell, would prevent accurate detection of cells. Nevertheless, hand-crafted features combined with machine learning optimization strategies have demonstrated adequate performance in more simplistic image processing tasks (Lin et al., 2020).

A family of models which is now commonly used for automated image analysis is the Convolutional Neural Network (CNN), which will be explored in detail in subsequent subsections. CNNs automatically learn features from images by optimizing sets of filters, which allows the network to learn relevant features at multiple scales and orientations. CNNs have been used in a wide range of applications, including classification, object detection, image segmentation and image generation.

Before going on to lay out the application of Machine Learning to the IFC images generated by the device presented in this thesis, it is useful to summarize the different categories of machine learning approaches in the context of image processing. Broadly, these categories can be defined as the following:

1) **Image Classification**

Image classification problems involve assigning a class to an image based on a set of learned mappings between images and class labels. There may be one or more classes and each image may be assigned either a single-class (binary) probability, a class based on the highest inter-class probability, or labelled with multiple classes based on the per-class probabilities exceeding a threshold value. An example with hypothetical outputs of each is shown in Figure 62, where for an input photograph of a dog sat in a field of bluebell plants, a binary classifier gives a [0-1] probability of the image being of a dog. A multiclass classifier gives [0-1] probabilities for the image content being of one of several classes, and only the class (dog) with the highest probability is selected for the output. In the multilabel classifier, each of the probabilities is independently checked and those classes assigned probabilities above a threshold value (dog **and** plant) are labelled to the image.

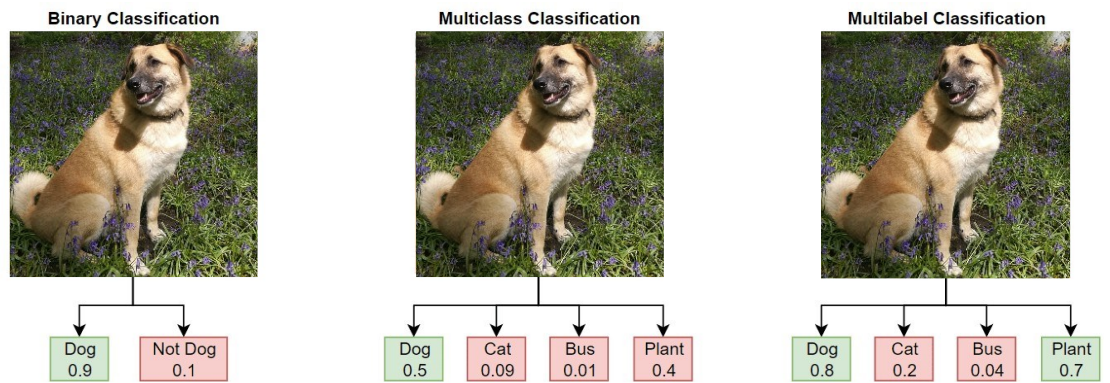


Figure 62 - An image of a dog in a field of bluebells is used to demonstrate three kind of image classification. In Binary Classification, a model which has been trained to recognise images of dogs outputs a 0-1 probability of the image containing a dog. In Multiclass Classification, the model is trained on multiple classes (Dog, Cat, Bus, Plant) and for each class outputs a probability of the image containing an example of that class. The class with the highest probability score, in this case Dog, is selected as the output. In Multilabel Classification, the model can assign more than one label to the image. In this example, the model outputs both Dog and Plant as both classes have probability scores which pass some threshold value. Adapted from (MathWorks, 2023).

In the context of imagery of phytoplankton which are produced by the acoustophoretic IFC presented in chapter 3, there are multiple potential use of each of the aforementioned image classification modalities. Firstly, for a given full-frame image of the presented IFC flow cell, a binary classifier could output a probability of the image containing cells or being empty, as demonstrated by the hypothetical output predictions presented in Figure 63. This type of classifier could be used to automatically discard 'empty' frames in IFC experiments; a setup like this could reduce storage use by only saving images which have cells present.

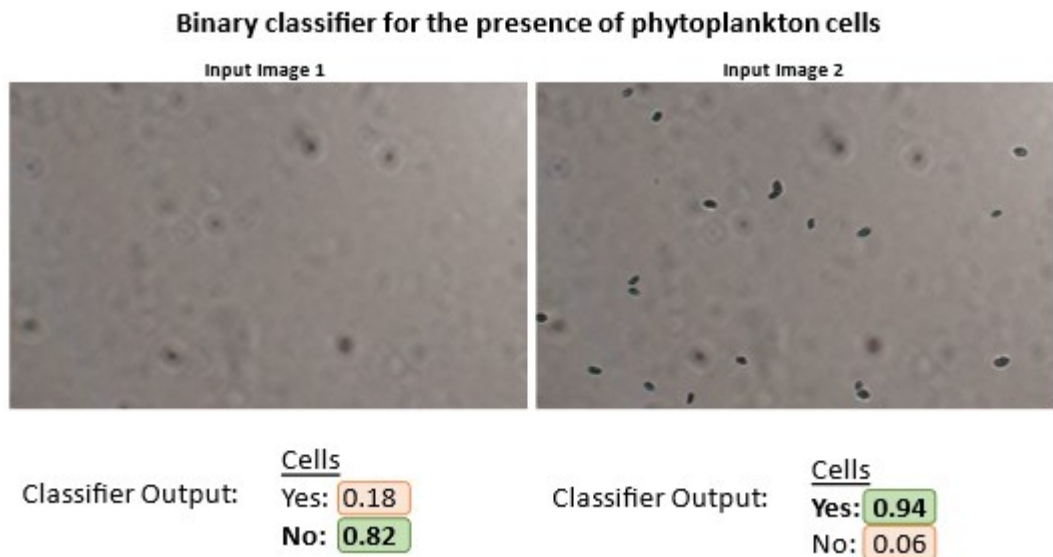


Figure 63 - Two images from the acoustophoretic IFC device are presented. For each of the two input images, a binary classification model which has been trained to detect cells could output a probability score of the image containing cells. Based on this score, frames which do not contain cells, like the first input image, could be discarded without being saved, saving memory and processing time.

If instead of whole-frame images, the classification algorithm was trained and used on cropped images of single cells, as in Figure 64, a hypothetical binary classifier could output the probability of the cell being healthy. A classifier of this type could be used in industrial phytoplankton cultivation to automatically monitor the health of a culture.

A multiclass classifier could be used to automatically determine the type of phytoplankton in the image from a list of possible genus/species. A hypothetical multiclass classifier output is also presented in Figure 64. Such a system would be useful for automatic monitoring of community composition for applications including HAB monitoring and rapid phytoplankton assemblage analysis. A multilabel classifier could output predictions of the species of multiple phytoplankton cells within a single image, which would have utility in analysing more complex samples not of a monoculture and with multiple cells per image.

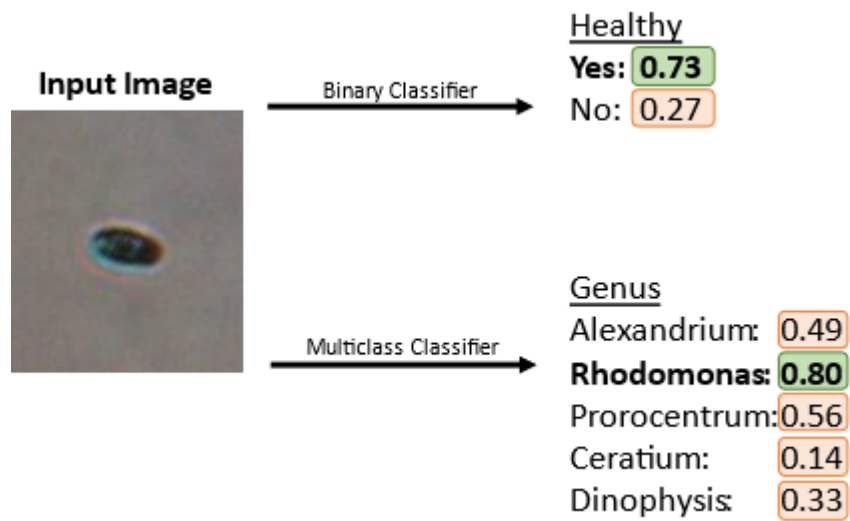


Figure 64 – An image of a single phytoplankton cell is used as an input for hypothetical binary and multiclass classifiers. The binary classifier can predict whether a cell is healthy or not, whereas a multiclass classifier could be trained to distinguish between several genus of phytoplankton and output a probability score of the input image being an example of each, selecting the highest probability.

2) Object Detection

As demonstrated above, a binary classification algorithm can take an input image and output, for instance, a probability of that image containing phytoplankton cells. If instead of a probability for the whole image, the algorithm was trained to output a *per-pixel* likelihood of the pixel being a part of a phytoplankton cell, the cell could be *localised* to a given region. An example of this kind of single-stage object localisation algorithm is presented in Figure 65, where for an input IFC image of a phytoplankton cell against an empty background, a model outputs a probability of each pixel being a cell in an output known as a probability density map (or pixel mask). From this map a variety of simple algorithms could be used to determine the exact extent of the object of interest, e.g. drawing a rectangular 'bounding box' around the cell or counting the number of pixels contained within it. This could allow software to measure the cell or to crop the cell from the input image, allowing the storage of only the region of interest. After cropping, the cell image can be passed to a secondary algorithm (e.g. species classifier); etc. This method is an example of the process known as Object Detection.




Figure 65 – An object detection algorithm based on a binary classification of each individual pixel within an input image. The algorithm generates a probability score (0-1) of the pixel containing a cell, after which a number of algorithms can be used to, for instance, count, localise or measure cells within an input image.

A more common approach to object detection is to use a regression model, which takes an image as an input, and predicts zero, one or more (depending on the number of objects present) sets of 4 numerical values which define a bounding box around an object of interest. An example of a hypothetical image detection model of this type is shown in Figure 66, where, for the same input image of a single phytoplankton cell as used above in

Figure 65, an object detection model outputs values which define the corners of a bounding box around the cell.




Figure 66 – Example of a typical input and output from a machine learning regression model for object detection. An IFC image of a single phytoplankton cell is input, and the model predicts values for the coordinates of the corners which could be used to draw a box around the cell(s) in the input image.

Further, rather than simply predicting a bounding box (or identifying pixels belonging to an object, as carried out by the previously introduced object detection type), an object detection system could make use of multiclass or multilabel classifiers. By extending these algorithms from binary to multiclass classifiers, each positive detection can be assigned to one of several classes based on the highest probability. In this fashion, the system can simultaneously localise and assign labels to cells (e.g., species). These simultaneous localisation and classification object detection algorithms are commonly used for medical image analysis (Karaman et al., 2023), pedestrian detection (e.g. (Dollar et al., 2011)), automated satellite imagery analysis (e.g. (Hussain et al., 2013, Bakirman, 2023)) and robotics (Terven and Cordova-Esparza, 2023).

Examples of simultaneous classification and localisation of objects within images for the two types of object detection algorithms explored are explored below.

In Figure 67, an object detection algorithm takes as its input an MRI image of a patient's torso and outputs one probability density map for each class of interest, in this case, the heart, kidney, liver and spleen. From these probability density maps, the organs can be segmented from the original MRI, measured in terms of pixel area, etc.

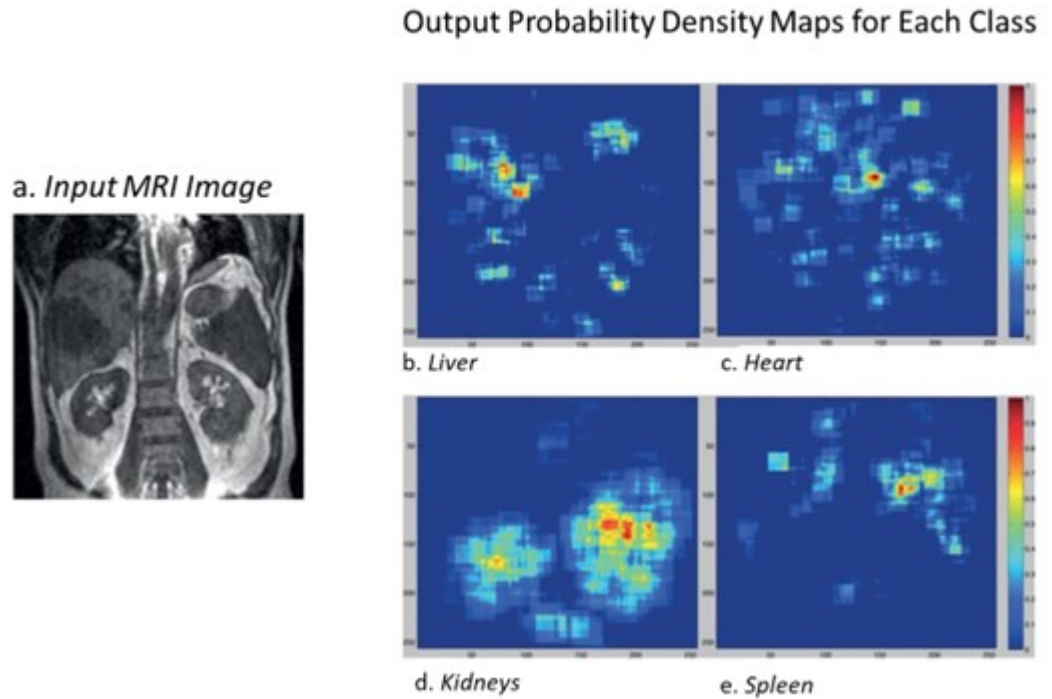


Figure 67 – This diagram, adapted from Shin et al. (2016a), demonstrates the ability of a machine learning model to output probability density maps which correspond to predicted locations of various classes of objects of interest. In this example, an abdominal MRI input image has 4 corresponding output density maps, one for each of the organs of interest.


Figure 68 presents the second type of object detection model. As opposed to the prior model which outputs a probability density map per class, this model simultaneously uses a trained regression model to predict the pixel coordinates defining a bounding box around objects and uses multiclass classification to assign a label value the objects. The label value can then be compared with a list of class names to draw labelled bounding boxes on the original image.

a. *Input Image*



b. *Object Detection output*

```
[xmin, xmax, ymin, ymax, class]
[130, 410, 300, 850, 1]
[90, 700, 100, 710, 2]
[680, 890, 40, 120, 3]
```



c. *Predicted bounding boxes and classes labelled on input image*

Figure 68 – Example of object detection via machine learning, where coordinates defining bounding-boxes are predicted via a regression model. In this example, the model has been trained to detect 3 different classes, and can simultaneously output a list (b.) of bounding boxes for an input image (a.) with a predicted class for each box (final column of the output list b.). Finally, the boxes can be drawn onto the input image to visually show the predictions of the model (c.).

3) Image Segmentation

Image Segmentation is conceptually similar to Object Detection but involves the assignment of a label to every pixel in the image. The two types of segmentation approaches, namely Semantic and Instance segmentation, are presented in Figure 69 below.

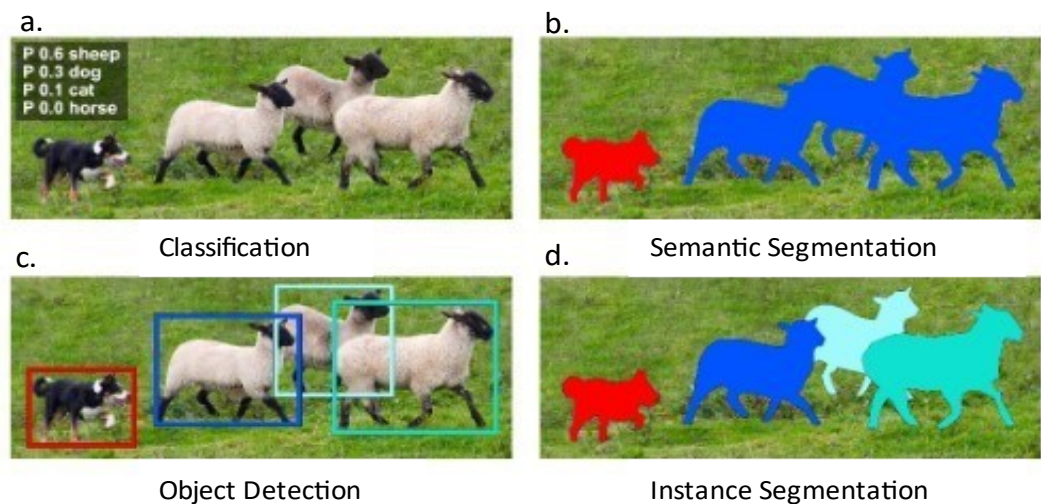


Figure 69 – Adapted from <https://manipulation.csail.mit.edu/segmentation.html>. An image of a dog and three sheep is used as an example of image classification, object detection, semantic segmentation and instance segmentation. In semantic segmentation, each pixel is assigned to the class with the highest probability. In the given example, red is used for the ‘dog’ class, and blue is used for the ‘sheep’ class. Pixels in the background are not separately coloured in this case, but would have been assigned the ‘background’ class. There is no distinction made between different instances of the same class, unlike in instance segmentation, which is shown below. In instance segmentation, each pixel is assigned a class as before, but also assigned an instance variable representing the number of that object. In this manner, the image pixels belonging to the three different sheep are separately labelled and can be extracted or analysed individually.

Semantic segmentation refers to the task of assigning a semantic class label, such as ‘sheep’, ‘dog’, ‘cat’, ‘background’, etc., to every pixel in an image. The purpose of semantic segmentation is to divide an image into different regions, each representing a unique object or background class. As shown in Figure 69 c. and b., while Object Detection approaches draw a bounding box around each instance of a detected class (dog and sheep), semantic segmentation assigns the label ‘dog’ or ‘sheep’ to every pixel comprising those instances.


Instance Segmentation extends semantic segmentation by assigning a unique identifier to each instance of each class. As demonstrated in Figure 69 d, this allows the same label to be assigned to pixels belonging to the same object, and different objects of the same class are given a unique label.

Image segmentation can be achieved based on, for instance, the pixel probabilities within a probability density map like that illustrated in Figure 65 being above a threshold value, with those labelled as belonging to a cell and those below the threshold labelled as background.

4) Image Generation

In image generation problems, the goal of the machine learning system is to generate a new image based on an input. The input may take the form of a prompt (*Text-to-image*), as demonstrated in Figure 70, where the Image generation algorithm *DALL·E 2* (<https://openai.com/product/dall-e-2>) was used to generate images based on short descriptive prompts defining the desired image content.

'A light-brown, medium-sized Anatolian shepherd dog sitting in a field of bluebells'



'A phytoplankton cell floating in seawater'




Figure 70 – two images which were generated by the DALL·E 2 image generation model. In each instance, a prompt, which is displayed here above the image, was used to generate a new image using a Generative Adversarial Network. The model is trained on many

Alternatively, image generation may take as an input one or more images, and generate a new image based on the inputs, as demonstrated in Figure 71, where DALL·E 2 has been used to generate two new images which are variants of an image given as an input.

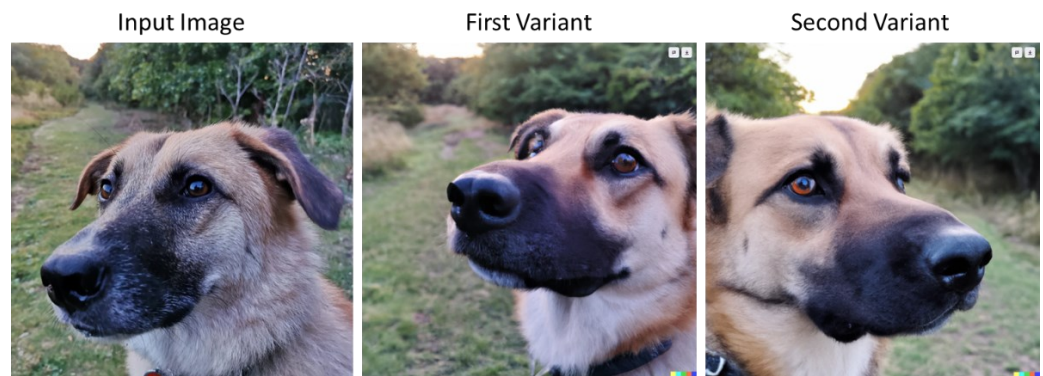


Figure 71 – Two variants of an input image generated by the generative adversarial network model DALL·E 2. Neither of the two variants are real images but instead were generated using a combination of previously learned features and features from the input image.

While image generation is not within the scope of this thesis, it holds potential in the generation of synthetic data that could be used to augment the training datasets used within other machine learning techniques relevant to IFC of phytoplankton. For example, image generation of unusually shaped phytoplankton cells could be used to train an object detection network to more reliably detect previously unseen phytoplankton within images from the flow cell.

The primary requirement of an automated image analysis system for the presented IFC is to count the number of cells present per image, as this produces the most fundamental biological measure of aquatic primary productivity. The secondary requirement is to localise the cells and store the cell coordinates, thus allowing later analysis of secondary biological characteristics such as cell volume, species, and health. Though not strictly required for the use of the IFC for basic phytoplankton science, it would be beneficial if the analysis software was able to measure these secondary properties online; that is, at the time of capture of the image, rather than after capture. This problem formulation most closely aligns with the description of Object Detection as given above, so for the remainder of the chapter the primary focus will be on the Object Detection problem of locating cells within IFC imagery.

As discussed previously, algorithms for image analysis rely on a set of features; while for classification problems an analysis is formed based on the presence or absence of these features

Chapter 4

within an image, object detection involves searching across an entire image for these features. It was also noted that hand-crafting these features is time-consuming, complex and potentially unreliable, but early object detection algorithms were based upon these feature sets and achieved moderate success; e.g. ((Papageorgiou et al., 1998, Rodenacker et al., 2006, Häder, 1995)).

Neural networks, particularly deep convolutional neural networks, which will be explored in detail in the following section, have revolutionised the field of image analysis by removing almost entirely any manual selection of features within the processing algorithms. By learning not just the optimal values for a given set of features but instead which features themselves are relevant, neural networks are able to model very complex relationships between the image data and the labels of interest (e.g., cell coordinates, species). In addition to determining, through the process of machine learning, the optimal feature set for a given problem, neural networks have several other benefits, including the speed with which they can be executed on computational hardware through parallel processing; and the ability to be robust to noisy input data. Neural networks will be the focus of the remainder of the chapter, and will be described in the next section.

4.3.2 Neural Networks

Neural networks are machine learning models which comprise matrix operations carried out by *layers* of interconnected nodes called *neurons*, each of which perform a weighted sum of their inputs and add a *bias* value. An example of the operation of a single neuron is presented in Figure 72, where a vector of inputs, $a = [a_1, a_2, a_3, \dots, a_N]$, is multiplied by a vector of *weights*, $w = [w_1, w_2, w_3, \dots, w_N]$, has a *bias*, b , added, and outputs the result.

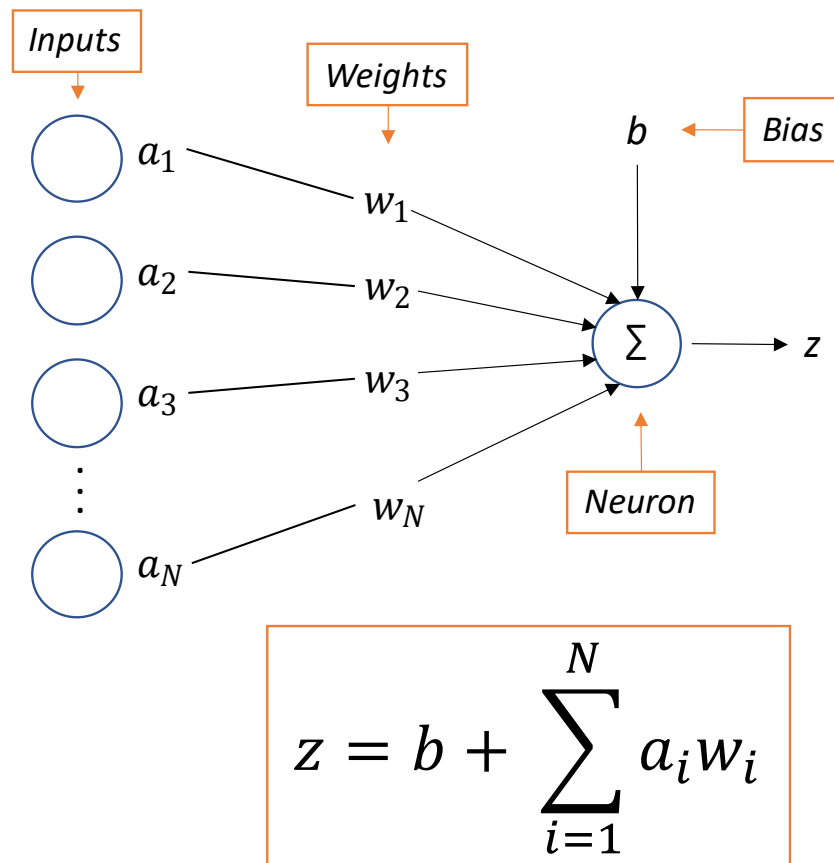


Figure 72 – Illustration of the operation of one neuron. The neuron computes a weighted sum of its inputs, $[a_1, a_2, a_3, \dots, a_N]$, using associated weight values, $[w_1, w_2, w_3, \dots, w_N]$. A bias value, b , is added and the output, z , is passed on to the next connection.

In a neural network, a *hidden layer* is an intermediate layer of neurons that sits between the input layer and the output layer. Its name derives from the fact that its neurons are not directly connected to the network's input or output layers. Instead, a hidden layer receives inputs from the previous layer and sends outputs to the next layer. Figure 73 demonstrates the simplest form of neural network; an *Input Layer* is multiplied by the matrix of weights within one single hidden layer, and the output is combined (using a weighted sum) to produce a result. Note that not all nodes of the previous layer need to be connected to every neuron in a given layer.




Figure 73 – An illustration of an extremely simple neural network. Two input neurons (blue) are connected to 4 neurons (grey) within a single hidden layer, which are in turn connected to a single output neuron (orange). Each connection represents one of the inputs to a matrix operation.

By modifying the weight and bias values for each of the neurons within a network, the neural network presented can learn linear relationships between the input and output variables and is therefore equivalent to a linear regression model. A practical example of regression using these neural networks could be modelling an expected population of plankton value based on numerical inputs, such as the water temperature, available photosynthetically available radiation, etc.

Within neural network terminology, values which are optimized through the training process (such as the weights and bias values of neurons) are called parameters, and those which are set manually (such as the number of hidden layers, or number of neurons per layer) are called *hyperparameters*. Note that hyperparameters can be either hardcoded directly or computationally optimized through the processes of network architecture exploration and hyperparameter tuning.

Chapter 4

An addition to the model described so far which enables the modelling of much more complex relationships between input and output values is that of nonlinearity. Nonlinear operations are added in the form of *activation* layers, which modify the output values of neurons in a nonlinear way.

Commonly employed activation functions include the *sigmoid function*, which enables neurons to perform an operation equivalent to logistic regression, and the *Rectified Linear Unit* function (ReLU), which sets all negative neuron outputs to zero, effectively creating a threshold that must be surpassed in order for the neuron to activate. This allows the network to model non-linear decision boundaries and capture more complex patterns in the data (Krizhevsky et al., 2017).

Figure 74 shows plots of these functions.




Figure 74 – The sigmoid and Rectified Linear Unit (ReLU) functions. By passing the output value of a neuron through these functions, the neural network is no longer performing a simple linear regression and is able to represent complex nonlinear relationships between input and output data. The specific activation function used influences the performance of a given model.

Chapter 4

In neural networks used for regression problems (including those modelling complex nonlinear relationships), the output layer consists of one or more neurons, each with a linear activation function outputting a continuous value. Different types of output layers can be used to produce various output types depending on the nature of the problem. For example, in a multiclass classifier, the output of the final hidden layer is often passed through a *softmax* activation function (Bridle, 1989), which produces a probability distribution over the different classes. The predicted class is then identified as the one with the highest probability. Other types of output layers, such as recurrent layers, can be used for sequential data such as time series.

In a *deep neural network*, multiple additional hidden layers are connected between the input and output of the network, and the operations of each layer are performed sequentially. A simple example of a deep neural network with 3 hidden layers is presented in Figure 75. Modern neural networks demonstrating success on visual recognition, natural language processing and other complex tasks often have dozens or even hundreds of hidden layers (Zhu et al., 2018). Deep neural networks have been shown to be highly effective at learning complex patterns in data, having achieved state-of-the-art performance on a wide range of tasks, including image classification, object detection and speech recognition (Sharma and Guleria, 2022).




Figure 75 - A deep neural network with 3 hidden layers is presented. In this network, each layer is fully-connected; that is, each neuron is connected to every neuron in the prior layer. By introducing additional hidden layers of neurons, the deep neural network is able to model increasingly complex relationships between input and output data. (Strauß, 2018)

Chapter 4

By setting the weight and bias terms for each neuron within a network of an appropriate architecture, the network can be enabled to provide an accurate label (\hat{Y}) for a previously unseen input (Y), providing the input shares similarities with the input data on which the network was trained (i.e., the weights and biases determined). The training process as a whole represents the optimisation of a nonlinear function to best map input to output data, according to some criterion.

One advantage of neural networks is that they comprise many small, simple mathematical operations compared to other machine learning algorithms, which can require complex singular operations (Bishop, 1994). This aspect allows computers to make use of *parallel processing*; that is, to simultaneously calculate the results of every neuron within a layer. Although traditional central processing units' (CPU) architectures are not efficient parallel processors; generally having only between 2 and 128 computing cores, graphics processing units (GPUs) are comprised of very high density, fast but simple computing cores (Hu et al., 2022). A modern consumer GPU (e.g., NVIDIA's RTX 3080 graphics card released in 2020) has over 8000 cores, each of which can operate at a rate up to 1.7 GHz.

As discussed previously, machine learning algorithms must go through a process of training to determine optimal parameter values. Neural networks are trained using many labelled data examples assembled into a training dataset. The number of training data required depends on the complexity of the mapping function between inputs and labels. Once sufficient training data is obtained, the weight and bias terms of each neuron must be iteratively optimized through the following process:

- a) A random input-label pair (Y, \hat{Y}) is selected from the training dataset.
- b) The input is fed to the first layer of the neural network and each neuron performs its $\alpha x + \beta$ operations, including the activation function (if present). The output is then passed to zero, one, or more neurons in the next layer of the network. This process continues until the final layer of the network, where the combination of neuron output values corresponds to a particular label or regressed value for the input data. This process is known as *forward propagation*.

Chapter 4

- c) A *Loss* value is calculated by comparing the label output by the network with the real label. This Loss can take the form of *Mean Squared Error* for regression problems, *cross-entropy loss* for categorisation problems, and so on. Many such loss functions have been proposed within the literature and the loss function used will depend on the task type and priorities (accuracy vs precision, etc).
- d) Starting from the last layer of the network, the backpropagation algorithm is used to calculate the gradients of the Loss function with respect to the weights and biases of each neuron in the network. Backpropagation uses the chain rule of calculus to compute the derivative of the Loss function with respect to each parameter in the network, including the weights, biases, and any other learnable parameters. This process is repeated backwards through the layers until the gradients of all the parameters have been computed.
- e) Finally, the parameters of the entire network are simultaneously updated using a *gradient descent algorithm*. The simplest of these is the stochastic gradient descent algorithm:
- $$Wx = \widehat{Wx} - \alpha \left(\frac{\delta L}{\partial \widehat{Wx}} \right),$$
- where Wx is a new weight, \widehat{Wx} is an old weight, α is a *learning rate* parameter and L is the loss value. Biases and another other learnable parameters are calculated similarly. For deep neural networks, more advanced optimization algorithms such as *Adam* or *RMSProp* may be used instead. These advanced algorithms use techniques such as *momentum*, *adaptive learning rates*, and *gradient normalization*, which are outside the scope of this thesis.

After training on a specified number of input-label pairs (*labelled data*), the neural network is typically tested on labelled data it has not been trained on. This process is known as *validation* and requires withholding a portion of the labelled data from the training dataset to form a validation dataset. The loss values calculated for the operation of the network on these validation data are used to ensure that the network does not fit its weights and bias matrices too specifically to the input data, a phenomenon known as *overfitting*. When overfitting occurs, the network has become too specialized to the input data and will not generalise well to unseen data. The process of training on the training dataset and then validating on the validation dataset is typically

Chapter 4

referred to as an epoch, and the total process of training may require several to thousands of epochs to model a complicated input-output relationship.

In modern neural network training algorithms, batches of training data are used to update the trainable parameters of the network, rather than updating the parameters based on a single input-label pair (or the entire dataset). By randomly sampling batches from the training data at each iteration of the training algorithm, the model is trained on a diverse set of examples, which can help prevent overfitting and improve the generalisation of the trained model. Generalisation, in the context of machine learning models, refers to the ability of a trained model to effectively perform and make accurate predictions on new, unseen data rather than just memorising the training data. Without batch training, the neural network would be more susceptible to learning incorrect mappings based on noise or the features of individual inputs. The size of the batches used is a hyperparameter, with the maximum possible batch size being determined by the amount of memory available to the training hardware.

It is important to ensure that a network has good performance (accuracy, precision and any other metrics) while not being overfit to its training data. The results of validation tests over time help quantify the performance of the model as it is trained and can be used to terminate the training process when a certain level of performance, according to chosen metrics, is achieved. A brief overview of the most fundamental performance metrics, many of which are used in the following sections, follows:

1) Inference Time

The time taken to process a single input into a corresponding output, typically presented in milliseconds (ms). This time is hardware-dependent but can be used for comparing the runtime efficiency of different models on the same computer.

2) Confidence Score

A confidence score is the degree of certainty that a machine learning model has in its prediction or classification. It is expressed as a value between 0 and 1, where higher scores indicate greater confidence. Confidence scores are useful for decision-making when accuracy

is critical, and low confidence scores can indicate cases where additional verification or correction may be necessary. Typically, a confidence threshold value is set such that predictions with confidence scores above the threshold are accepted, while those below are rejected or flagged for manual review. The choice of threshold can be adjusted to balance the trade-off between accuracy and the risk of false positives (or false negatives).

3) Accuracy, Precision, and Recall (*Classification models*):

Each of these metrics can be defined in terms of True Positives (*TP*), False Positives (*FP*), True Negatives (*TN*), and False Negatives (*FN*).

The *accuracy* of a classifier refers to the number of correctly classified samples out of the total number of samples in the dataset. The formula for accuracy is therefore: $(TP + TN) / (TP + FP + TN + FN)$. *Precision* refers to the proportion of correctly predicted positive samples out of all the samples that were predicted as positive by the model. The formula for precision is given by: $TP / (TP + FP)$. *Recall* measures the ability of the model to identify all positive samples correctly out of all the actual positive samples in the dataset, and is computed using the formula $TP / (TP + FN)$. Recall is often used in combination with precision to evaluate the performance of a classification model, as together they provide a more complete understanding of the model's performance in terms of both positive and negative predictions.

4) *F1 Score* (*Classification models*)

This metric combines both precision and recall to give an overall measure of a classification model's performance. The F1 score is the harmonic mean of precision and recall and is calculated as follows:

$$\text{F1 Score} = 2 * ((\text{Precision} * \text{Recall}) / (\text{Precision} + \text{Recall})).$$

The value of F1 ranges from 0 to 1, with a value of 1 indicating perfect precision and recall, and a value of 0 indicating that the model is not able to correctly classify any of the samples.

5) *Mean Absolute Error* (MAE) and *Mean Squared Error* (MSE) (*Regression models*)

MAE measures the mean absolute difference between the ground-truth values and the model-predicted values of the target variable. Also referred to as the L1 Loss function, it can be expressed as shown in Equation 7:

Equation 7 – Equation for Mean Absolute Error (MAE), used to calculate error in a regression machine learning model.

$$MAE = \frac{\sum_{i=1}^n |y_i - \hat{y}_i|}{n}$$

Where n is the total number of samples in the dataset; y_i is the real value of the target variable for the i -th sample; \hat{y}_i is the predicted value of the target variable for the i -th sample. In the implementation of Object Detection neural networks, MAE can be used as a Loss function for the regression of the coordinates of objects of interest during training, and for evaluating performance of the model when testing.

MSE, which is the mean squared difference between the true value and model-predicted value of a target variable, is more sensitive to outliers than MAE and is useful for object detection tasks, where small errors in detected coordinates are much less important than larger ones. The equation for MSE is given below in Equation 8.

Equation 8 - Equation for Mean Squared Error (MSE), while like MAE can be used to calculate the error in a regression-based machine learning model.

$$MSE = \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{n}$$

Where n is the total number of samples in the dataset; y_i is the real value of the target variable for the i -th sample; \hat{y}_i is the predicted value of the target variable for the i -th sample.

MSE will be used when implementing an object detection neural network in Section 4.4.3, where a model is trained to predict output pixel values according to the confidence of the network that the input pixel represents a phytoplankton cell when presented with input images from the IFC.

- 6) *Intersection over Union* (IoU) (Bounding-box regression models)
Intersection over Union (IoU) (Figure 76) is a commonly used evaluation metric in object detection tasks. It measures the overlap between a predicted bounding box and the ground-truth (human labelled) bounding box of an object in an image.

To calculate the IoU between two bounding boxes, one first computes the area of their intersection, which is the area where they overlap, and the area of their union, which is the total area covered by both boxes. The IoU is then defined as the ratio of the intersection area to the union area.




Figure 76 – Intersection over Union (IoU) is a performance metric which calculates the accuracy of a bounding box prediction by dividing the intersection (overlap) by the union (area enclosed by both boxes), for the predicted and the ground truth bounding boxes.

- 7) *Average Precision* metrics (e.g., AP50, AP95) (Bounding-box regression models). These metrics refer to the average precision (AP) calculated at a given IoU threshold, with this threshold value indicating the percentage of overlap required between predicted and ground truth bounding boxes for a detection to be considered correct. Threshold values of 50% and 95% are commonly used to compare the performance of different models, and are abbreviated to AP50 and AP95 respectively.

In object detection, each ground-truth object in the dataset will have multiple predicted bounding boxes with varying confidence scores. To evaluate the model's performance, all model predictions are sorted by confidence, from highest to lowest. For calculating AP at a given IoU threshold (e.g., AP50), the model's predictions across the entire testing dataset are first sorted, and then precision values are computed for every recall level, considering the chosen IoU threshold. After these precision values have been calculated, the results are averaged. This process helps to evaluate the overall performance of the model, as it captures the trade-off between precision and recall.

Hence, the AP50 metric is the average of the precision scores at all recall levels for all the objects in the testing set, with an IoU threshold of 50%. A higher AP50 value indicates that the model has better accuracy in detecting objects in the test set. In contrast, AP95 measures the average precision at the higher IoU threshold of 95%, which means that the predicted bounding boxes must have a greater overlap with the ground truth bounding boxes to be considered valid detections. AP95 is therefore more rigorous than AP50 and provides a stricter evaluation of the model's performance.

AP50 and AP95 will both be used in Section 4.4.6, where an object detection neural network will be trained to predict the coordinates of bounding boxes enclosing phytoplankton cells when presented with images from the IFC system.

4.3.3 Convolutional Neural Networks

Convolutional Neural Networks (CNNs) are a type of neural network particularly adapted for multidimensional input data, such as images, which are stored digitally as matrices of dimensions [Height x Width x Channels]. Typically, colour images are represented in 3 channels (Red, Green, Blue) whereas grayscale images are stored in a single channel (brightness).

CNNs build upon the neural networks described in Section 4.3.2 by introducing the convolution operation, which is conceptually identical to applying a sliding window filter to the input. The filter is a matrix of dimensions width x width x channels, and in a CNN each element of the filter matrix is a neuron whose weight and bias values can be optimised through the backpropagation process. The convolution operation calculates the dot product of the filter and the image at each position to produce a new image, known as the convolved image. This principle is demonstrated in Figure 77, which shows how the destination pixel in a convolved image is a product of multiplications of source pixels with the values of a convolutional filter.




Figure 77 – A visual representation of a convolution operation on an image represented as a 2D matrix (left). Convolution computes the value of each destination pixel in an output matrix (right) by multiplying each value in the convolution filter (centre) by the corresponding value in the input image, and then sums the results. This process is repeated for every pixel in the input image, resulting in a new output image where each pixel is the result of applying the convolution filter to the corresponding pixel in the input image. Within a Convolutional Neural network, the inputs to a convolution, the outputs from the operation and the filter itself can be of any number of dimensions, with CNNs having many layers potentially having hundreds of dimensions within some layers.

An important mathematical property of CNNs is that the convolution operation in the spatial domain is equal to the multiplication of the image and the filter in the frequency domain, and vice versa. This property allows CNNs to be efficiently executed by a computer.

Convolutional Neural Networks (CNNs) have been demonstrated to provide accurate, robust, and fast object detection within images (Girshick et al., 2014, Shin et al., 2016b). Compared with the traditional image processing techniques described in section 4.2, CNNs are more accurate and robust to different lighting and background conditions (Bhatt et al., 2021). They can also be run at

Chapter 4

high speed on specialized hardware such as graphics processing units (GPU) and tensor processing units (TPU), making them a powerful tool for image analysis.

In addition to convolutional layers, CNNs that achieve state-of-the-art results on common image analysis benchmarks, such as ImageNet (Deng et al., 2009). These will be explored below.

In addition to convolutional layers, CNNs often use pooling layers. Max pooling (Ranzato et al., 2008) is a sample-based downsampling method that reduces the size of the feature map in the output of a CNN layer, as demonstrated in Figure 78, which shows how a 4x4 matrix is downsampled to 2x2. The size of the max pooling operator defines the scaling factor of the output, which is calculated by retaining only the highest value from each region of the input. This not only reduces the number of parameters in the model, but also helps to prevent overfitting. As Max Pooling layers reduce the spatial dimensions of the input layer, they also have the benefit of increasing computational efficiency for the subsequent layers of the CNN.

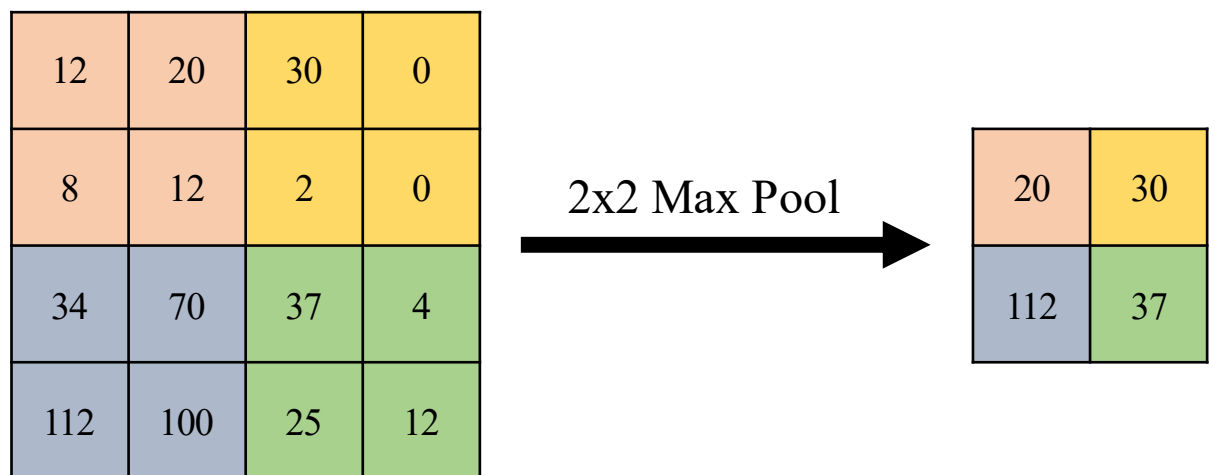


Figure 78 – An illustration of a 2x2 Max Pooling operation on a 4x4 matrix. The original matrix is divided into 2x2 regions and only the maximum value from each region is retained in the output matrix. Therefore, the output is of shape 2x2, and the input has been downsampled.

Other types of pooling, such as average pooling, are occasionally used and are conceptually very similar to max pooling, except that they use a mean (or other function) of each region instead of the maximum value. These types of pooling layers can also be useful for reducing the size of the

feature map and improving computational efficiency in CNNs. However, max pooling is more commonly used in practice due to its superior performance on many image analysis tasks (Bieder et al., 2021).

Another commonly used technique in modern CNN architectures is the *normalisation* of neuron outputs. One method for this is *Local Response Normalisation* (LRN), which normalises the output of a single convolutional layer. LRN adjusts the output of neurons by taking into account the values of adjacent neurons using the following formula shown in Equation 9.

Equation 9 – Equation for Local Response Normalisation (LRN). LRN is a normalisation technique used in neural networks to enhance convolutional layers' outputs by normalising activations across channels. The formula below computes the normalised output, promoting competition among neighbouring neurons and improving generalisation.

$$b_{x,y}^i = a_{x,y}^i / \left(k + \alpha \sum_{j=\max(0,k-n/2)}^{\min(N-1,i+n/2)} (a_{x,y}^j)^2 \right)^\beta$$

Where: $a_{x,y}^i$ is the input activation of a neuron in a feature map (the output from the activation function of the previous layer); $b_{x,y}^i$ is the output activation of the same neuron after LRN; x, y are the spatial coordinates of the neuron in the feature map; i is the channel (layer dimension) of the neuron; k is a bias for the LRN (a hyperparameter); N is the total number of feature maps in the layer; n is the size of the normalisation window (a hyperparameter); α and β are hyperparameters that control the strength of the normalisation.

The output of the LRN operation for a given neuron is thus simply the output of the previous layer, divided by a scaled summation of the neurons within the column formed by moving through the depth dimension of the input feature map at given spatial coordinates. The rationale for using LRN is to model a form of lateral inhibition similar to that found in real neurons, and using LRN has been demonstrated to improve model generalisation and hence overall accuracy (Krizhevsky et al., 2017).

LRN can help to improve the generalisation of CNN models by increasing their ability to recognize patterns across different images, and they can also help to prevent overfitting. However, more recent techniques such as *Batch Normalisation* (BN) have been shown to be more effective for normalizing the outputs of CNNs, and have largely replaced LRN in modern CNN architectures.

Batch normalisation is a technique used in deep CNNs that relies on batch training, where the training data is divided into small subsets, called batches, and the model is trained on each batch in turn. Unlike Local Response Normalisation (LRN), which normalises the neuron response across the layers in a feature map generated while processing a single input sample, batch normalisation works to normalise responses across an entire batch of training data. Many state-of-the-art models for image analysis use batch normalisation because it reduces the sensitivity of the network to the choice of hyperparameters, such as learning rate, allows for larger learning rates that can speed up training (Ioffe and Szegedy, 2015), and is more effective than LRN at preventing the model from learning features based on noise and outliers in the training data (Samir et al., 2020).

The formula for BN is given in Equation 10:

Equation 10 – Formula for Batch Normalisation (BN). is a technique used to improve the training process of neural networks by normalising the input features in each batch. The formula below calculates the normalised values, enhancing the model's stability and generalisation.

$$\text{BN}(x_i) = \gamma \frac{x_i - \mu_B}{\sqrt{\sigma_B^2 + \epsilon}} + \beta$$

The formula is for a single neuron, where: (x_i) is the input activation from the preceding layer for a given sample in the batch; μ_B is the mean value across the batch; σ is the standard deviation of the batch; γ and β are learnable parameters for scale and shift, that allow the general model to learn the optimal scale and mean for each feature; ϵ is a small constant (typically set to a very small value like 1e-5) to ensure numerical stability.

4.4 Implementation of CNN Processing of Phytoplankton Images

4.4.1 Introduction

This section introduces CNNs for processing and analysing phytoplankton images obtained using the novel acoustically-focussed IFC. As discussed in section 4.3, CNNs are a machine learning technique that has demonstrated impressive performance in a wide range of image classification and object detection tasks. By training a CNN on a large dataset of annotated IFC images, it is possible to create a model that is capable of accurately localising and/or classifying phytoplankton cells within new images.

In the context of this thesis, the purpose of using CNNs is to analyse the complex and varied images of phytoplankton cells captured by the IFC. While CNNs can be trained to analyse complex attributes of the phytoplankton, such as cell species and volume, the primary interest here is using them for the simplified problem of accurately counting the number of cells within water samples. Cell count is a fundamental measure of algal productivity and is critical for deriving other data such as growth rate over time. By developing an automated system that can accurately count the number of phytoplankton cells, it is possible to increase the number of samples that can be analysed compared with traditional techniques; automatically monitor HAB events and manage aquatic ecosystems.

The following subsections describe the implementation of CNN-based processing of phytoplankton images and present the results. After a discussion of the computer hardware on which the CNNs were trained and run, a description of the initial approach to investigating a novel CNN-based processing method is presented, followed by an implementation of the YOLOX model and its training and performance evaluation. Overall, it is demonstrated that CNN-based processing represents a promising approach for automated analysis of IFC images, with the potential to significantly improve accuracy and speed compared to traditional image processing techniques.

4.4.2 Hardware

After an image frame has been captured by the camera in the IFC, the image data must be either stored for later analysis (offline processing) or analysed in real-time (online processing). So long as sufficient storage capacity is available and there is enough time to process images before they are overwritten, offline processing is preferable as the additional processing time made available allows the employment of more sophisticated algorithms. *In situations* where storage or file-transfer bandwidth is limited (e.g., devices deployed at sea for extended periods), online processing is a necessity as it allows the storage or transfer of processed numerical results, which are of a far lower file size than complete raw images.

While the networks which will be introduced in sections 4.4.3 and 4.4.5 were trained on a high-performance, GPU-enabled desktop PC (Intel i7-9700k, NVIDIA RTX 3080), the IFC platform presented within this thesis is intended to be flexible to either *in situ* deployment or lab-based use, where such hardware is prohibitively bulky, power-demanding and expensive. Therefore, a small, low-cost, low-power computing platform which allows both the online analysis of image data and the storage of a high volume of images was sought after. This resulted in the selection of the NVIDIA Jetson Xavier AGX platform, which is a powerful and energy-efficient system-on-module (SoM) that is designed specifically for running machine learning models. It features an 8-core CPU, an NVIDIA Volta GPU with 512 CUDA cores, and 32GB of RAM, making it capable of running complex CNN-based models for image analysis.

The Jetson Xavier AGX is designed for use in autonomous systems and is small enough to be integrated into compact and portable devices. This makes it an ideal platform for developing a portable and robust system for in-situ analysis of IFC images. Such a system could be used to analyse water samples in the field, without the need for sending samples back to a lab for analysis. The Xavier AGX also costs under £1,000, which is significantly less than the hardware needed to attain equivalent inference performance in a conventional computer.

NVIDIA also produces a Jetson SoM with a lower cost and lower power draw than the Xavier AGX, at the cost of reduced performance. This SoM, the Jetson Nano, costs around £100, making it ideal for enabling low-resource scientific environments to utilise cutting-edge IFC for

phytoplankton analysis. A table comparing the full specifications of the Jetson Xavier AGX and Nano is presented in Table 4.

Table 4 - Comparison of hardware features and cost of two NVIDIA devices in the Jetson family of embedded computers designed for machine learning:

Specification	Jetson Nano	Jetson Xavier AGX
GPU	NVIDIA Maxwell architecture with 128 NVIDIA CUDA cores	NVIDIA Volta architecture with 512 NVIDIA CUDA cores
CPU	Quad-core ARM A57	8-Core ARM v8
Memory	4 GB RAM	32 GB RAM
Storage	microSD	32 GB eMMC
Price (£ GBP 2022)	£100	£980

4.4.3 CNN Investigation for Phytoplankton Detection

With the primary objective to maximise analysis quality and the secondary objective to minimize inference time, a review of the literature uncovered a lightweight CNN-based object detection algorithm presented by Heo et al. (2017). The method proposed is a high-performance processing pipeline designed to count and track cells within images from a comparable IFC system, and was demonstrated (on their grayscale, 100x500 pixel IFC images) to be both fast, with inference at up to 500 frames per second, and accurate, with an error relative to human counts of 0.128%. The method proposed by Heo et al. is a simplified version of the *Fully Convolutional Regression Network* (FCRN) first demonstrated by Xie et al. (2018a). The pipeline which will be used in this investigation for IFC image analysis consists of a modified FCRN modelled after Heo et al., and a custom algorithm for processing the FCRN output to bounding boxes and cell counts.

For this work, this FCRN architecture was modified by integrating the modern batch normalisation technique explored in Section 4.3, and the final network consists of the layers presented in Table 5. An improved gradient training algorithm was also implemented, which will be discussed later. The input to the network is a 3-channel (RGB) image of pixel dimensions 512x512. The input is fed

to a first convolutional layer, with a ReLU activation function. The first convolutional layer has the same spatial dimensions as the input, but increases the channels to 32. The data is subsequently passed through a Max Pooling operation, which halves the spatial dimensions. Next, another convolution + ReLU layer increases the channels to 64, before a third convolution + ReLU decreases the channels to 32. At this point, a *Deconvolution* layer increases the spatial dimensions to 513x513; the deconvolution is simply the inverse of a convolution and can be used to up-sample the input. Finally, 2 more convolution + ReLU layers decrease the channels down to 1; the output from the network is a single-channel confidence map (or 'probability density map'). The spatial dimensions of the output are the same as the input image, such that there is a direct pixel mapping between the input and the generated probability density map. In total, the CNN presented has 52,147 trainable parameters.

Table 5 – Full CNN architecture used to localise and count phytoplankton cells within images from the acoustically-focussed IFC. The network consists of 5 convolutional layers, one deconvolutional layer and one max pooling layer. The input is 3-channel (full-colour RGB) and the output is a single channel probability density map where each pixel value represents the confidence of the network of the presence of a phytoplankton cell.

Layer Type	Dimensions
Input	(3*512*512)
Conv + ReLU	(32*512*512)
Max Pooling	(32*256*256)
Conv + ReLU	(64*256*256)
Conv + ReLU	(32*256*256)
Deconv	(32*513*513)
Conv + ReLU	(16*511*511)
Conv + ReLU	(1*512*512)

Heo et al, in their implementation, used a local response normalisation layer after the final convolutional layer, an addition to the original FCRN presented by Xie, Noble and Zisserman. As previously discussed, LRN has been shown to be less effective at solving the problem of overfitting

Chapter 4

and is more computationally expensive than batch normalisation; so the novel, improved approach taken here instead employs BN following each convolution operation.

The model, when applied to IFC images, solves the regression problem of, for each pixel, determining a 0-1 likelihood of it being the centre coordinate of a phytoplankton cell. The output from this neural network takes the form of a *probability density map* where each cell in the input image has a corresponding Gaussian distribution in the density map. The entire probability density map is represented as a mixture of Gaussians as in Equation 11:

Equation 11 – Equation describing the probability density map output of a regression-based object detection model.

$$Y = \sum_{k=1}^K N((i, j) | \mu_k, \Sigma_k)$$

Where: (i, j) is the pixel index of the input grayscale image matrix X ;

Y is the output probability density map;

K is the number of Gaussians in the density map.

$N((i, j) | \mu_k, \Sigma_k)$ represents a bivariate Gaussian distribution with mean $\mu = [\mu_x, \mu_y]^T$ and isotropic covariance $\Sigma = \sigma^2 I_{2 \times 2}$.

Model implementation, training and testing was carried out using Python 3.9, Pytorch 1.9.1 and CUDA 11.1.

To train the network, a training dataset was required. To obtain adequate accuracy, precision, and generalisation (the ability of the network to correctly process inputs it has not ‘seen’ before), a large number of labelled training examples is usually required. For instance, Shahinfar et al. (2020) found that a minimum of 150 training images was required to achieve good performance in the six different CNNs they investigated. In the case of our IFC images, 150 images were therefore labelled, with images chosen in such a way as to represent various possible background states, lighting conditions, cell presence/type/number/size, etc. The images were taken during IFC testing using a culture of *Rhodomonas salina* phytoplankton, and the dataset was curated such as

Chapter 4

to have examples of images with few cells (min $N = 3$), many cells (max $N = 55$), and numbers in between.

The training dataset consists of IFC images, and label images comprising pixels with the value 0 everywhere other than the centre of a cell, where that pixel is set to 1. Creating the training dataset was performed using LabelMe image annotation software (<https://github.com/wkentaro/labelme>). Of the 150 IFC images used for the training dataset; 22 (~15%) were withheld for validation during the training cycle.

Reducing the size of the images is a common practice in deep learning, as it enables batch training by decreasing the computational requirements and memory usage. Images were therefore resized to 512x512 pixels, which allowed training with a batch size of 16 on a PC with an Intel Core i7-9700K CPU and NVIDIA RTX 3080 graphics card with 10GB memory. In this application, the reduction in image size is unlikely to significantly impact the model's performance, as the essential features of the cells within the images can still be effectively recognized by both a human and the CNN at the lower resolution. Images were normalised such that the mean and standard deviation of pixel values across each of the 3 colour channels were the same for every image.

As discussed in section 4.3.2, a gradient descent algorithm is used in the backpropagation stage of the learning process in order to update the weights and biases of the neural network. Rather than the stochastic gradient descent algorithm (SGD) used by Heo et al. in their presentation of the model, for this implementation the Adaptive Movement Estimation algorithm, (Adam) algorithm (Kingma and Ba, 2014) was used as it has been demonstrated to reliably improve the speed with which the model converges to an optimal solution (Soydaner, 2020). Other than the previously mentioned improvement from using Batch Normalisation as opposed to Local Response Normalisation, this is the only other change compared to their implementation.

An example input image and output density map is presented in Figure 79. The three cells visible in the IFC image have corresponding gaussian probability distributions in the output image, demonstrating a successful inference by the object detection algorithm.




Figure 79 – Example of an input image to the network (top): a 1920x1200 pixel, colour image of the IFC flow cell with 3 phytoplankton cells (a close up of 1 is provided on the right). On the bottom is the probability density map output by the CNN, with each pixel having a value which represents the confidence of the CNN that there is a cell there.

Training was performed for 150 epochs, with the MSE loss for training and validation datasets recorded at each epoch. The resulting graph of MSE loss against training epoch is presented in Figure 80.

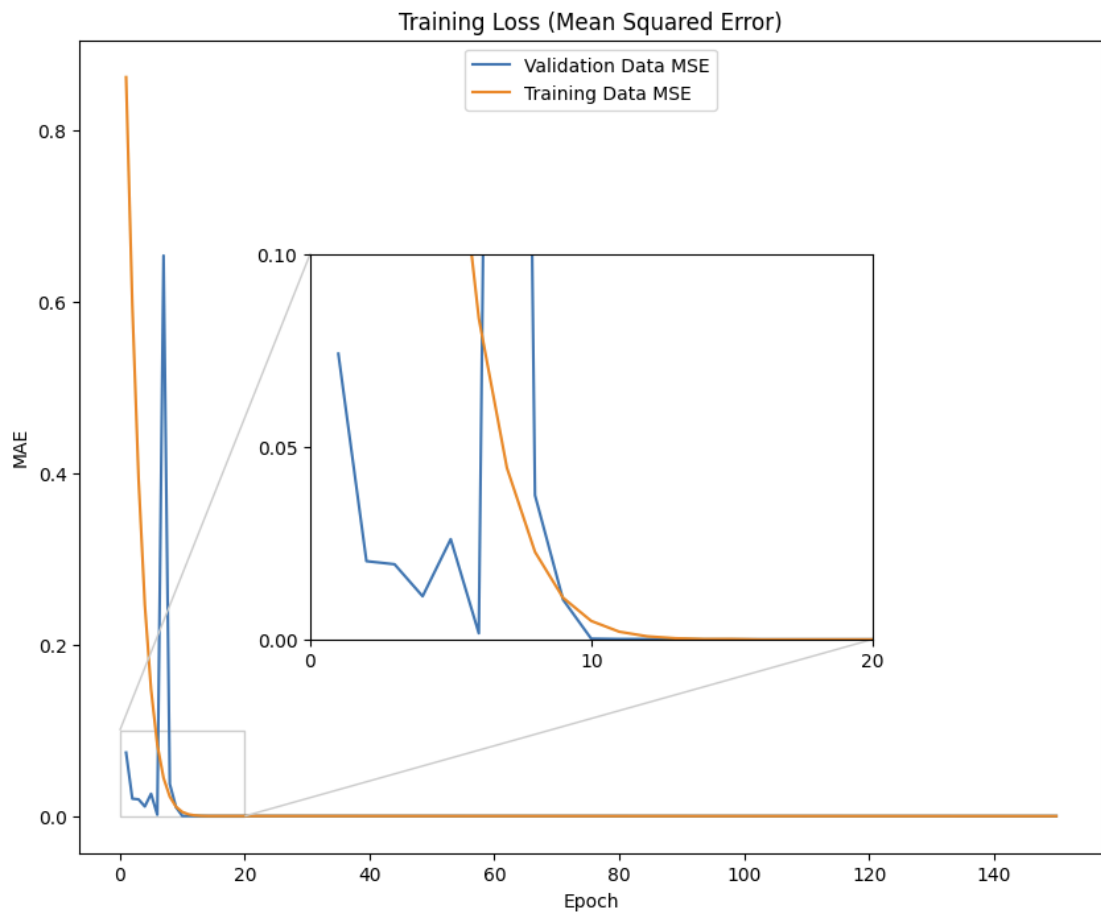


Figure 80 - Evolution of training (orange) and validation (blue) loss (Mean Squared Error, MSE) across 150 epochs for an object detection CNN, showing the prompt minimization of training loss, a transient surge in validation loss around epoch 8, and subsequent convergence to a value close to 0.

Several features of the training graph stand out. Firstly, there is a smooth decrease in the MSE on the training dataset during the first 15 epochs, after which the training loss is extremely small. The validation loss also reaches a very low value within these first epochs, which suggests that the model reaches a good level of generalisation without needing more than 15 epochs. There is, however, a significant spike in the validation dataset MSE at around epoch 8, which could be caused by several factors, including random fluctuations in the data, initialization of network parameters, or the learning rate used to update the model weights.

It is not uncommon to see spikes in an otherwise decreasing loss metric, especially when the number of training samples is relatively small compared to the complexity of the model, as is the case for the IFC training dataset. This can be caused by the model having overfit to prior batches

Chapter 4

of training data and then being tested a validation batch without the same characteristics or patterns as the previously seen training batches. As the training progresses and the model is exposed to more varied image batches, it learns to generalize better, which results in a gradual decrease in the loss metric. The fact that the validation MSE rapidly approaches zero after this spike suggests that the model has moved away from overfitting to any particular subset of the training data, and has instead learned the general features as desired.

In order to minimise overfitting to the data, the commonly employed (Mahsereci et al., 2017) 'early-stopping' strategy was employed, and so the mode state after 10 training epochs was saved. At this point in training, the model generalisation is good (as evidenced by the low validation dataset loss), but the training loss has not reached its asymptote, which would indicate possible overfitting to training data. This corresponds well to the findings of Heo et al. in their presented model, which converged to an asymptote after 7 epochs.

Chapter 4

After training, in order to perform cell detection on an IFC image, the following procedure is carried out:

- 1) *Resize image to 512x512 pixels.* To process an image, the input image size must match that used during training, which was 512x512 pixels.
- 2) *Normalise the image.* To account for changes in illumination brightness between different images (for instance, if incident light from the environment onto the IFC increased resulting in brighter images), the algorithm normalises the image to have the same mean and standard deviation of pixel values in each of the 3 colour channels as those on which it was trained.
- 3) *Perform inference using the trained neural network.* The normalised image is fed forward through the CNN using the weights and biases learned during training, producing an output probability density map.
- 4) *Apply a noise cancellation threshold to the entire density map.* Background pixels (negative detections) are typically weakly positive rather than zero in the density map, so we first set all pixels below a noise threshold value to zero. The noise threshold was selected to be the 5th percentile of the pixel values of the density map to capture positive detections but exclude very low-confidence *false negatives*.
- 5) *Binarize the density map.* A threshold must be applied to the probabilities such that those below a detection confidence threshold are set to zero and those above set to 1. The threshold must be manually set by the user, but (in theory) should not have to be changed once an optimum value has been determined, since as the input images were normalised, the probability density maps should have similar pixel values regardless of small changes to illumination intensity, etc.
- 6) *Count connected regions.* Cells are counted based on the principle of pixel connectivity, i.e., all pixels in a connected component share similar pixel intensity values and are in some way connected with each other.

Figure 81 demonstrates the probability density map output by the trained model and the subsequently binarized output image for a typical IFC image input, which has been resized to 512x512 pixels.




Figure 81 – Input image (left), probability density map representing model confidence of the presence of a cell (middle) and binarized object mask after thresholding the density map (right).

If rather than simply counting cells, one wished to identify their locations on the original image, allowing one to subsequently draw bounding boxes, crop the cells out of the full-frame and/or use a secondary algorithm such as a species classifier on them, one can scale the binarized probability density map to the same dimensions as the input image (1200x1920 pixels) and overlay it onto the original image. Figure 82 illustrates this concept; for an input image with 3 cells, the binarized density map overlaid onto the input correctly indicates the locations of the cells.




Figure 82 – Input IFC image (a.) and the same image overlaid with a binarized output mask of the CNN-predicted cell locations (b.), illustrating how the 3 cells visible can be easily localised based on the probability density map output by the model.

4.4.4 Evaluation of CNN Performance

All tests were carried out using the same hardware setup used during training. Images were loaded from a Western Digital Black SN750 1TB Solid State Drive (SSD). The model was run in GPU-optimised (CUDA) mode and was tested on 500 1200x1920 pixel RGB images from the IFC. The program completed in 4.27 s, achieving an average framerate of 117.01 frames per second including image pre-processing. This is significantly slower than the inference speed of the model as presented by Heo et al., but the source of this performance decrease is not the model itself but the resizing of the images from 1200x1920 to the 512x512 pixels required by the model, combined with the fact our images are 3-channel (colour) compared to their single-channel images. This decreased in speed could potentially be mitigated (if required) if the images were saved at a lower resolution at the time of capture; while the FLIR camera used in the IFC has maximum full-frame pixel dimensions of 1200x1920, it is capable of capture at lower resolution (which also improves the framerate).

To assess the accuracy of the model at counting phytoplankton cells, a comparison with manual counting of the cells within images was performed. *Rhodomonas salina* cells were counted by hand from a set of 50 IFC images containing varying numbers of cells (hence referred to as *Image Set 1*), and these images were then passed through the CNN to automatically predict cell counts.

The descriptive statistics summary of result of this comparison is presented in Table 6. The comparison demonstrates a large difference between the total number of cells counted manually (593) and those counted by the CNN system (476); difference = 117 (~20.0% of manual counts).

Table 6 – Table of descriptive statistics of the cell counts from 50 IFC images of *R. salina* cells, with cells counted either manually or by the CNN presented within this section.

<i>Image Set 1</i>		
<i>N=50</i>	Manual	CNN
Sum of all cells across all images:	593	476
Mean cells per image	11.86	9.52
Minimum cells per image	7	6
Maximum cells per image	20	17

To statistically test the correlation between the two methods, the Pearson's Correlation Coefficient was calculated for the per-frame counts generated by each counting technique. The resulting correlation coefficient (r) was 0.881 (3 s.f.), with a p-value of $3.13e-17$ (3 s.f.). The high correlation coefficient indicates a strong positive linear relationship between the cell counts obtained by the manual and CNN methods. Moreover, the extremely low p-value (significantly below the common threshold of 0.05) demonstrates that the observed correlation is statistically significant and unlikely to have occurred by chance.

However, it is important to note that a strong correlation does not necessarily imply good agreement between the methods, and the significantly larger total number of cells counted across the entire Image Set by the CNN suggests the model is not performing well.

Therefore, the evaluation also makes use of Bland-Altman analysis (Martin Bland and Altman, 1986) to evaluate the agreement between cell counting using the CNN and manual counting methods. Bland-Altman analysis is a widely accepted and powerful statistical technique for comparing two measurement techniques (Doğan, 2018), and involves plotting the differences between paired measurements from the two methods against the average of those paired measurements. It is therefore possible to examine the presence of any systematic bias, proportional errors, or other patterns in the data that may suggest disagreement or limitations in the methods being compared. The resulting Bland-Altman Analysis plot for comparing the CNN and manual cell counting techniques across Image Set 1 is presented in Figure 83.




Figure 83 - Bland-Altman plot comparing cell counts from the machine learning (CNN) method and manual counting of cells within Image Set 1. The mean difference is -2.34, indicating a lower mean count by the CNN method. The upper and lower limits of agreement are 0.56 and -5.2, respectively, illustrating the range within which 95% of the differences between the two methods are expected to lie.

When interpreting a Bland-Altman plot, the *mean difference* represents the average of the differences between the paired measurements from the two methods, providing an estimate of the systematic bias between the techniques. The limits of agreement, calculated as the mean difference ± 1.96 times the standard deviation of the differences, indicate an interval within which 95% of the differences between the paired measurements are expected to lie, assuming a normal distribution of differences.

The Bland-Altman plot presented in Figure 83 demonstrates a mean difference of -2.34, which suggests that, on average, the CNN method counts 2.34 cells fewer than the manual method per image. This negative value indicates a systematic bias, where the CNN method tends to provide lower cell counts compared to the manual counting method, and is significant compared to the total number of cells in each image (ranging from 7 to 20 counted manually). The limits of agreement, with an upper limit of 0.56 and a lower limit of -5.2, define the range within which 95% of the differences between the paired measurements from the two methods are expected to

Chapter 4

lie. The implication of this is that for a given image, the CNN method's cell count can be anywhere from 0.56 cells more to 5.2 cells fewer than the manual count in 95% of the cases, assuming a normal distribution of differences.

To further validate the above finding, a further test on 50 more images, this time using cells from a different culture of the same species of phytoplankton (hence referred to as *Image Set 2*), was carried out. The results of this comparison are presented in Table 7. Once again, the difference between the number of cells counted manually (189) and by the CNN (348) was large (difference = 159; ~84.1% of manual counts), providing further evidence of the CNN's inability to count cells in the IFC images with a statistically similar performance to a human.

Table 7 – Table of descriptive statistics of the cell counts from 50 additional IFC images of *R. salina* cells, with cells counted either manually or by the CNN presented within this section.

Image Set 2	Manual	CNN
N=50		
Sum of all cells across all images:	189	348
Mean cells per image	3.78	6.96
Minimum cells per image	1	2
Maximum cells per image	9	13

The same statistical tests as used previously were carried out to compare the CNN method's performance with that of a manual counter. The Pearson's Correlation Coefficient statistic was computed to 0.794 (3 s.f.), with a p-value of 5.96e-12 (3 s.f.) this again shows a linear correlation between the manual and CNN-based cell counts, with a high degree of probability.

The Bland-Altman Analysis of the two techniques used in the images from Image Set 2 is presented in Figure 84. The mean difference (+3.18) between the CNN and manual methods, and the significantly higher total cell count across the entire image set, demonstrates that the CNN systematically detected many more cells per image within Image Set 2 than a manual counter.

The discrepancy between the CNN model's performance on Image Set 1, where it significantly under-counted cells, and on Image Set 2, where it significantly over-counted cells compared to the manual counter, raises concerns about the model's consistency, generalisability, and accuracy across different datasets.

These contrasting results suggest that the CNN model has not been adequately trained to handle the variability in cell appearance, image quality, or other factors that may be present in different sets of images. It is possible that the model has learned specific features or patterns from the training data that do not generalise well to other datasets, leading to inconsistent performance.

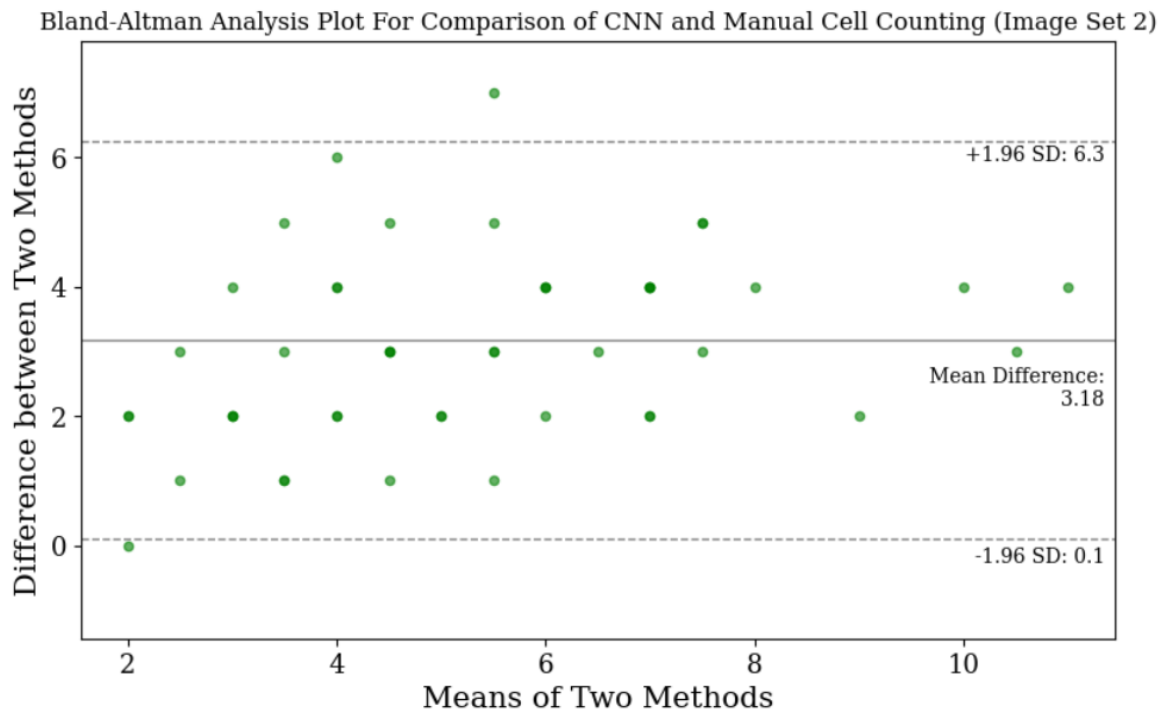


Figure 84 - Bland-Altman plot comparing cell counts from the machine learning (CNN) method and manual counting of cells within Image Set 2. The mean difference is 3.18, indicating a significantly higher mean count by the CNN method. The upper and lower limits of agreement are 6.3 and 0.1, respectively, illustrating the range within which 95% of the differences between the two methods are expected to lie.

Chapter 4

While trying to understand the discrepancy between the results of the algorithm on the first and second sets of manually counted IFC images, the tuning of the Masking Threshold value was identified as a significant source of variability. As discussed previously, the masking threshold is a manually set value which must be subjectively tuned to eliminate false positives while avoiding false negatives. In order to demonstrate this concept, an input IFC image after resizing and normalisation is presented in Figure 85, along with the output of the generated probability density maps after masking using two different masking threshold values ($M=0.01$ and $M=0.05$).




Figure 85 – Input image after resizing (top) and masked CNN output probability density maps. The different masking threshold values selected (0.01 and 0.05) produce outputs with significantly different numbers of cell detections, with a threshold of 0.01 resulting in many false positives.

Chapter 4

In the case of the image presented in Figure 85, a masking threshold of 0.05 results in a threshold mask with the same number of cells as in the input image as determined by manual count, indicating that in this instance the algorithm can successfully count the number of cells with the same accuracy as a human. The lower threshold value of 0.01, however leads to many false positive detections, particularly in the area with maximum background complexity. This is due to the neural network having some confidence of the presence of plankton cells there despite the visual artefacts actually resulting from cells adhered to the sides of the flow cells, rather than those in flow.

The appropriate masking threshold should, when determined, work regardless of the brightness, contrast and background state of the image input to the network, as the probability density map should have a similar range of confidence values for positive detections (and for background noise) regardless of the exact characteristics of the input image used. The fact that the same masking threshold value works significantly better for the first set of 50 testing images than for the second set is convincing evidence that the model is not able to predict the locations of cells with a consistent confidence across input images.

To further investigate this issue, the confidence of the model was assessed by analysing the means and standard deviation of pixel values within the probability density map, in regions of positive detections (cell presence) and background (no cell) using the following procedure:

1. Manually draw bounding boxes around the cells in several images from each set of 50 test IFC images.
2. Run the images through the object detection network to generate corresponding predicted probability density maps.
3. Extract the pixels from the density map which fall within the bounding boxes drawn in step 1 (with the coordinates of the bounding boxes appropriately scaled for the reduced image dimensions of density map compared with original image). Calculate the mean and standard deviation of these pixel values.
4. Set these pixels to 0, so that their values will not be included within the corresponding calculation for background pixel values.
5. Calculate the mean and standard deviation of all the pixels left within the image.

Chapter 4

In this manner, the value of positive detections and background noise can be analysed. The procedure was carried out on 3 images from the first unseen set of 50 images, and 3 images from the second.

The results are presented in Table 8. For each image set, the mean pixel values of pixels in the predicted density maps that are in the location of cells is approximately 10 times higher than the mean value of background pixels, which is good and demonstrates that the model can distinguish between cells and background pixels with at least some degree of confidence.

Table 8 – Pixel values in the probability density map for each of the two sets of 50 images, split by whether the pixels are contained within a manually labelled bounding box for a cell or are in the background. Confidence values for cell detections should be much higher than those in the background, which they are, at approximately 10 times higher. However, if the model was very effective, confidence values would be expected to be far higher than presented, at around 5%. Furthermore, the mean confidence of detections in Image set 1 are different than those in image set 2, which indicates inconsistent performance of the model across different input image conditions.

	Probability Density Map (Image Set 1)		Probability Density Map (Image Set 2)	
	Cells	Background	Cells	Background
Mean Pixel Value	0.0537	0.0061	0.0402	0.0048
Pixel Value SD	0.1264	0.0308	0.0939	0.0254

The values, however, indicate why a masking threshold value that works well for one image set may not be appropriate for another; the pixel values of cells within the first image set are significantly higher than those of the second image set. If the model had correctly learned a general mapping between input and the appropriate outputs, the confidence of pixels within cell regions should be comparable across each Image Set on which it is tested.

To confirm the discrepancy between confidence values on CNN outputs from Image Sets 1 and 2, Student's T-Tests were carried out to compare the means of the pixel values between the first and second image sets, for regions with cells and background regions respectively. The null hypotheses H_0 are that the pixel values for the same regions (cell or background) are the same between the two image sets. For the background pixels, the T-test generated an effect size of **14.6** and p-value of **3.00×10^{-49}** ($\ll 1\%$). For the pixels within cells, the T-test generated an effect size of **4.69** and p-value of **2.76×10^{-6}** ($\ll 1\%$). Therefore, the alternative hypotheses, that the pixel values for both background and cells are different in each of the two image sets, must be accepted.

It has been conclusively demonstrated, therefore, that although the neural network model presented should be expected to predict the location of cells with similar (and ideally high) confidence levels across different image sets, it is inconsistent in its predictions, with this inconsistency leading to unreliable cell counts. While cell counts could in theory be made more accurate by regularly checking and altering the masking threshold value such that cells are positively detected and background pixels are not, in practice this disadvantage nullifies the sought benefits of using this particular CNN configuration to perform object detection on the IFC images. The difference in performance described here with the good performance in the initial implementation by Heo et al. is a result of the significantly increased image complexity; the use of full colour and the complexity of both the imaged cells and background are too great a challenge for this algorithm.

In order to address the deficiencies of the explored object detection approach, a second neural network-based model was investigated, which will be detailed in the following sections.

4.4.5 YOLOX for Phytoplankton Detection

YOLO (You Only Look Once) is an object detection model that was introduced in 2016 by Redmon et al. (2016). It is a *single-stage* (predictions are made without subsequent thresholding of a confidence map) object detection model based on a CNN that divides an image into a grid and predicts bounding boxes and class probabilities for each grid cell, as shown in Figure 86. The model also uses anchor boxes to improve the accuracy of the bounding box predictions. YOLO is known for its real-time object detection capabilities and is often used in self-driving cars, robotics and other applications where inference speed is important (Terven and Cordova-Esparza, 2023).

Since its introduction several iterations have improved upon the YOLO architecture. *YOLOv2* (aka *YOLO 9000*) (Redmon and Farhadi, 2017) introduced an improved CNN ‘backbone’, batch normalisation and several other improvements.

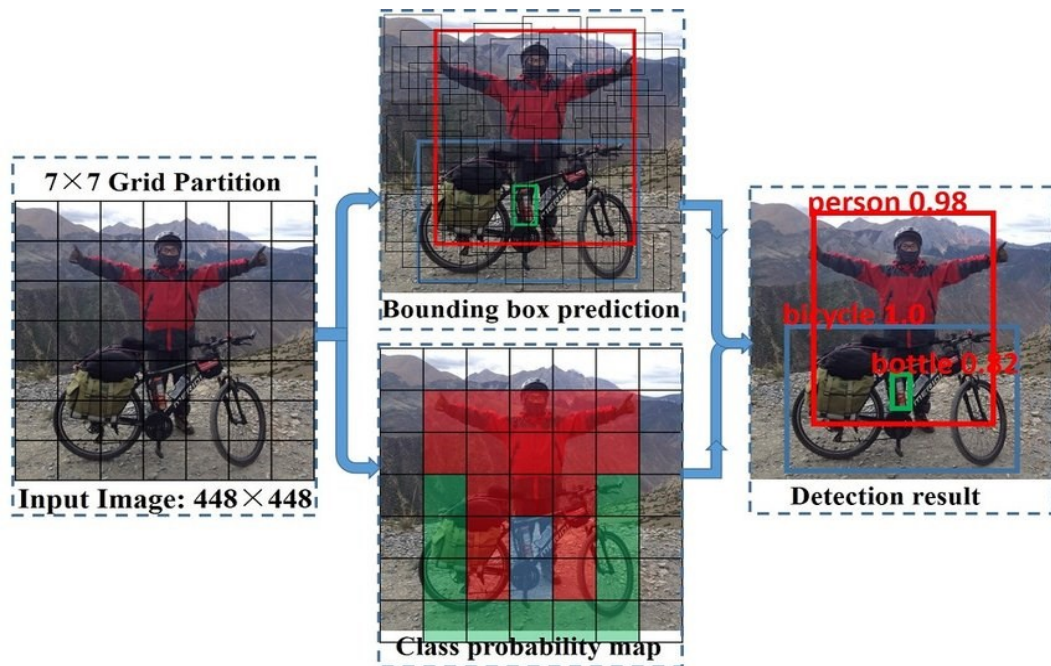


Figure 86 - Top-level overview of the behaviour of a YOLO object detection model when detecting objects within an image. The image is divided into a grid, then the network predicts bounding boxes and confidence scores for potential objects within each grid cell, along with their associated class probability maps. These predictions and class probability maps are combined to create a detection result, which identifies the locations and classes of objects within the image. (Wu and Zhou, 2019)

YOLOv3 (Redmon and Farhadi, 2018) again introduced a new backbone, *Darknet53* (Figure 87), as well as other small improvements resulting in better performance overall and particularly for detection of objects at different scales. *Darknet53* makes use of several *Residual Layers*, layers which improve the ability of information to flow through the network by introducing *skip connections* that allow information to bypass one or more layers in the network.

Chapter 4

	Type	Filters	Size	Output
1x	Convolutional	32	3 × 3	256 × 256
	Convolutional	64	3 × 3 / 2	128 × 128
	Convolutional	32	1 × 1	
	Convolutional	64	3 × 3	
	Residual			128 × 128
2x	Convolutional	128	3 × 3 / 2	64 × 64
	Convolutional	64	1 × 1	
	Convolutional	128	3 × 3	
	Residual			64 × 64
8x	Convolutional	256	3 × 3 / 2	32 × 32
	Convolutional	128	1 × 1	
	Convolutional	256	3 × 3	
	Residual			32 × 32
8x	Convolutional	512	3 × 3 / 2	16 × 16
	Convolutional	256	1 × 1	
	Convolutional	512	3 × 3	
	Residual			16 × 16
4x	Convolutional	1024	3 × 3 / 2	8 × 8
	Convolutional	512	1 × 1	
	Convolutional	1024	3 × 3	
	Residual			8 × 8
	Avgpool		Global	
	Connected		1000	
	Softmax			

Figure 87 - The Darknet53 architecture, which consists of 53 convolutional layers and is used as the backbone for YOLOv3, a modern object detection algorithm. Diagram from (Redmon and Farhadi, 2018).

YOLOX is a state-of-the-art extension of YOLO that was introduced in 2021 (Ge et al., 2021). It improves upon prior YOLO derivatives in several ways which result in superior accuracy, speed, and efficiency. *YOLOX* uses an improved version of the CNN backbone used by YOLOv3 called *CSPDarknet* (Figure 88) that has a higher efficiency and allows for more parameters to be trained. This efficiency derives from the Cross-Stage Partial connections employed in the CNN which give it its name. As in a Residual Layer, in *CSPDarknet*, each layer is connected not only to the next layer but also to others across the network in what is known as a cross-layer connection. A cross-layer connection is highlighted in the green box of Figure 88, which illustrates the complete structure of *CSPDarknet*. This partial connection strategy allows the network to share information more efficiently across layers, which can reduce the number of parameters that need to be trained and can make the network more computationally efficient.

CSPDarknet also utilizes *channel splitting*, where the channels of the input feature map are split into two groups, one of which passes through a convolutional layer while the other group is passed through a *cross-stage partial connection layer*, which concatenates the feature maps from multiple layers. This operation further reduces the number of model parameters thus improving computational efficiency, and is highlighted in the right-hand 'CSP Block' section of Figure 88.

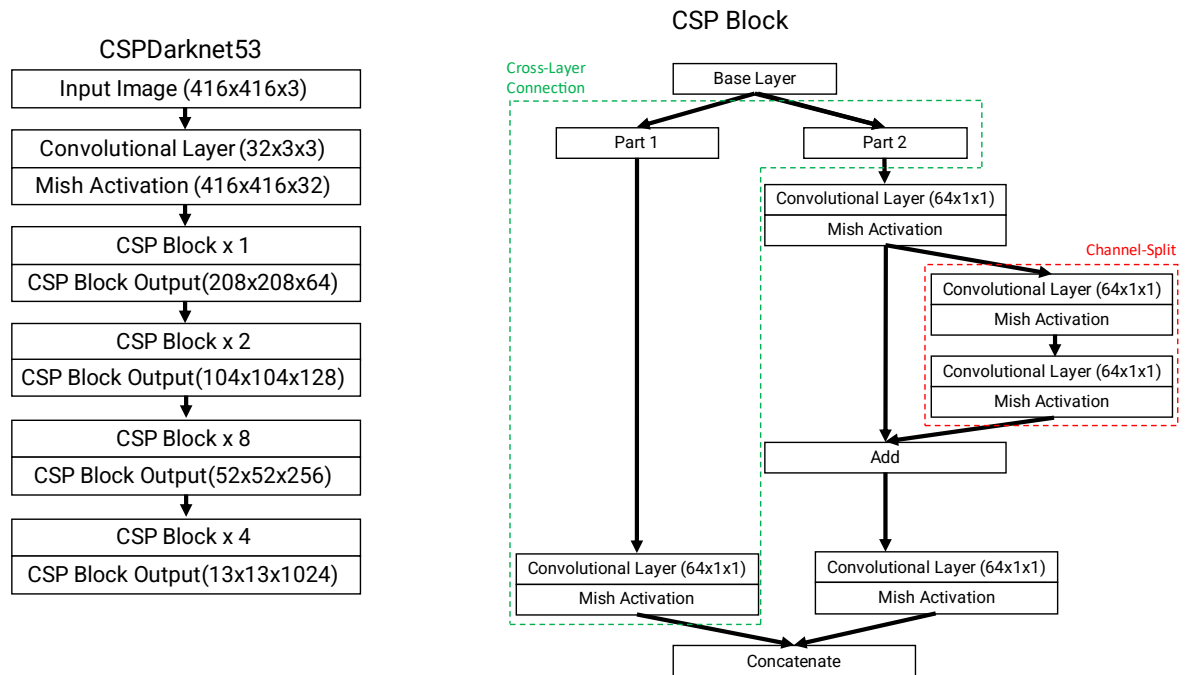


Figure 88 - The CSPDarknet53 neural network architecture, first introduced in YOLOv4. Used as the backbone for advanced YOLO object detection models, CSPDarknet53 employs Cross-Stage Partial (CSP) connections, facilitating superior feature extraction capabilities and enhanced object detection performance across a range of scales in the input image. The schematic on the right illustrates the structure of a CSP block, which divides the input features into two streams: one stream passes through a sequence of convolutional layers, while the other bypasses these layers. Subsequently, both streams are merged, resulting in an efficient and effective combination of extracted features. Adapted from (Xu et al., 2021).

In addition to the Backbone previously discussed, YOLOX comprises a *Neck* consisting of the *Path Aggregation Network* (PANet) introduced by Liu et al. (2018). The purpose of the neck is to concatenate the feature maps from different layers of the backbone network and send them as inputs to the *Prediction Head*, which takes the feature maps from the Neck and outputs bounding boxes and class predictions. YOLOX introduces a novel prediction head module which separates the classification and bounding box localisation into two separate tasks, which the authors found to improve accuracy compared to previous YOLO versions which performed coupled predictions.

YOLOX also introduces a *Spatial Pyramid Pooling* (SPP) module to capture features at multiple scales by applying pooling operations of different sizes to the backbone-predicted feature map, producing a new feature map that encodes information about objects at different scales. Using

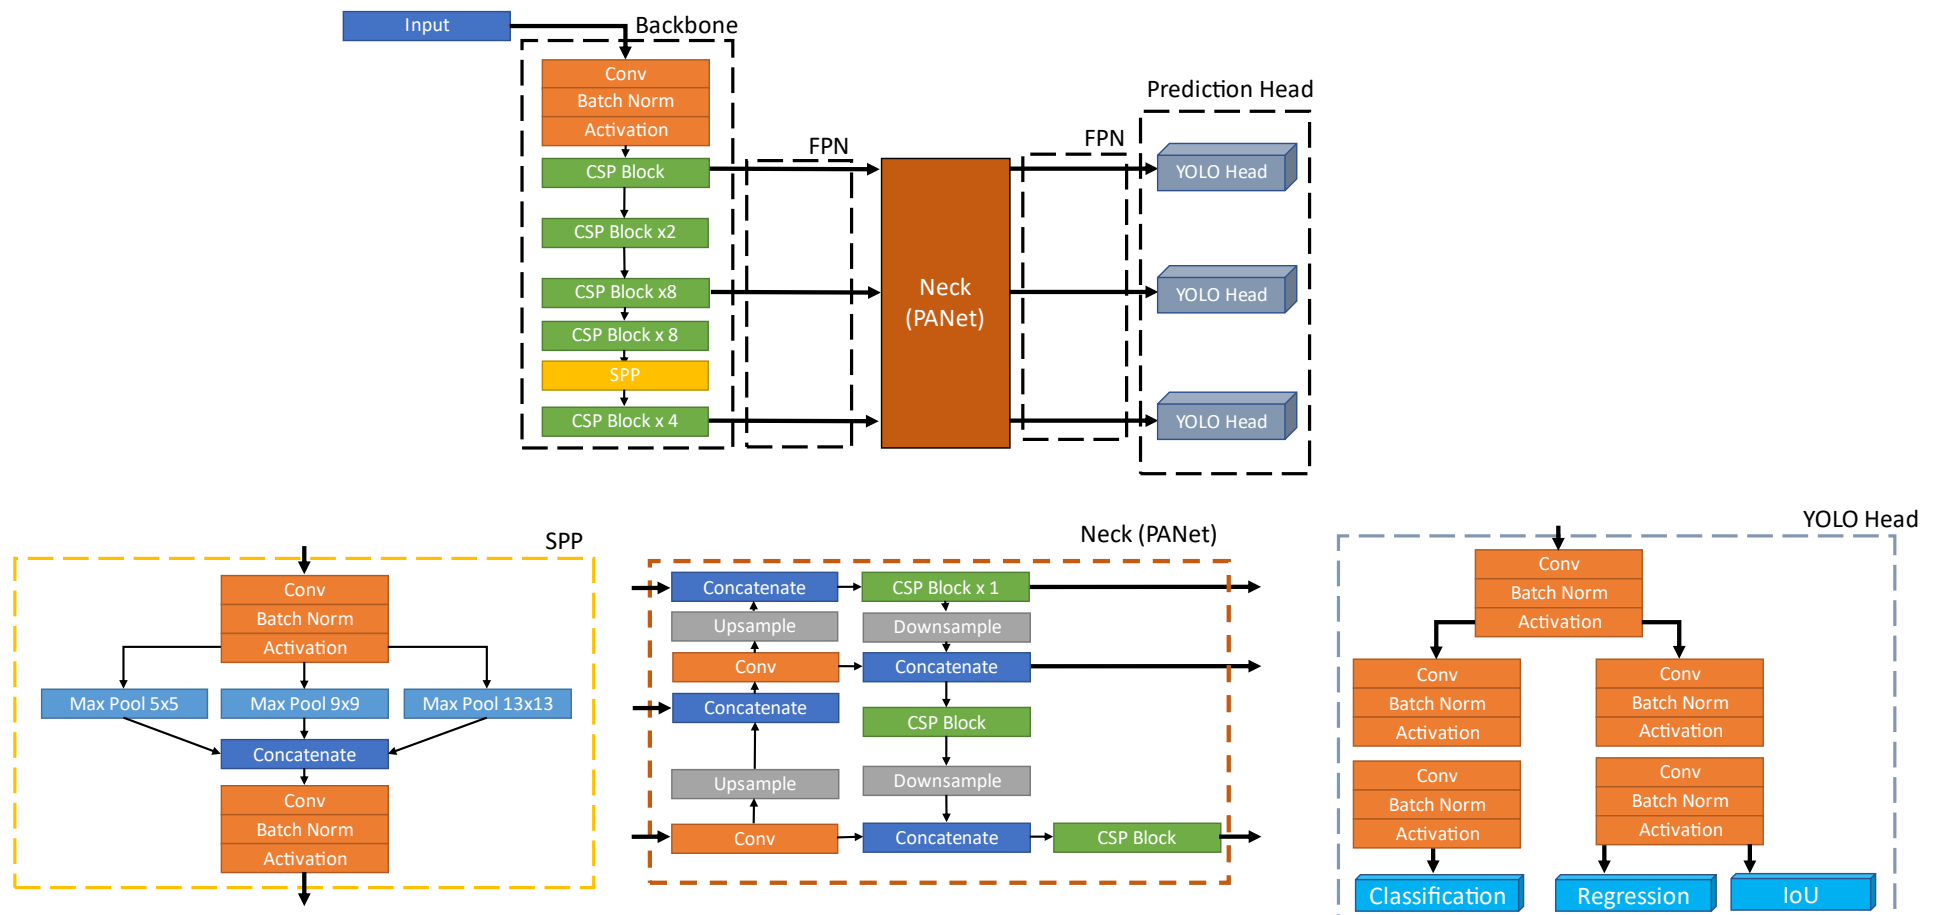
Chapter 4

Feature Pyramid Networks (FPN) adds allows the model to combine high-level features from the upper layers of the backbone or neck with lower-level features from the preceding layers. This produces a set of feature maps with different resolutions that are used to detect objects at different scales.

In combination, the SPP and FPN connections allow YOLOX to detect objects of different sizes and scales in an image more accurately and efficiently. In the context of IFC imagery of phytoplankton cells, these features combine to make YOLOX well suited to the task of detecting multiple plankton species of widely differing sizes, shapes and aspect ratios. YOLOX was also recently demonstrated to be capable of performing high-accuracy identification of plankton within fluorescence imaging (Wang et al., 2022), but has not yet been applied to brightfield, colour phytoplankton images such as those from the IFC developed here and described in Section 3.

Figure 89 illustrates the entire YOLOX architecture including the Backbone, Neck and Prediction Head. Compared to the previous Object Detection network investigated, YOLOX is a far more advanced model which is much better suited to the complexity of IFC imagery, as will be demonstrated in the testing which follows.

Figure 89 - The complete YOLOX network architecture, featuring inlaid diagrams of its core components: the Spatial Pyramid Pooling (SPP) module, the Neck (Path Aggregation Network, or PANet), and the YOLO Prediction Head. These components enable efficient feature extraction, hierarchical feature aggregation, and accurate bounding box predictions with class probabilities, ultimately resulting in cutting-edge object detection for objects of a wide range of scales. Adapted from (Chou, 2022)



YOLOX also has the useful attribute of having been developed in several different levels of parameter complexity; different versions of the model called *YOLOX-l* (large), *YOLOX-s* (small), *YOLOX-nano*, etc., have been published with a range of different numbers of parameters. The different versions of the model have a different number of layers in the neck and backbone, with the number of CSP blocks and complexity of the PANet neck dictated by version. The lighter-weight models, i.e., *YOLOX-s* and *YOLOX-nano* are small enough to be run on embedded platforms such as those discussed in section 4.4.2. Nguyen et al. (2022) conducted a comparison of different YOLO versions on embedded platforms (NVIDIA Jetson Nano and Jetson Xavier AGX) and determined maximum framerates (including image I/O and output processing). A table of model parameter counts and the results of their investigation is presented in Table 9 below. Compared to the previous model with ~52,000 parameters, even these small models are extremely complex. However, due to the parallel computing afforded by GPU-based running of the networks, the maximum framerate of the models is sufficient to allow real-time (>30 FPS) object detection on the Jetson Xavier embedded platform.

Table 9 – Comparison of the parameter complexity of the YOLOX-s and YOLOX-nano object detection models and the framerates at which they can process 512x512 images. Adapted from Nguyen et al. (2022)

Model	Trainable Parameters (millions)	NVIDIA Jetson Nano Framerate	NVIDIA Jetson Xavier AGX Framerate
YOLOX-s	9.0	8	32
YOLOX-nano	0.91	13	40

4.4.6 YOLOX Training

In order to train the YOLOX model, the same 150 images used for the previous neural network training were again curated into training and validation datasets. Similarly to the previous CNN trained, training YOLOX was carried out with images resized to 512x512 pixels, but due to the complexity of the model, a maximum batch size of 8 was possible using the same PC hardware as detailed previously.

Chapter 4

Unlike the previous model, however, rather than the point coordinates used to represent cells, bounding boxes were drawn around the entire cells; though the same image labelling software (*labelme*) was used. As discussed in section 4.3, object detection models generating bounding boxes around detected objects often use the *Average Precision* performance metrics *AP_{XX}*, where *XX* is the *IoU threshold percentage* at which precision is calculated. *AP₅₀* and *AP₉₅* are two of the most commonly reported AP metrics within the literature (e.g. (Liang et al., 2023, Xue et al., 2021, He et al., 2021)), so these were selected for reporting the performance of YOLOX on the IFC phytoplankton images.

The YOLOX-s model with 9 million parameters was selected as the version of YOLOX to use for the problem due to its aforementioned ability to be used on embedded platforms. The model was trained for the same number of epochs (150) as the previous model investigated, and the *AP₅₀* and *AP₉₅* on the validation datasets were recorded at each epoch. Graphs of *AP₅₀* and *AP₉₅* against epoch are presented in Figure 90. Similarly to the previous model, it is evident that there is a rapid increase in performance during the first 20 epochs as the model learns to recognise the important features of cells within the images, with noisy asymptotic behaviour continuing between 20 and 150 epochs. Unlike the previous model, performance of YOLOX-s continues to be noisy throughout the 150 epochs, which suggests that overfitting has not yet occurred during this training timeframe. This is further backed up by the fact that the *AP₅₀* values seem to still be increasing, on average, towards the end of training.




Figure 90 – Graph of AP50 (Blue) and AP95 (Orange) for YOLOX-s object detection model while being trained to detect and localise cells of *Rhodomonas salina* phytoplankton within images from the acoustically focussed IFC. After 150 epochs of training the AP50 value is close to 1.0, indicating a good overlap between predicted and ground truth bounding boxes. The lower AP95 score suggests that the model does not predict bounding boxes with a very tight alignment to those which were manually labelled to produce the training dataset, the 95% IoU threshold requires a large degree of overlap between the predicted and ground truth bounding boxes. This is unlikely to be a cause for concern as the human-labelled bounding boxes were not highly precise to a pixel level, especially in comparison to the size of the cells themselves, which each take up a very small proportion of the image at 10X magnification.

An Average Precision value of 1.0 would indicate that the model has achieved perfect object detection accuracy on objects in the validation dataset. This means that all objects in the validation dataset have been correctly detected by the model without any false positives or false negatives. In the training results of YOLOX-s presented above, AP50 values asymptotically approach 1.0 whereas AP95 values do not reach significantly above 0.4. This suggests that the model is able to detect objects with a high degree of accuracy for a lower degree of overlap with the ground truth bounding boxes, but may struggle to achieve the same level of accuracy for a higher degree of overlap. Due to the difficulty of drawing bounding boxes accurately for the small

cell sizes present within the training images, the lower AP95 score is not a major cause for concern. It is reasonable to expect a less precise overlap between the predicted bounding boxes and the manually labelled boxes when the latter are known not to be extremely precise.

4.4.7 Evaluation of YOLOX Performance

After training, the model performance was evaluated using the same challenge as for the first network investigated; 2 sets of 50 unseen (not in the training or validation datasets) IFC images had their YOLOX-s automated cell counts compared with manual counts. The descriptive statistics of this comparison on the first 50 images (Image Set 1) is presented in Table 10. For comparison, the results derived from the first CNN investigated in section 4.4.4 are included again.

Table 10 - Table of descriptive statistics of the cell counts from 50 IFC images of *R. salina* cells, with cells counted either manually or by the YOLOX-s object detection network presented within this section. For the purposes of comparison, the results of the same analysis using the CNN investigated in section 4.4.4 are also included.

<i>Image Set 1</i>	Manual	YOLOX	<i>First CNN</i>
<i>N=50</i>			
Sum of all cells across all images:	593	593	476
Mean cells per image	11.86	11.86	9.52
Minimum cells per image	7	7	6
Maximum cells per image	20	22	17

A Pearson's Correlation Coefficient test was carried out and the computed test statistic (r) was 0.981 (3 s.f.), with a p-value of 6.70e-36 (3 s.f.). This test statistic is far closer to the ideal value of 1.0 than achieved by the previous CNN, which was 0.881 (3 s.f.), with a p-value of 3.13e-17 (3 s.f.).

As in the investigation of the previous CNN (Section 4.4.4), a Bland-Altman analysis was conducted on the per-image cell counts from Image Set 1 to compare the YOLOX and manual counting strategies. Figure 91 shows the plot of this analysis. With a mean difference of 0.0 and

95% confidence intervals of ± 1.3 , the model performance is clearly vastly superior to that of the simple CNN, and compares very well to the manual cell counts with no systematic bias.




Figure 91 - Bland-Altman plot comparing cell counts from YOLOX and manual counting of cells within Image Set 1. The mean difference is 0.0, indicating a lack of systematic bias. The upper and lower limits of agreement are +1.3 and -1.3, respectively, illustrating the range within which 95% of the differences between the two methods are expected to lie. These values further demonstrate the absence of systematic bias.

In order to confirm the above finding, the same second set of 50 images as used for the first model was fed through the trained YOLOX-s model and the cell count recorded. The results of this comparison with manual cell counts is presented in Table 11, and the same methodology as above was used to compare the two sets of counts.

Table 11 - Table of descriptive statistics of the cell counts from the second set of 50 IFC images of *R. salina* cells, with cells counted either manually or by the YOLOX-s object detection network presented within this section. For the purposes of comparison, the results of the same analysis using the CNN investigated in section 4.4.4 are also included.

Image Set 2 <i>N=50</i>	Manual	YOLOX	<i>First CNN</i>
Sum of all cells across all images:	189	198	348
Mean cells per image	3.78	3.96	6.96
Minimum cells per image	1	1	2
Maximum cells per image	9	12	13

For images within the second Image Set, the Pearson's Correlation Coefficient between manual and YOLOX-counted cells was computed with a statistic 0.962 and a p-value of 1,22e-28. This statistic indicates that the YOLOX performed slightly worse on the second Image Set compared to the previous, a finding which is further backed up by the total cell count across all images being 4.8% higher when counted by YOLOX compared to a human.




Figure 92 - Bland-Altman plot comparing cell counts from YOLOX and manual counting of phytoplankton cells within Image Set 2. The mean difference is 0.18, indicating a slight systematic bias where YOLOX detects more cells per image. The upper and lower limits of agreement are +1.5 and -1.1, respectively, illustrating the range within which 95% of the differences between the two methods are expected to lie.

A Bland-Altman Analysis was again carried out to compare the cells counted in each image by a human and by the YOLOX model, and the plot is presented in Figure 92. Unlike previously, there is now a small systemic bias, with a mean difference of 0.18 cells per image. The 95% confidence intervals of 1.5 and -1.1 are skewed in the direction of over-counting, and bar the outliers at low mean cell counts, the plotted points show that YOLOX tended to predict more cells per image than a human across the entire set of images.

In order to determine the source of this discrepancy, which was not presented within the previous test, a qualitative analysis of the input images within Image Set 2 was carried out. Image Set 2 was found to have significantly increased detritus within the images compared to the previous set, and while these clumps of debris were not counted as cells by the human, several of them were mistakenly identified as plankton by the algorithm. An example of one of these misclassifications is presented in Figure 93, which shows three correctly detected plankton cells and one clump of detritus which has been classified as a cell (with 63.5% confidence).




Figure 93 – An illustrative example of a misclassification of detritus as a plankton cell by YOLOX. While the three cells present in the image are correctly identified with >70% confidence, the clump of detritus has been falsely labelled as a cell by the algorithm with 63.5% confidence.

Despite this, the very close degree of alignment between the cell counts generated by YOLOX and by human indicate that the model has correctly learned a general mapping between input and output data, and is reliable enough to be used for the automatic detection, localisation and counting of plankton cells within IFC images. The amount of cellular detritus within the samples being processed should, however, be closely monitored to ensure that the algorithm does not incorrectly label debris as cells.

Based on the above results of the two comparisons between human and YOLOX cell counting, it can be claimed that YOLOX-s can be used to count cells to a statistically similar performance to that of a human analyst, unlike the prior neural network investigated in Section 4.4.4.

In addition to the ability of the network to correctly count cells within the IFC images, YOLOX-s has the distinct advantage of simultaneously classifying and predicting the bounding box of the

Chapter 4

detected objects, as discussed previously. In this experiment only cells of a single species (a monoculture) were counted, so the model only distinguishes between cells and the background. The model also performs classification on any objects it detects, however, so it could be easily adapted to identify different types of cell, provided enough labelled examples were made available during training. Additionally, though this further work was not within the scope of this thesis, it should be possible to improve the performance of the algorithm on monoculture samples by labelling cell detritus as a second class independent from the labelled cells. This would enable the model to learn to distinguish between cells, debris and background pixels.

4.5 Conclusion

To address the need to provide fast and accurate automated cell detection in images generated by the IFC investigated within this thesis, traditional digital image processing using handcrafted features has been shown to be inadequate. The complexities, which include cells adhered to the sides of the flow cell, different cell geometries and cells in contact with each other, pose too much of a challenge to image processing based on brightness, contrast and other features which would be sufficient on other IFC systems which produce less visually complex imagery.

Machine learning offers the potential to algorithmically determine the set of features which optimally distinguish the cells within an image from the flow cell background, and to be far more robust than traditional techniques. Convolutional Neural Networks have been demonstrated to be very well-suited to the problem of object detection within images and were shown by Heo et al. (2017) to perform well in another IFC system. They are also well-disposed to be used on small, embedded platforms such as the NVIDIA Jetson SoMs due to their parallel computing architecture, an advantage for an IFC platform which may need to be deployed away from existing computational infrastructure. A model following the implementation of Heo et al. (2017), but with improvements based on more recent developments in Computer Vision research, was created and trained on 150 pairs of images and labelled cell coordinates. The model unfortunately did not perform adequately upon testing on images on which it had not been trained, as there was a statistically significant difference between the cell counts generated automatically by the algorithm and those manually counted.

In order to create a more robust and accurate automated cell detection system, the cutting-edge object detection CNN *YOLOX-s* was trained on the same 150 images. *YOLOX-s* demonstrated a far greater ability to distinguish cells from the flow cell background, including successfully detecting cells in flow even partially obscured by adhered cells. In comparison with the manually counted numbers of cells within IFC images, *YOLOX-s* did not have a statistically significant difference with the manual counts.

The utility of *YOLOX-s* at solving the problem of detecting phytoplankton in IFC images even when the flow cell background is highly complex has therefore been demonstrated. This is the first implementation of *YOLOX-s* being used for phytoplankton detection in colour IFC images, and the

Chapter 4

first time a CNN has been used for localising objects within images from an acoustically-focussed IFC. This approach has the potential to drastically improve upon existing automated cell analysis of IFC by existing instruments such as the *FlowCam* and IFCB. YOLOX-s has been demonstrated to run at over 30 frames per second on the NVIDIA Jetson Xavier AGX platform, a low-cost (~£1000) computing system designed for AI applications, opening the possibility of the IFC being deployed in a fully automated sample-and-analyse configuration without any need for human input.

Chapter 5 Use of the Acoustically-Focussed IFC in Phytoplankton Experiments

5.1 Introduction

This chapter presents the testing and verification of the developed autonomous, acoustically-focussed imaging flow cytometer in two experiments with phytoplankton cells, highlighting its potential as a low-cost, high-throughput imaging method to address diverse research questions within phytoplankton ecology and aquaculture.

Section 5.2 focuses on a using the IFC to automatically count cells with the image processing techniques presented in Chapter 4. Preserved phytoplankton cells of a mono-culture, a simplified analogue to samples those collected and preserved at sea during oceanographical research cruises, were analysed and results compared to the gold standard of manual counting by microscopy. This section aims to evaluate the performance, efficiency, and accuracy of the complete IFC system including automatic cell identification in quantifying cell concentrations. The outcomes of this comparison provide essential information regarding the IFC's potential as a viable alternative for traditional microscopy of phytoplankton cells, particularly in settings where high throughput and autonomous operation are desirable.

In section 5.3, the IFC is used in a long-term experiment with live plankton cells, monitoring the growth of phytoplankton cultures every few hours from their initial seeding to the end of the growth phase after more than a week. By measuring cell count at regular intervals, the IFC enables the tracking of the dynamics of the phytoplankton population over time and at a high temporal resolution, capturing critical information about growth rates and response to diurnal lighting conditions. This experiment showcases the IFC's capacity to provide continuous, in-depth, real-time data on phytoplankton populations. Continuous imaging over a long period is critical as it offers the capability to autonomously monitor microalgal populations for applications in ecological studies, phytoplankton biotechnology, aquaculture management, HAB early detection and environmental monitoring. This continuous application is not possible using most other IFC available commercially or released in open-source papers due to their discrete sampling regimes and requirement for supervision, so represents a major advantage of the system laid out within this thesis.

Chapter 5

Together, these experiments demonstrate the versatility and potential of the acoustically-focused IFC in addressing a range of research questions related to phytoplankton ecology and biology. By exploring the IFC's applications in both preserved and live cell experiments, this chapter highlights the instrument's capacity to contribute to a deeper understanding of phytoplankton communities and their role in aquatic ecosystems. It is demonstrated that the IFC presented in this thesis has the potential to provide a novel, unique and low-cost approach to high-throughput autonomous analysis of algal populations in a manner that would previously be inaccessible to all but the small number of oceanographic laboratories with the funding for commercial instrumentation.

5.2 Experimental Validation of Complete System

5.2.1 Experimental Objectives

For the reliable use of the presented IFC as a replacement for a human microscopist, it is critical that the entire system, including image analysis algorithms, acoustic focussing and imaging hardware, can autonomously count cells within a sample with the same level of accuracy as a human. This means that not only does the counting of cells within images need to be accurate, the number of cells which are imaged prior to counting needs to be a known and consistent proportion of all the cells present in a given volume of liquid. Only if this is this case can one confidently extrapolate the overall population of cells within a given volume from an IFC-analysed sample.

Therefore, in order to validate the cell-counting performance of the presented image analysis system and the acoustically focussed IFC hardware as a complete system, a cell counting experiment was devised. A preserved sample of phytoplankton is injected into the IFC, imaged, and then collected after passing through the complete system. This collected sample will then be analysed under light microscopy, with the number of cells counted. The captured images will be analysed using the YOLOX algorithm presented in Chapter 4, and the cell count from each method compared.

5.2.2 Materials and Methods

Cells of *Rhodomonas salina*, a motile cryptophyte originating from a brackish, eutrophic, and lower latitude water body, were obtained from the *Culture Collection of Algae and Protozoa* (CCAP), UK. *R. salina* is a flagellated cryptophyte phytoplankton with a typical cell equivalent spherical diameter (ESD) of approximately 10 μm . *R. salina* was selected as it is a robust, easily-cultured and well-studied phytoplankton species and is widely used as an aquaculture feedstock (Thoisen et al., 2018).

A 5mL sample of *R. salina* was preserved with Lugol's Iodine 2%, agitated, and introduced to the IFC using a syringe pump at a volumetric rate of 0.65 mL/min.

Chapter 5

2500 image frames, for a total imaged volume of ≈ 0.769 mL (3s.f.) were captured using the IFC, and the sample collected from the flow cell outlet. After imaging, 1mL of the collected sample was pipetted onto a *Sedgewick Rafter Cell* and given 10 minutes to fully settle. Subsequently, the cells in 20 μ L of this sample were counted manually under a brightfield microscope, at 10X magnification. 20 μ L was selected for the microscope analysis as it is the volume commonly analysed in the literature when cell abundance is high (Menden-Deuer et al., 2020, Gutiérrez-Rodríguez et al., 2016).

5.2.3 Results and Discussion

Each individual “count”, which for the IFC represents the cells in 0.30751 μ L, and for the manual microscopy represents the cells in a single 1 μ L well of the *Sedgewick Rafter Cell*, - was scaled up to cell count per mL. The mean cell count of the automated IFC approach was **222,000 cells/mL** (3 s.f.), standard deviation 25,300 (3 s.f.). The mean of the manual cell counts was 228,000 cells/mL (3 s.f.), standard deviation 21,500 (3 s.f.). The 95% confidence intervals for the means of each method were 219,000– 238,000 (3 s.f.) for microscopy and 221,000 – 223,000 (3 s.f.) for IFC. A box and whisker graph was plotted to show the two sets of counts, and is presented in Figure 94.




Figure 94 - Box and whisker plot of the cell counts generated by: IFC & YOLOX (left/blue), the fully automated system involving imaging using acoustically focused IFC and processing using the YOLOX object detection model; and manual cell counts (right/orange) using a microscope. The deviations from the mean do not indicate inaccuracies in the counting method, but rather highlight the non-uniform distribution of cells in each imaged volume.

The graph in Figure 94 shows that the interquartile range for the automated cell counts was significantly larger than that of the manual cell counts, and there are more outliers in the IFC technique. As samples are not completely homogeneous even after agitation, it is unsurprising that some images contain comparatively more or fewer cells than would be expected if every image was an equal, uniform sample of the overall population. Comparing the 2500 individual cell counts generated automatically with the 20 manual counts is therefore not straightforward and inherently requires a trade-off between analysis speed and ease versus the potentially higher sample variance and bias of smaller volume samples. It is promising that the entire 95% confidence interval of the mean of IFC counts is within the same confidence interval of the microscopy counts, and serves as good evidence that with a greater number of manual counts the two methods would produce identical means.

A statistical test was used to find the quantitative likelihood of the two cell counting methods being equivalent. Due to the significantly different number of counts generated by the manual and automated techniques (20 vs 2,500), the variances between the two groups might be unequal

and a simple 2-sample T-test would be inappropriate. Therefore, the *Mann-Whitney U test* (Mann and Whitney, 1947), a robust non-parametric statistical test that does not assume equal variances between the two groups, was used, with a null hypothesis, H_0 , that the two methods are samples of the same underlying population.

The test statistic was 21805.0, with a p-value of **0.324** (3s.f.) (>0.05), therefore the null hypothesis is not rejected. It follows that the counts from the two techniques are statistically equivalent, and it has been demonstrated that the acoustically-focussed IFC in conjunction with YOLOX can automatically count phytoplankton cells in a sample with an equivalent accuracy to a human using a microscope.

This experiment demonstrates the potential of the novel IFC developed within this thesis to be used to automatically analyse fixed phytoplankton samples. This is particularly important as much phytoplankton research, such as that conducted during research cruises, involves capturing cells, preserving them, and analysing them at a later stage (Santhanam et al., 2019). The novel IFC developed in this thesis offers several key advantages over traditional methods:

Firstly, the IFC enables rapid processing and analysis of numerous samples, significantly improving the throughput of phytoplankton studies compared to manual microscopy or existing IFC instrumentation. Additionally, by automating the analysis process, the IFC reduces reliance on human expertise and subjectivity, enabling more consistent results across different samples and studies.

In conclusion, the successful application of the acoustically-focussed IFC for automatically analysing preserved phytoplankton samples has significant implications for aquatic ecosystem research, biotechnology and HAB monitoring. By offering a cost-effective, high-throughput and automated alternative to traditional methods and commercial IFCs, the acoustically-focussed IFC has potential to advance phytoplankton research.

5.3 Use of Acoustically-Focussed IFC in a Live Phytoplankton Experiment

5.3.1 Introduction

In order to demonstrate the utility of the acoustically-focussed IFC and image analysis system at counting not only 'fixed' (preserved) populations of dead phytoplankton, but also to analyse in real-time the population of live algae as cells move and reproduce over time, a second experiment was devised and will be reported in the following sections. The use of the IFC for this purpose is important as it demonstrates a capability that is currently not possible using many traditional instruments such as FlowCam or traditional microscopy, both of which require the analysis of discrete samples of plankton cultures, rather than being directly connected to and able to sample from the cultures as they grow.

Therefore, an experiment was carried out with the aim of monitoring the cell count of growing cultures of phytoplankton over time, with the aim of demonstrating the IFC's ability to provide a time-lapse of the fluctuations of cells due to growth and diurnal effects. Culturing vessels of growth medium were seeded with populations of live phytoplankton and connected to the IFC via a computer-controlled fluidic system in order to allow the IFC to draw and image samples from the cultures, count the cells within those images using the previously presented algorithms and extrapolate those counts to determine the total cell population within the cultures. Additionally, manual samples were drawn once per day to verify the cell counts from IFC.

Phytoplankton Growth Rate

Phytoplankton are a critical component of aquatic ecosystems and play a vital role in global biogeochemical cycles. As primary producers, they fix carbon dioxide and produce organic matter through photosynthesis, supporting higher trophic levels and driving the biological carbon cycle. Determining the growth rates of individual species is essential for predicting the responses of marine and freshwater ecosystems to environmental change.

Chapter 5

Growth rate is the rate at which phytoplankton reproduce over a given period. The specific growth rate exhibited by individual species is influenced by a variety of factors, primarily nutrient availability, light intensity, temperature, and grazing pressure.

Measuring the growth rate of individual phytoplankton species is essential for understanding their responses to environmental changes, such as pollution, ocean acidification and climate change. For example, as discussed in Chapter 2, some species are more sensitive to changes in nutrient availability than others, leading to shifts in community composition and ecosystem function in response to eutrophication and other anthropogenic changes to nutrient content.

To measure the growth rate of phytoplankton, researchers often use culture-based methods. These involve isolating individual species of phytoplankton and culturing them in laboratory conditions, under controlled environmental conditions. The growth of a culture can be monitored over time by measuring changes in cell numbers or biomass, using techniques such as cell counting or optical density (OD) measurements.

The growth dynamics of phytoplankton are complex and involve distinct phases. Generally, phytoplankton growth can be divided into four stages, which are the *lag phase*, *exponential growth phase*, *stationary phase*, and *death phase*. These phases are briefly summarized below.

In the lag phase, the population is adjusting to environmental conditions, and growth is either slow or non-existent. Cells synthesize new proteins and enzymes to optimize resource utilization, such as nutrients and light, and prepare for growth. The duration of this stage may vary from a few hours to several days, depending on species and environmental conditions (Vonshak, 1985).

Once adapted to the environment, the phytoplankton population enters the exponential growth phase. In this stage, the rate of cell division increases rapidly, and the population size doubles at an approximately constant rate. This phase is the peak of growth, where the phytoplankton population attains its maximum abundance and biomass.

Chapter 5

As the population grows, it eventually enters a stationary phase due to resource limitation, leading to a plateau in population growth. During this phase, the rate of cell division slows down, and the number of new cells produced is eventually balanced by the number of cells dying.

In the final phase, known as the death phase, the phytoplankton population declines due to various factors such as resource depletion and environmental stress. The rate of cell death exceeds the rate of cell division, leading to a decrease in population size and biomass.

A key measurement of phytoplankton productivity reported in the literature is the maximum growth rate of a given species under specific environmental conditions. This value is also called the *maximum specific growth rate* (μ_{\max}) or *intrinsic growth rate* (Sun and Ning, 2005). To determine the maximum growth rate, experiments are conducted that involve monitoring a phytoplankton monoculture over time. This process typically entails measuring the cell count or optical density (OD) of the culture at 24-hour intervals until the end of the exponential phase is reached. By analysing the rate of change in cell count during this phase, it is possible to accurately estimate the maximum growth rate, which serves as a crucial indicator of the species' potential for population expansion and overall productivity in its natural environment. This information can be invaluable for understanding the dynamics of phytoplankton communities, particularly the speed with which they can bloom (Kremer et al., 2017).

Diel Vertical Migration

Motile phytoplankton can move up and down within the water column in response to changing environmental factors such as exposure to photosynthetically available radiation, nutrients, and the presence of predators. Diel vertical migration (DVM) is a common behaviour observed in many phytoplankton species, which involves the daily movement of these organisms up and down the water column (Olli, 1999).

During the day, phytoplankton species exhibiting DVM stay near the surface of the water, where they can most efficiently absorb sunlight for photosynthesis. As light levels decrease, they sink towards deeper waters, where the nutrients such as nitrate and phosphate, which are needed for growth and reproduction, are often more abundant (Eppley et al., 1968).

A unique capability of a device which can continually monitor the cell count within a culture such as the acoustically-focussed IFC is the indirect observation and potentially quantification of DVM. As the inlet to the IFC is fixed within the culture and the cells are free to move vertically, our instrument will sample from only those cells which are approximately at the same vertical level as the inlet tube.

5.3.2 Materials and Methods

In the following sections, the experimental plan will be outlined, including the cell cultures, incubation environment, sampling protocol and data analysis procedures.

Cell culture

The cell culture used was the same *Rhodomonas salina* as used in the previous experiment discussed in Section 5.2. The cells were cultured in 500 mL of f/2 medium, a seawater-based medium enriched with nitrate (742 μM), phosphate (36.2 μM), vitamins, and trace metals (Hammer et al., 2002) in order to maintain nutrient-saturated growth throughout the course of the experiment. 4 separate cultures were grown in 2 L conical flasks. The cultures were inoculated into fresh, sterile medium while in their exponential growth phase, and were placed in an incubator maintained at a constant temperature of 24°C. The initial cell count was determined by agitation and sampling of the culture vessels immediately after seeding.

Light conditions

The cultures were illuminated with LED lamps emitting Photosynthetic Photon Flux Density (PPFD) of 9.6 $\mu\text{mol m}^{-2}\text{s}^{-1}$. Cultures were exposed to light in a 14-hour on, 10-hour off cycle to simulate natural lighting conditions.

The light intensity was measured using an Ocean FX UV-VIS Spectrometer (Ocean Insight, UK), and was kept at a constant level throughout the light phases for the duration of the experiment.

Sampling Procedure

The growth of *R. salina* cells was monitored over 8 days by automatic sampling and analysis using the acoustically-focussed IFC, at 4-hour intervals, with manual samples taken once per day. Manual samples were not taken at exact 24-hour intervals due to lab availability, but were ± 1 hour of the time of the first sample, other than the final sample which was taken 22 hours after the penultimate sample. Tables of all sample times (manual and IFC) are presented in Appendix E, Table 20 and Table 21.

Chapter 5

The automated samples were drawn directly from the culture flasks through LDPE tubing of 0.58 mm Inner Diameter, 0.96 mm Outer Diameter (*Smiths Medical Portex, Fisher Scientific, USA*) with an inlet approximately 2 cm from the base of the flask. Samples were pumped using a peristaltic pump (*MiniPuls 3, Gilson, USA*) connected to an electronic valve selector (*EV750-107, Rheodyne, USA*). Automated samples were of 1 mL, of which 153.755 μL (500 frames) was imaged using the acoustically focussed IFC. Samples were not recirculated to the culture vessels after analysis.

Before each automatic sampling event, 2 mL of sterile F/2 medium was automatically pumped through the IFC system in order to reduce the chance of imaging cells which had previously been left within the tubing, and to reduce the likelihood of bubbles persisting within the flow cell. The IFC was set to capture 50 images of the flow cell during this flushing in order to verify the absence of contamination. Between imaging each phytoplankton sample from the various culture vessels, a further 1 mL of sterile medium was pumped through the flow cell to further reduce the chance of cells which might have adhered to the tubing or flow cell incorrectly being imaged twice. After all 4 cultures were imaged, a further 2 mL of medium was again pumped through the flow cell for the same reasons given above. The flow cell was kept occupied with sterile medium when not sampling from the cultures as repeated wetting and drying of the flow cell increases the risk of bubble formation. When sampling from each culture, enough fluid from the culture was passed through the IFC prior to initiating imaging to be sure that only cells sampled directly from the culture would be imaged. This excess culture was discarded, along with medium used for flushing.

A diagram depicting the experimental layout is presented in Figure 95.

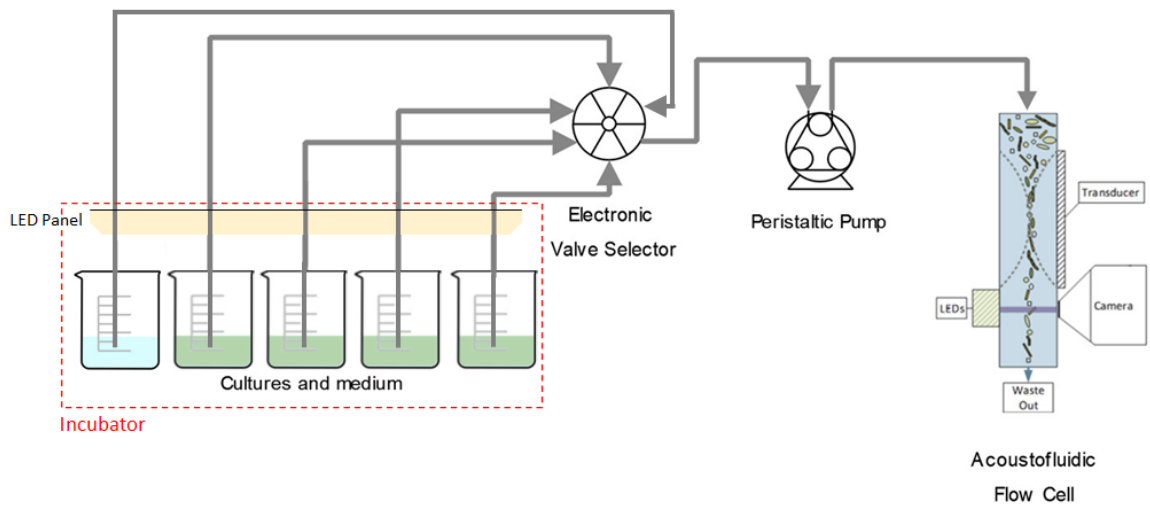


Figure 95 – Layout of system used in the growth rate experiment to measure cell count using acoustically focussed IFC.

Manual samples were drawn from the culture vessels using a standard 1 mL pipette with sterile tip, and prior to this sampling the cultures were agitated by gentle swirling for 10 seconds. After manual sampling, samples were preserved with Lugol's Iodine solution at 2% volume and stored for later microscopy analysis.

Data analysis

Phytoplankton growth rate was calculated from manual cell counts using the equation given in Equation 12:

Equation 12 – Formula for specific growth rate of phytoplankton based on cell counts (Levasseur et al., 1993).

$$\mu = \ln\left(\frac{c_n}{c_{n-1}}\right)/\Delta t$$

Where: μ is specific growth rate, c_n is the measured cell count on day n , c_{n-1} is the measured cell count on the previous day, and Δt is the time between samples (1 day). μ_{max} , the maximum growth rate, is the highest daily specific growth rate.

Chapter 5

For the automated cell counts generated by IFC, the growth rate was calculated by comparing the cell counts across different days, grouped by the time of sampling. This approach was taken so that the growth rate would only be calculated between times when the plankton would be at approximately the same vertical distribution, as comparing cell counts from the night and day would introduce error due to DVL. As there were samples processed by IFC every 4 hours during the experiment, there will be multiple separate growth rate values per sample per day. As the most important metric of cell growth is the maximum growth rate during the exponential phase, only the maximum of the 6 daily growth rates will be reported here as this will also allow a straightforward comparison with the growth rates of manually processed samples.

5.3.3 Results

After manually counting cells within 5 μL of each of the verification samples under a microscope, the cell counts per mL were extrapolated. Figure 96 presents these cell counts across the duration of the experiment. There was generally a clear S-shaped growth curve observed for each of the cultures, with a lag phase and exponential phase until nutrient limitation slows growth. Some of the cell counts (e.g. Culture #2 on Day 8) appear to be outliers.




Figure 96 - Cell counts measured through manual microscopy of pipetted and preserved *R. salina* throughout the growth rate experiment.

Chapter 5

The specific growth rates for each day of growth were calculated from these manually counted cell densities for each of the cultures, and the results are presented in Table 12. The maximum growth rates, μ_{max} , which are bolded in the table, were 1.78 for Culture #1, 0.978 for Culture #2, 0.823 for Culture #3 and 1.236 for Culture #4. These values mostly correspond well with the literature, in which a range of maximum growth rates of *R. salina* have been reported between 0.75 and 1.2 per day (Latsos et al., 2021), though Culture #1 appears to be an outlier, possibly due to an erroneously low cell count on Day 2. The maximum specific growth rate averaged across all 4 cultures is 1.20, which fits within the upper end of reported growth rates for this species.

Table 12 – Measured specific growth rate (μ) of cultures of *R. salina* determined using the cell count measured with manual microscopy after pipette sampling.

Day	Culture #1 Cell Count	Culture #1 Growth Rate (μ) (day ⁻¹)	Culture #2 Cell Count (per ml)	Culture #2 Growth Rate (μ) (day ⁻¹)	Culture #3 Cell Count (per ml)	Culture #2 Growth Rate (μ) (day ⁻¹)	Culture #4 Cell Count (per ml)	Culture #4 Growth Rate (μ) (day ⁻¹)
1	39,000	-	33,400	-	17,200	-	17600	-
2	21,600	-0.591	78,000	0.848	37000	0.766	33600	0.647
3	128,000	1.78	122,800	0.454	118400	1.16	108933	1.18
4	238,600	0.623	326,400	0.978	256600	0.773	225400	0.727
5	568,800	0.869	782,000	0.874	584200	0.823	775666	1.236
6	753,000	0.281	1,100,250	0.341	853000	0.379	868000	0.112
7	824,000	0.090	1,352,750	0.207	1042000	0.200	1087500	0.225
8	895,000	0.082	869,500	-0.441	1096000	0.050	1066500	-0.020

Chapter 5

The cell count measurements as determined by YOLOX analysis of the IFC images is presented in Figure 97. This graph is an exciting result as it serves as an excellent visual representation of the diurnal variation of the number of cells at the same vertical height as the sample inlet tube within the culture. The grey vertical bars on the graph demonstrate the times at which the samples were agitated prior to the manual sampling of the cultures using a pipette for verification under microscopy. Clearly, this agitation will affect the number of cells drawn into the IFC by homogenising the vertical distribution of *R. salina* throughout the culture vessels, but it is apparent that the cell counts rise before these mixing events and are still falling over 4 hours post-agitation, by which time the cells should have returned to a non-agitated state due to their motility. The cell counts within the L1 medium (sterile culture) used to flush the flow cell are extremely close to zero throughout the course of the entire experiment, as would be expected, which is good evidence that contamination and/or adhered cells within the fluidic network or flow cell were not a problem during the experiment.




Figure 97 - Automatically measured cell counts based on IFC images of the 4 cultures of *R. salina* during the growth rate experiment. L1 is sterile medium, and acts a control.

For a direct comparison between the verification cell counts measured using microscopy and the automated cell counts using IFC, and to relate the diurnal variation of cell counts observed within the IFC data to the lighting condition within the incubator, a combined plot is presented in Figure 98. This plot demonstrates the DVL behaviour induced by light availability, which is shown with

Chapter 5

the orange bars on the graph. The graph also makes it clear that the manual cell counts do not correspond well with the automated cell counts, an issue which will be explored in detail in the following section.




Figure 98 – Combined graph of manual (dots) and automatic IFC (lines) cell counts measured throughout the growth rate experiment. Orange bars represent the times when the incubator’s lights were on, and black vertical lines are the times at which the samples were agitated and manual samples drawn off for verification. L1 is sterile medium, and acts a control.

The full table of specific growth rates calculated using the IFC-measured cell counts, grouped by sampling time, is presented in Table 22 in Appendix F. Note that a problem with the imaging system caused by a computer error resulted in only one cell count being determined at 09:00 and 11:00 sample times, hence no growth rate could be calculated for those time points.

The maximum calculated growth rates, regardless of sampling time, using IFC cell counts, are as follows:

Culture #1: **1.75**; Culture #2: **1.34**, Culture #3: **1.47**, Culture #4: **1.56**. (day^{-1})(All to 3 s.f.)

For Culture #1, the mean maximum specific growth rate is 0.700 day^{-1} , with a standard deviation of 0.408. Culture #2 has a mean growth rate of 0.654 day^{-1} , with a standard deviation of 0.251. Culture #3 exhibits a mean growth rate of 0.913 day^{-1} , with a standard deviation of 0.457. Lastly, Culture #4 has a mean growth rate of 0.994 day^{-1} , with a standard deviation of 0.265. (All to 3 s.f.).

While these calculated maximum growth rates align with the literature for the study species, the range of values obtained depending on when the sample was taken, demonstrated by the large standard deviations and differences between the mean and maximum values, illustrates the challenge of calculating growth rates based on continuous measurement of unagitated samples. Furthermore, the difference of these values calculated using the IFC cell counts with those obtained from manual cell counting under microscopy, though unsurprising given the difference between manual and automated counts shown in Figure 98, provides additional evidence of a problem with the measurement protocol. The possible causes of these differences will be explored in the following section.

5.3.4 Discussion

In the results presented in the previous section, it was observed that cell counts obtained through traditional microscopy were consistently and significantly greater than those derived from automated IFC for all samples after the initial acclimatisation phase of the cultures. This also caused significant differences between the growth rates calculated using measurements from each method. Our systematic approach to investigating the cause of this difference is detailed below. Firstly, several potential causes are ruled out, then, evidence is presented for a hypothesis that the discrepancy is caused by active motion of the plankton actively avoiding being drawn into the inlet tubing.

Failure of Acoustic Focussing?

In order to make sure that the reduced cell count derived from the IFC images was not caused by a failure of the acoustic focussing mechanism, which would have resulted in cells being out of focus and hence not counted by the object detection algorithm, a manual investigation of the captured images was carried out. The successful functioning of acoustic focusing was confirmed

Chapter 5

by the high percentage of in-focus cells captured in the images across the 8 experiment days. If the acoustic focusing had not been functioning properly, only a small proportion of the cells would have been in focus at any given time, with the other cells present in the imaging region randomly distributed through the imaging axis. This qualitative check gives high confidence of the functional operation of the acoustic focussing when the number of cells in each captured frame is high, as it is easy to subjectively see the difference between all cells being near to the same focus plane as opposed to randomly distributed through the thickness axis of the flow cell.

Analysis revealed that across all the images collected during the experiment, only a very small number of cells were out of focus. This result provides strong evidence that the IFC's acoustic focusing mechanism was effective in concentrating cells into a single plane throughout the entire duration of the experiment. Several example images taken from different days of the experiment were analysed and a typical frame is presented in Figure 99 which clearly demonstrates that even when there are many cells present in a frame, they are all focussed acoustically.




Figure 99 - A single frame taken from the IFC during imaging of *R. salina* cells during the growth rate experiment, demonstrating successful acoustic focussing of all cells.

Cell loss in Fluidic System?

In addition to the correct operation of the acoustic focussing demonstrated in the IFC, we note that our experiments have also shown that the cells were not lost due to adherence to the flow cell or connected tubing (i.e. between the culturing vessels and the IFC). This was determined through the experiment detailed in Section 5.2, where a sample of fixed cells was imaged using the IFC and subsequently collected for manual inspection under a microscope. The sample passed through several centimetres of the flow cell after being imaged, then through a significant length (20cm) of tubing before being collected for manual assessment. If cells were able to be lost due to adherence or deposition within the tubing or flow cell, this would have resulted in discrepancies between the cell counts obtained from the IFC images and those observed manually. The fact that this was not observed in our experiment suggests the cause of the observed discrepancy in cell counts within this experiment is not due to cells being pumped out of the culture vessel and then depositing within the fluidic system before imaging.

Difference in Sampled Populations due to Diel Vertical Migration?

Manual sampling was performed after sample agitation to ensure a uniform distribution of cells, and should accurately represent the population of phytoplankton within the culture vessels. Automatic IFC sampling which occurred without prior agitation of the culture, represents only a count of the cells suspended near the inlet of the tubing connected to the peristaltic pump, and would therefore the two measurements would be expected to differ if cells occupied discrete vertical layers. The cyclical fluctuations in cell concentration observed within the IFC-analysed samples were therefore anticipated due to the fixed position of the sampling tube inlet relative to the vertical motion of phytoplankton due to Diel Vertical Migration (DVM), but these fluctuations fail to account for the observed discrepancies. Notably, DVM was ruled out as the sole cause of the disparity between manual microscopy and IFC cell counts due to the capture of one set of samples within 10 minutes after agitation of the cultures in preparation for microscopy (the final day of the experiment, Day 8). Had solely DVM explained the difference between manual and automated counts, these samples should still have been homogenised due to agitation and IFC-derived counts should therefore have been identical to manual microscopy.

Best Hypothesis: Active Avoidance of Sampling

In order to identify possible remaining causes of the observed discrepancy between IFC and microscopy counts, literature on cell motility was reviewed. Although it was known that *Rhodomonas salina* cells were motile, using their flagella to propel themselves vertically within the culture as part of DVM behaviour, the possibility of them intentionally evading a source of suction had not been considered. Evidence from literature is presented below showing that phytoplankton such as the *R. salina* cells studied in this experiment are capable of rapidly responding to external mechanical stimuli as a predator-avoidance mechanism, and are able to swim away from inlet tubing when they sense the accompanying fluid shear stresses. A hypothesis that this is the source of the reduced automatic counts is therefore supported.

This hypothesis also extends to the observed variability in manual cell counts. As discussed previously, the cell count of a growing phytoplankton population typically occurs in an S-shaped curve. After an initial acclimatisation phase, cell count grows exponentially, before levelling off as nutrient limitation occurs at high population density. This S-shape is mostly observed in the manually counted data, but several unexpected outliers (e.g. Culture #2 on the final day has a recorded cell count which appears erroneously low compared to previous measurements of the same population. This and other inconsistent manual cell counts could plausibly have resulted from the escape of cells from the pipette suction (which also causes shear stresses) while sampling. The possibility of pipetting causing damage to cells by inducing fluid shear stress has been investigated in the research (e.g. (Lund, 2016, M. Art et al., 2023)), and is recognised to trigger the escape response of zooplankton (Singarajah, 1969), but does not appear to be generally considered within phytoplankton research. The greater variability in manual counts would correspond to variability of draw rate; unlike the automated sampling which used a fixed volumetric flow rate of 0.68 mL/min, manually pipetting can result in different draw rates depending on how quickly the plunger is released.

If this hypothesis is correct, the population of cells imaged by the IFC could have consisted only of those cells with reduced motility, possibly due to compromised physiological state or even cell death. It is therefore reasonable to hypothesise that this selective sampling of a small subsection of the overall population ultimately led to the observed underestimation of cell counts when utilizing automated imaging flow cytometry in comparison to manual microscopy.

Chapter 5

Before assessing whether the cells could have avoided the sample inlet tube while pumping was underway, it is important to consider whether cells sampled for imaging by IFC would have been able to swim away from the acoustic focus plane after passing the acoustic transducer, as this would lead to them being out of focus by the time they were imaged. Previous experimental work investigating the effect of acoustic focussing on motile phytoplankton has demonstrated that the focussing mechanism does not disable their motility (Kim et al., 2021), though while in the acoustic focussing region above the transducer, the acoustic forces can be strong enough to overcome their motion away from the acoustic focus plane.

In the IFC presented there is, however, a gap of approximately 2mm between the edge of the acoustic transducer (where cells exit after acoustic focussing) and the imaging region (in front of the objective lens), in the direction of flow. This gap is due to the size of the objective making it impractical to position the imaging area closer to the edge of the transducer. Although it was previously noted that cells were, apart from a small percentage of outliers, in focus within the images captured during the experiment, the likelihood of having too low a cell count due to the organisms swimming out of focus of the objective is also considered below.

Blackburn et al. (2022) determined that phytoplankton generally swim at speeds equivalent to 20 times their body length per second, which, for *R. salina*, translates to approximately 200 $\mu\text{m/s}$. Previous research using the exact species studied in this experiment measured swimming velocities of over 150 $\mu\text{m/s}$ (Jakobsen et al., 2006). A separate study involved the measurement of the swimming velocities of *R. salina* cells under light and dark conditions to determine the effect on motility of photosynthetically available radiation and measured $82 \pm 19 \mu\text{m/s}$ and $104 \pm 22 \mu\text{m/s}$ in the light and dark conditions respectively (Kaňa et al., 2019).

At the volumetric flow rate utilised in this study (0.68 ml/minute), the velocity of fluid through the flow cell within the acoustic focus plane (in the centre) is 11.1 mm/s, giving the phytoplankton approximately 0.18 seconds between leaving the acoustic focussing region 5mm above the imaging region, and being imaged. The depth of field of the 10X magnification, 0.25NA objective lens used in the IFC is approximately 8 μm (Liu and Hua, 2011). If it is assumed that the sampled cells exhibited swimming velocities of 100 $\mu\text{m/s}$, which would be supported by the literature reviewed above, cells could swim over 18 μm between leaving the transducer and presenting for imaging. Less than half of this distance would need to be swum perpendicular to the imaging

plane for cells to appear heavily out of focus within the captured images and hence not be successfully counted by the image processing algorithm.

The fact that we do not see many cells out of focus within the captured images suggests that the cells that entered the IFC were not as mobile as the literature suggests. Given that the swimming velocities of *R. salina* cells are expected to follow a normal distribution with a mean around 100 $\mu\text{m/s}$ (Kaňa et al., 2019), it follows that a significant population of cells within the images would be expected to be out of focus if a representative sample of the complete population was drawn up for IFC imaging. This supports the hypothesis that cells were escaping the flow field of the inlet tubing as in this case it would be only the less motile cells that were taken up and imaged.

Research has shown that motile phytoplankton are capable of detecting predation threats by sensing shear in the fluid flow around their bodies, and actively swim away from high shear forces (Barry et al., 2015). In addition to swimming away from these forces, some motile phytoplankton can "jump" - rapidly accelerating to high swimming velocities greater than 5-fold their usual swimming velocities - in an alternate direction. Jakobsen, Everett, and Strom (2006) examined the predation of *R. salina* by predatory ciliates, measuring the mean swimming speed of *Rhodomonas* cells at 153 microns per second. Upon detecting the presence of predators, however, the *Rhodomonas* cells accelerated to speeds of up to 950 microns per second (mean), effectively preventing their capture by the ciliates.

Jakobsen (2001) conducted an investigation into the escape behaviours exhibited by plankton exposed to fluid mechanical signals induced by a siphon flow (a steady flow into a narrow tube), in a setup closely mirroring the inlet arrangement in the present work. The study focused on various plankton species of a comparable size to the *R. salina* investigated within this experiment. Fenchel and Juel Hansen (2006) also conducted an experiment that paralleled Jakobsen (2001), with both studies employing narrow siphons (of comparable diameters to the tubing used within this experiment to pump fluid from the culture vessels to the IFC) while simultaneously observing the behaviour of motile plankton in response. In both cases, the researchers found that the phytoplankton cells were capable of detecting and evading the siphon's suction at flow rates similar to those employed in the present study.

Chapter 5

In order to determine whether these escape swimming velocities would have allowed the main population of healthy, motile cells to avoid capture by the sampling pump, a comparison the flow rates used within this experiment with those investigated in the above studies will be carried out. Fenchel and Hansen (2009) witnessed escape behaviour caused by a siphon of 0.15 mm diameter, measuring flow velocity of 1 cm/s just outside the tip of the siphon. Jacobsen (2001) used a siphon of 0.25 mm inner diameter with a velocity of approximately 1 cm/s outside the siphon tip. In the experiment presented within this chapter we used a tube with 0.58 mm inner diameter tip. At a flow rate of 0.58 mL/min, this would result in a fluid velocity of 4.23 cm/s (3s.f.) occurring at the tube inlet, which is greater than but similar to the values utilized in the aforementioned research.

It is therefore a logical conclusion that the shear forces exerted by the pumping of phytoplankton culture from the culture vessels into the IFC were almost certainly great enough to trigger the escape response of the *R. salina* cells, and this is the reason for the significant difference between manual and automated sampling approaches is supported. Furthermore, the inconsistencies in the manual counts are likely due to a fraction of the phytoplankton successfully escaping the pipette tip, with outliers caused by different draw rates.

Although it is not within the scope of this experiment, future experimental work to quantify the shear rates at which this escape behaviour occurs could be carried out using different draw rates, either using a similar automatic sampling setup or an electronically controlled pipette.

The motility and escape responses of zooplankton, the heterotrophic organisms which typically predate phytoplankton, have been well-understood for decades (Singarajah, 1969, Fleminger and Clutter, 1965), and some sampling instruments such as the *Zooplankton Sampler* (McLane, USA) are specifically designed to induce low fluid shear stress to avoid triggering escape behaviour. However, the possibility of the same behaviour regularly occurring in motile phytoplankton species does not seem to be widely considered within the research.

Our literature review uncovered only a single experiment on phytoplankton in which researchers had discussed the potential introduction of bias due to phytoplankton motility. In Olson et al.'s investigation of the addition of acoustic focussing to the *Imaging Flow Cytobot (IFCB)*, the authors

Chapter 5

note that motile cells may manage to escape the acoustic focus plane before imaging (Olson et al., 2017), though escape of the initial sampling prior to acoustic focussing does not appear to have been considered. Moreover, in-situ direct sampling from aquatic environments, such as by the IFCB, is increasingly utilised as autonomous technology matures. When samples are pumped continuously from the environment, there may not be sufficient agitation (e.g., by waves, wind) for the phytoplankton to be unable to detect and avoid the shear induced by the pump systems used. IFCB (Olson and Sosik, 2007b), for example, has a comparable inlet tube inner diameter to the inlet used in our experiment (0.762 mm) and operates at 15 ml/hour (Blackburn et al., 2022). At these flow rates, it is plausible that the IFCB pumping system could also trigger escape responses in motile phytoplankton and hence lead to their underrepresentation within the imaged population.

In conclusion, our findings and the lack of attention to phytoplankton motility within much literature in the field suggests a possible systemic problem which could lead to incorrect assessments of the percentage of some motile species within aquatic environments, and a bias in the cell count in monoculture experiments such as that carried out here.

Chapter 6 Conclusion & Future Work

6.1 Conclusion

This thesis has presented the development, characterization and application of a novel, low-cost, high-throughput, acoustically-focussed Imaging Flow Cytometer for real-time, continuous monitoring of phytoplankton cells.

In Chapter 2, the importance of phytoplankton to the climate, aquatic food chain and global economy was highlighted. Phytoplankton provide around half of all carbon fixation and are the primary producers on which the overwhelming majority of marine and freshwater life depends. Phytoplankton cause Harmful Algal Blooms which can drastically harm ecosystem health and have severe health and financial effects on coastal communities.

Many research challenges can be addressed using imaging flow cytometry. Without early-detection by *in situ* measurement platforms, it is not possible to predict the formation of algal blooms. Bulk measurement techniques such as satellite imaging and fluorometry do not have the species-level taxonomic resolution needed to tell the difference between benign and harmful blooming species.

Counting the number of cells of each phytoplankton species within a water sample is a fundamental method within oceanographic microbiology and is commonly used to provide an assessment of ecosystem health and carbon fixation. This process commonly requires the labour-intensive use of manual microscopy and specialised taxonomic knowledge, professions which are known to be in decline. Furthermore, industrial applications such as aquaculture and algal biotechnology require an in-depth understanding of the biomass, health and composition of algal cultures, which again requires the use of imaging techniques.

Recent global reviews have highlighted the need for far more widespread use of *in situ* imaging sensors for addressing some of these research questions. The limiting factors slowing the uptake of existing IFCs include their prohibitive cost and low throughputs, challenges which are sought to be addressed with the IFC developed in this thesis.

Chapter 6

Chapter 3 detailed the development process of a low-cost, disposable flow cell which uses an ultrasonic transducer to induce an acoustic standing wave within the fluid flow. This acoustophoretic approach allows the focussing of cells into a plane at a known and consistent distance from a microscope objective such that all cells can be imaged at high speed and quality. Work presented in Section 3.3 involved the modelling and experimental evaluation of the acoustically focussed flow cell and build confidence that the acoustic wave would reliably focus phytoplankton cells across a wide size spectrum. Finally, a discussion of the limitations of our developed experimental approach and possible alternatives was presented.

The development of a robust protocol to measure the imaging resolution of the IFC was presented in Section 3.4. Having a reliable method by which to measure the minimum cell size which could be imaged by the IFC is critical to assessing the phytoplankton which can be analysed using this device. Some phytoplankton are extremely small ($< 1 \mu\text{m}$) and thus water samples will typically always contain plankton below the imaging resolution of optical systems. That said, that the $\approx 3\mu\text{m}$ resolution of our device compares extremely favourably with far more expensive commercial instruments and allows the analysis of a wide range of plankton species.

Section 3.5 details the design and evaluation of the integrated system, including flow cell, imaging and illumination optics. The design was aimed to be low-cost and easily modifiable platform which uses predominantly off-the-shelf components in order to create a device which is maximally accessible to resource-constrained research laboratories. A suitable camera and pulsed illumination optics such that images will be captured at a high resolution without motion blur, even at high flow rates, were used. The application of the aforementioned imaging resolution measurement protocol is described, and the limitations of the approach discussed.

In Chapter 4, the development of an approach to the automatic analysis of the images captured by the IFC, was presented, comparing several different approaches. First, traditional image processing-based approaches are summarised and their limitations explored. Next, the development of an improved version of a convolutional neural network-based object detection model used for a similar cytometer by another research group was described, but unfortunately it was incapable of counting the cells within images with a performance equal to that of a human – a critical metric for the reliable use of an autonomous algorithm. Finally, Sections 4.4.6 and 4.4.7 present the new use of a cutting-edge Convolutional Neural Network-based object detection

Chapter 6

model, YOLOX, which successfully allows the automatic counting of cells within images with a performance equal to a human operator. This novel use of YOLOX for IFC images has a significant potential to allow the *in situ*, reliable and fully autonomous measurement of cells in aquatic environments as it can run on an embedded GPU computing system designed for deployment. Not only is it demonstrated that can YOLOX successfully count phytoplankton cells, there is also the potential for it to be trained to simultaneously classify species or other classifications of interest by extending the dataset on which it was trained.

As IFC instruments are increasingly developed to be run continuously over long periods without human intervention, it is critical that reliable and consistent algorithms process the vast amounts of image data they generate. Our implementation of YOLOX successfully addresses this important research challenge in a form that is suitable for remote deployment and real time operation.

Chapter 5 describes two experiments that were undertaken to evaluate the performance of the IFC on real phytoplankton research problems. The first, presented in Section 5.2, involves the analysis of a simple, single-species sample of preserved phytoplankton cells. This work is important as many research scenarios involve the need to analyse preserved samples and it is critical to have confidence in the ability of the complete IFC system combined with the autonomous image analysis algorithm described above, to accurately count cells in a sample. It was found that the IFC successfully counted cells with the same accuracy as a human using the gold standard of manual microscopy.

In Section 5.3, experiment was introduced in which the IFC was used to measure the cell count of growing cultures of phytoplankton, at a high temporal resolution, throughout the course of the growth phases of the cultures. This experiment demonstrated the ability of the IFC to operate continuously and autonomously over a significant period of time, a capability not shared by most commercial IFCs such as the FlowCam. Though the aim of the experiment had been to demonstrate the accuracy of counting of plankton cells using the IFC by comparing the cell counts with manual verification samples analysed under microscopy, the experiment highlighted a sampling issue (discussed below) with potentially substantial implications across varied phytoplankton research scenarios.

Chapter 6

Instruments including IFCs, flow cytometers, and other phytoplankton measuring devices such as fluorometers are increasingly developed to operate *in situ* on living cells in addition to analysing preserved samples taken at a prior date as has historically been more common. As the sophistication and autonomy of these devices progressively expands, allowing them to generate useful scientific data with low or even zero manual verification of their measurements, it is crucial that these instruments accurately quantify the abundance of plankton across the range of species they sample.

Our experiment clearly demonstrates the risks of incorrectly measuring the number of phytoplankton cells within a water sample due to the avoidance behaviour of the cells toward sample inlets. There is existing literature in which experiments have measured the motility of some flagellated microalgae, and the ability of cells of species such as the *Rhodomonas salina* used in our experiment to detect and escape from regions of high fluid shear. There does not, however, appear to be any research into the possible bias that such predator avoidance behaviour could introduce into sampling of natural phytoplankton assemblages.

It is possible, therefore, that instruments such as the *Imaging Flow CytoBot*, in addition to the many autonomous phytoplankton measurement instruments either recently released or under development, incorrectly report the number of motile plankton due to them avoiding being drawn into the device. It is also of great significance that even manual pipetting of these cells may fail to accurately capture a representative sample of their population density. These systemic problems could already have led to incorrect data in multiple research projects, a risk which will only increase as sensor development continues. The following section will explore possible future work to quantify the problem and suggest possible solutions.

In conclusion, this thesis has presented a novel, low-cost, easy to assemble and easily modifiable imaging flow cytometer based around a disposable, acoustically-focussed flow cell and demonstrated its potential to address research questions involving measurement, counting and classification of phytoplankton. This instrument has the potential to significantly advance the rate at which images of plankton can be captured around the world by providing a higher-throughput and cheaper alternative to the small number of currently available devices.

6.2 Future work

In order to further explore sampling bias due to phytoplankton motility, future experiments could be devised to assess the magnitude of the problem. These could be straightforward to carry out – for example, by collecting a sample of a culture through a tube of fixed aperture and pumped at a specified rate; counting cells within that pumped sample; then preserving the entire culture and measuring the cell density. By comparing the cell counts of the same culture where one sample has been pumped and one has not, it would be trivial to measure the percentage of cells which avoid sampling. The experiment could be repeated with various pumping velocities and using different inlet tube diameters, including conditions identical to those used on existing instrumentation. Other experiments could be carried out to further investigate the extent to which the escape behaviour is present across all phytoplankton species, which would allow the determination of the magnitude of the problem in real-world *in situ* sampling applications.

It is also possible to devise experiments to test potential solutions to the problem. Since it has been demonstrated experimentally that detection of fluid shear around the cell body is the trigger for phytoplankton escape behaviour, it is likely that affixing a funnel to the end of the sample inlet, hence reducing the shear rate, would allow sampling without triggering escape jumps. Therefore, future experiments should test various inlet geometries. Experimentally this would also be straightforward to achieve by using a similar experiment design as above, where replicate cultures would be sampled by different inlet types and another sample preserved before counting in order to elucidate the contribution of each inlet type on measured cell count after sampling.

Section 5.3.4 discussed the fact that motile plankton cells would have been able to swim out of the focus plane of the instrument if they had been sampled in the first place. To address this limitation, experiments with motile cells could be conducted where the voltage across the acoustic transducer was increased or the flow rate increased to reduce the time between focussing and imaging such that cells cannot swim out of focus.

In the justification of this project (Chapter 2), it was noted that there is a lack of *in situ* sensors for measuring phytoplankton abundance and diversity. In its current state, it is only possible to use the IFC presented within this thesis in an indoor environment. In order to further develop the IFC to allow its use in the field, there are several additional engineering challenges which could be

addressed. To begin with, the power supply to the device, which is currently a mains (230V AC) cable, could be replaced with a battery. This would allow use of the IFC in environments without access to electrical infrastructure. The peristaltic pump that was used in this work (*Gibson MiniPuls3*) is heavy and not designed for operation outside of a laboratory. This pump could be replaced by an off-the-shelf lightweight pump or a custom 3D-printed pump similar to that used by the *PlanktonScope* discussed in Section 2.5.2. For the IFC to be used while exposed to the elements, water ingress into the device would need to be prevented. This could be achieved by means of an external enclosure surrounding the flow cell, optics and electronics, with an inlet and outlet for providing the necessary fluid path for the water to be sampled.

To enable autonomous operation, a biocide reservoir could be added and the flow cell automatically purged and cleaned on e.g. daily intervals, to prevent or delay biofouling of the fluidic system. Filtering of the input to prevent the ingress of large particles which may obstruct the flow cell would likely also be required and could be achieved by means of a standard e.g. 100 μm nylon mesh in the input pipe. Due to the extremely variable number of phytoplankton cells present in natural water samples, pre-processing of input water may be needed; concentration of the samples could be achieved using a sedimentation step before imaging the sample, though this would require work to ensure settling does not create new biases towards less motile species. Telemetry and image data from the device could be enabled by wired or wireless network connection to allow the remote monitoring and data connection without the need for a monitor and keyboard to be connected to the device.

Finally, further training of the machine learning models implemented in Section 4.4 would allow the use of our IFC for a greatly expanded range of phytoplankton research problems. While the utility of our models on single-species cultures grown in a laboratory was demonstrated, using freely available phytoplankton microscopy dataset such as PMID (Li et al., 2020), it would be possible to train our model to classify cells within mixed samples of plankton communities down to the species level. Doing so would be relatively straightforward and increase the scientific value of the data generated by the IFC, but care would need to be taken to ensure the phytoplankton species represented in the training data was appropriate for the environment from which samples analysed by the device originated.

6.3 Summary of Novel Contributions

This thesis presented a body of work involving the design, modelling and testing of a new acoustically focussed flow cell, building on a previously published design. The flow cell was improved by modifying the layers to incorporate a cover slip for reduced optical thickness, and therefore improved imaging performance. This layered design was characterised for its acoustic resonances and their associated acoustic focussing performance. Finite Element Modelling was used to predict the performance of the flow cell for different phytoplankton sizes and types, demonstrating its applicability across a wide range of cells, even at high flow rates compared to existing instruments. The work of previous authors was built upon by carrying out novel experiments to quantify uncertainties in the MTF method for measuring image quality and, applying this knowledge, the work here shows that the optical performance of our device is sufficient to capture images of even very small (< 5 μm diameter) phytoplankton. This is not possible using commercial instruments such as the FlowCam, which has a reduced minimum cell size due to its depth-of-focus-increasing optical setup.

A complete, low-cost IFC system around the acoustically-focussed flow cell, a 3D-printed flow cell holder and a rail system constructed from off-the-shelf parts was designed and presented. This ensures that our system will be accessible and easily modifiable, thereby increasing access to high-quality IFC for oceanographic microbiology.

A novel convolutional neural network was developed to detect and localise cells within the images from the IFC, which although based on an existing approach, integrated more modern developments in computer vision. After extensive testing this model was replaced by an implementation of a cutting-edge object detection model which was demonstrated to have the ability to count cells within images with a performance comparable to a human operator.

Finally, the complete IFC system was tested, which showcased the ability of the IFC to accurately measure the cell density of a preserved plankton sample with the same performance as manual microscopy. This highlights the potential of the developed instrument to replace or improve the traditionally labour-intensive counting of cells using traditional techniques. Furthermore, given the ability of the machine learning image processing technique utilised to be trained on images of multiple different species, the complete system has the potential to allow rapid and automated

Chapter 6

analysis of preserved samples of real-world natural phytoplankton assemblages to characterise the complete aquatic population of microalgae within its operating limits.

The ability of our system to carry out high temporal resolution, autonomous operation over a significant period was demonstrated by measuring the growth of live cell cultures in an incubator. This application demonstrates the potential of the IFC to be used in aquaculture and biotechnology applications, with potentially significant improvements in monitoring and cost reductions. While further experiments would be required to optimise the device with respect to motile cells which have been shown to possess the ability to actively avoid sampling, the instrument presents a novel contribution to aquatic research and satisfies the aims and objectives laid out at the beginning of this thesis.

Appendix A Parameters for KLM Model of Layered Resonator Flow Cell

Transducer:

Table 13 – Table of parameters defining the transducer used within the KLM model of the acoustically-focussed flow cell.

Parameter	Value	Unit
Thickness	1.0	[mm]
Voltage (p-p)	10	[V]
Qm	100	
Length	2	[cm]
Width	5	[cm]
Dielectric Loss	0.003	
Piezoelectric Pressure Constant	2.37×10^9	[NC ⁻¹]
Permittivity	6.195×10^{-9}	[Fm ⁻¹]
Sound Velocity	4529.8	[ms ⁻¹]
Density	7700	[kgm ⁻³]
Terminating Material (air) Sound Velocity	331.6	[ms ⁻¹]
Terminating Material (air) Density	1.293	[kgm ⁻³]

Appendix A

Particles:

Table 14 - Table of parameters defining the particles used within the KLM model of the acoustically-focussed flow cell.

Parameter	Value	Unit
Sound Velocity	1962	[ms ⁻¹]
Radius	5	[μm]
Density	1055	[kgm ⁻³]

Device Layers:

Table 15 - Table of parameters defining the layers of the flow cell used within the KLM model of acoustic focussing.

Layer	Thickness [μm]	Density [kgm ⁻³]	Sound Velocity [ms ⁻¹]	Q
Glue Gap	1	1080	2640	100
Matching Layer	170	2540	5510	100
Fluid Layer	390	(Defined by T, S)		100
Reflector	1100	2500	5872	100
Terminating Material		1.293	331.6	

Fluid Channel:

Table 16 - Table of parameters defining the fluid channel inside the flow cell, used within the KLM model.

Width	4	[mm]
Length	6	[cm]

Appendix B COMSOL Finite Element Model Parameters and Variables

Parameters

Table 17 - Table of parameters used in Finite Element Modelling of the acoustically-focussed flow cell.

Parameter	Expression	Value	Description
η	0.89 [mPa * s]	8.9E-4 Pa·s	Dynamic Viscosity of Medium
a	5[um]	5E-6 m	Particle radius
c_0	1500[m*s ⁻¹]	1500 m/s	Sound Velocity of Medium
ϕ	0.279	0.279	Acoustic Contrast Factor
k	$\frac{2 * \pi * f}{c_0}$	8055.4 1/m	Wavenumber
f	$\frac{c_0}{2 * height}$	1.9231E6 1/s	Frequency
q	1[mL/min]	1.6667E-8 m ³ /s	Volumetric Flow Rate
A	height*width	1.56E-6 m ²	Cross-sectional area of channel
height	390 [um]	3.9E-4 m	Channel height
width	4[mm]	0.004 m	Channel width
length	5[cm]	0.05 m	Channel length
ρ_0	1023 [kg/m ³]	1023 kg/m ³	Density of Medium
E_{ac}	$\frac{1}{4} * \frac{p^2}{(\rho_0 * c_0^2)}$	29.369 Pa	Acoustic energy density
p	520 [kPa]	5.2E5 Pa	Acoustic pressure amplitude

Variables

Table 18 - Table of variables used in Finite Element Modelling of the acoustically-focused flow cell.

Variable	Expression	Unit	Description
F	$4\pi a^3 * E_{ac} * k \sin(2ky) \varphi$	N	Acoustic Radiation Force
G	$\frac{q}{width} * \frac{12\eta}{height^3}$	N/m ³	Pressure Gradient (dP/dx)
vx	$\frac{height^2}{2\eta} * G * \left(\frac{y}{height} * \left(1 - \frac{y}{height} \right) \right)$	m/s	Particle velocity, x component
vy	$\frac{F}{6\pi\eta a}$	m/s	Particle velocity, y component

Appendix C Pseudocode for Bead Detection and Separation Measurement in Double-Exposure Images

```

Load Image
Gaussian Filter Image -> BlurredImage
Binarize BlurredImage -> BinaryImage
Erode BinaryImage -> ErodedBinaryImage
Open ErodedBinaryImage -> OpenedBinaryImage
Close Holes OpenedBinaryImage -> ClosedBinaryImage
Watershed Algorithm ClosedBinaryImage -> WatershedImage
Region Proposal Algorithm WatershedImage -> ImageStats
For Stat in ImageStats:
    If Stat.Area > ThresholdSize:
For Each Centroid:
    Label Centroid
    Prompt User to Input Labels
    While UserInput != "q":
        For UserInput1, Userinput2:
            Calculate Centroid1, Centroid2 Distance, Angle
            Store Distance, Angle
Load Next Image
  
```

Where: *ThresholdSize* is the minimum size (in pixels) a bead image occupies. This must be set to avoid spurious detections of noise and/or small detritus. A value of 8000 worked well for this analysis.

Appendix D Bill of Materials for IFC

Table 19 - Bill of Materials to construct the optical, imaging and computational hardware for the IFC.

Supplier	Component	Description	Quantity	Total Cost (UK£)
Thorlabs	CP20S	Cage System Iris Diaphragm	2	140.8
Thorlabs	ER10	Cage Assembly Rod, 10" Long, Ø6 mm	4	39.24
Thorlabs	ER12	Cage Assembly Rod, 12" Long, Ø6 mm	3	38.94
Thorlabs	CP33T/M	SM1-Threaded 30 mm Cage Plate, 0.50" Thick	2	34.9
Thorlabs	SM1ZM	SM1 Zoom Housing for Ø1" Optics	1	132.29
Thorlabs	SM1V15	Ø1" Adjustable Lens Tube, 1.31" Travel Range	1	27.62
Thorlabs	SM1V05	Ø1" Adjustable Lens Tube, 0.31" Travel Range	1	23.37
Thorlabs	LCP02/M	30mm to 60 mm Cage Plate Adapter, M4 Tap	5	158.25
Thorlabs	SM1L40	SM1 Lens Tube, 4.00" Thread Depth	1	35.46
Thorlabs	SM1L15	SM1 Lens Tube, 1.50" Thread Depth	1	12.13
Thorlabs	RMSA1	Adapter with External M25 x 0.75 Threads and Internal RMS Threads	1	15.74
Thorlabs	SM1A39	Adapter with External C-Mount Threads and External SM1 Threads	1	15.91
Thorlabs	ER90C	90 Degree "T" Extension	4	38.64
Thorlabs	ACL2520U-DG6	Aspheric Condenser Lens w/ Diffuser, Ø25 mm, f=20.1 mm, NA=0.60	1	14.69
Edmund Optics	47-637	25mm Dia. x 50mm FL, VIS 0° Coated, Achromatic Lens	2	161.5

Appendix D

Edmund Optics	33-438	20X DIN Achromatic Finite Intl Standard Objective	1	97.75
Mouser Electronics	LZ4-40CW08-0065	LED Engin High Power LED - White	1	15.43
FLIR	GS3-U3-23S6C-C	Grasshopper3 USB3 2.3 MP, 163 FPS, Sony IMX174, Color	1	842.4
DFRobot	Xavier NX	NVIDIA Jetson Xavier NX Developer Kit	1	395.41

Appendix E Growth Rate Experiment: Manual and Automated Sample Times

Table 20 - Table of sampling times of the *R. salina* cultures during the growth rate experiment (Section 5.3) which were preserved and manually counted using microscopy.

Date	Time
23rd March 2022	16:30
24th March 2022	16:00
25th March 2022	16:00
26th March 2022	15:30
27th March 2022	15:30
28th March 2022	17:30
29th March 2022	16:30
30th March 2022	14:30

Appendix E

Table 21 - Times at which *R. salina* samples were automatically pumped into the IFC for imaging.

The IFC was restarted on 26th March at 15:00 due to a connectivity issue, which resulted in an additional sample being processed that day and a change in the hours of sampling on subsequent days.

Day	Sample 1	Sample 2	Sample 3	Sample 4	Sample 5	Sample 6	Sample 7
23rd March 2022	17:00	21:00					
24th March 2022	01:00	05:00	09:00	13:00	17:00	21:00	
25th March 2022	01:00	05:00	09:00	13:00	17:00	21:00	
26th March 2022	01:00	05:00	09:00	13:00	15:00	19:00	23:00
27th March 2022	03:00	07:00	11:00	15:00	19:00	23:00	
28th March 2022	03:00	07:00	11:00	15:00	19:00	23:00	
29th March 2022	03:00	07:00	11:00	15:00	19:00	23:00	
30th March 2022	03:00	07:00	11:00	15:00			

Appendix F Calculated Specific Growth Rates from IFC Counts of *R. salina* During Growth Rate Experiment

Table 22 – Specific Growth Rate of each of the 4 cultures, calculated using the cell counts measured by IFC, using only those measurements taken at the same time in order to reduce the potential influence of diurnal changes in vertical distribution.

Bold=maximum growth rate.

	Calculated Maximum Specific Growth Rate across the 8 experimental days, grouped by time of sampling			
Sample Time	Culture #1	Culture #2	Culture #3	Culture #4
01:00	0.258134	0.507689	-0.09	1.201591
03:00	0.827393	0.708401	1.466809	0.674891
05:00	0.43438	0.498359	0.712013	0.916474
07:00	1.752674	0.647959	1.460398	1.564099
09:00	-	-	-	-
11:00	0.556888	0.491138	0.555596	0.769417
13:00	-	-	-	-
15:00	0.681092	0.685839	0.611873	0.745593
17:00	0.479863	0.387895	1.16038	0.880193
19:00	0.915759	0.683234	1.18074	0.850216
21:00	0.310155	0.590418	0.868966	1.27372
23:00	0.784225	1.344312	1.203302	1.060018

Bibliography

- ÁLVAREZ, E., LÓPEZ-URRUTIA, Á., NOGUEIRA, E. & FRAGA, S. 2011. How to effectively sample the plankton size spectrum? A case study using FlowCAM. *Journal of Plankton Research*, 33, 1119-1133.
- ANDERSON, D. M., GLIBERT, P. M. & BURKHOLDER, J. M. 2002. Harmful algal blooms and eutrophication: Nutrient sources, composition, and consequences. *Estuaries*, 25, 704-726.
- ANDERSON, D. M., HOAGLAND, P., KAORU, Y. & WHITE, A. W. 2000. Estimated annual economic impacts from harmful algal blooms (HABs) in the United States. National Oceanic and Atmospheric Administration Norman OK National Severe
- ANGLÈS, S., JORDI, A. & CAMPBELL, L. 2015. Responses of the coastal phytoplankton community to tropical cyclones revealed by high-frequency imaging flow cytometry. *Limnology and Oceanography*, 60, 1562-1576.
- ANGLÈS, S., JORDI, A., HENRICH, D. W. & CAMPBELL, L. 2019. Influence of coastal upwelling and river discharge on the phytoplankton community composition in the northwestern Gulf of Mexico. *Progress in Oceanography*, 173, 26-36.
- ANTFOLK, M. & LAURELL, T. 2019. Acoustofluidic Blood Component Sample Preparation and Processing in Medical Applications. *Applications of Microfluidic Systems in Biology and Medicine*. Springer.
- ASCH, R. G., STOCK, C. A. & SARMIENTO, J. L. 2019. Climate change impacts on mismatches between phytoplankton blooms and fish spawning phenology. *Glob Chang Biol*, 25, 2544-2559.
- BACH, J. S. & BRUUS, H. 2020. Theory of acoustic trapping of microparticles in capillary tubes. *Phys Rev E*, 101, 023107.
- BAKIRMAN, T. 2023. An Assessment of YOLO Architectures for Oil Tank Detection from SPOT Imagery. *International Journal of Environment and Geoinformatics*, 10, 9-15.
- BARRY, M. T., RUSCONI, R., GUASTO, J. S. & STOCKER, R. 2015. Shear-induced orientational dynamics and spatial heterogeneity in suspensions of motile phytoplankton. *Journal of The Royal Society Interface*, 12, 20150791.
- BATTEN, S. D., CLARK, R., FLINKMAN, J., HAYS, G., JOHN, E., JOHN, A. W. G., JONAS, T., LINDLEY, J. A., STEVENS, D. P. & WALNE, A. 2003. CPR sampling: the technical background, materials and methods, consistency and comparability. *Progress in Oceanography*, 58, 193-215.
- BAUTISTA-CHAMIZO, E., SENDRA, M., CID, A., SEOANE, M., ROMANO DE ORTE, M. & RIBA, I. 2018. Will temperature and salinity changes exacerbate the effects of seawater acidification on the marine microalga *Phaeodactylum tricornutum*? *Sci Total Environ*, 634, 87-94.
- BEHRENFELD, M. J. 2010. Abandoning Sverdrup's Critical Depth Hypothesis on phytoplankton blooms. *Ecology*, 91, 977-989.
- BENFIELD, M. C., GROSJEAN, P., CULVERHOUSE, P. F., IRIGOIEN, X., SIERACKI, M. E., LOPEZ-URRUTIA, A., DAM, H. G., HU, Q., DAVIS, C. S., HANSEN, A., PILSKALN, C. H., RISEMAN, E. M., SCHULTZ, H., UTGOFF, P. E. & GORSKY, G. 2007. RAPID Research on Automated Plankton Identification. *Oceanography*, 20, 172-187.
- BHATT, P. M., MALHAN, R. K., RAJENDRAN, P., SHAH, B. C., THAKAR, S., YOON, Y. J. & GUPTA, S. K. 2021. Image-Based Surface Defect Detection Using Deep Learning: A Review. *Journal of Computing and Information Science in Engineering*, 21.
- BHATTACHARYA, D., MEDLIN & LINDA 1998. Algal Phylogeny and the Origin of Land Plants. *Plant Physiology*, 116, 9-15.
- BIEDER, F., SANDKUEHLER, R. & CATTIN, P. 2021. *Comparison of Methods Generalizing Max- and Average-Pooling*.
- BISHOP, C. M. 1994. Neural networks and their applications. *Review of scientific instruments*, 65, 1803-1832.

- BLACKBURN, N., HAECKY, P., JURGENSON, I., GRINIENE, E., BRUGEL, S., ANDERSSON, A. & CARSTENSEN, J. 2022. The use of an automated organism tracking microscope in mesocosm experiments. *Limnology and Oceanography: Methods*, 20, 768-780.
- BLASCHKO, M. B., HOLNESS, G., MATTAR, M. A., LISIN, D., UTGOFF, P. E., HANSON, A. R., SCHULTZ, H., RISEMAN, E. M., SIERACKI, M. E., BALCH, W. M. & TUPPER, B. Automatic In Situ Identification of Plankton. 2005 Seventh IEEE Workshops on Applications of Computer Vision (WACV/MOTION'05) - Volume 1, 5-7 Jan. 2005 2005. 79-86.
- BOLANOS, L. M., KARP-BOSS, L., CHOI, C. J., WORDEN, A. Z., GRAFF, J. R., HAENTJENS, N., CHASE, A. P., DELLA PENNA, A., GAUBE, P., MORISON, F., MENDEN-DEUER, S., WESTBERRY, T. K., O'MALLEY, R. T., BOSS, E., BEHRENFELD, M. J. & GIOVANNONI, S. J. 2020. Small phytoplankton dominate western North Atlantic biomass. *ISME J*, 14, 1663-1674.
- BORN, M. & WOLF, E. 1999. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*, Cambridge, Cambridge University Press.
- BRIDLE, J. 1989. Training stochastic model recognition algorithms as networks can lead to maximum mutual information estimation of parameters. *Advances in neural information processing systems*, 2.
- BROTAS, V., BREWIN, R. J. W., SÁ, C., BRITO, A. C., SILVA, A., MENDES, C. R., DINIZ, T., KAUFMANN, M., TARRAN, G., GROOM, S. B., PLATT, T. & SATHYENDRANATH, S. 2013. Deriving phytoplankton size classes from satellite data: Validation along a trophic gradient in the eastern Atlantic Ocean. *Remote Sensing of Environment*, 134, 66-77.
- BROWNLEE, E. F., OLSON, R. J. & SOSIK, H. M. 2016. Microzooplankton community structure investigated with imaging flow cytometry and automated live-cell staining. *Marine Ecology Progress Series*, 550, 65-81.
- BRUUS, H. 2012. Acoustofluidics 7: The acoustic radiation force on small particles. *Lab on a Chip*, 12, 1014-1021.
- BURKHOLDER, J. M. & GLASGOW JR., H. B. 1997. Pfiesteria piscicida and other Pfiesteria-like dinoflagellates: Behavior, impacts, and environmental controls. *Limnology and Oceanography*, 42, 1052-1075.
- BURKHOLDER, J. M., GLIBERT, P. M. & SKELTON, H. M. 2008. Mixotrophy, a major mode of nutrition for harmful algal species in eutrophic waters. *Harmful Algae*, 8, 77-93.
- BURNS, P. A., IS, IS, IS & IS 2000. *Slanted-edge MTF for digital camera and scanner analysis*.
- BUSKEY, E. J. & HYATT, C. J. 2006. Use of the FlowCAM for semi-automated recognition and enumeration of red tide cells (*Karenia brevis*) in natural plankton samples. *Harmful Algae*, 5, 685-692.
- BUSSI, G., WHITEHEAD, P. G., BOWES, M. J., READ, D. S., PRUDHOMME, C. & DADSON, S. J. 2016. Impacts of climate change, land-use change and phosphorus reduction on phytoplankton in the River Thames (UK). *Sci Total Environ*, 572, 1507-1519.
- CAMPBELL, L., HENRICH, D. W., OLSON, R. J. & SOSIK, H. M. 2013. Continuous automated imaging-in-flow cytometry for detection and early warning of *Karenia brevis* blooms in the Gulf of Mexico. *Environmental Science and Pollution Research*, 20, 6896-6902.
- CAMPBELL, L., HENRICH, D. W., PEACOCK, E. E., FUTRELLE, J. & SOSIK, H. M. 2016. Imaging FlowCytobot provides novel insights on phytoplankton community dynamics. *MARINE AND FRESH-WATER HARMFUL ALGAE*, 74.
- CAMPBELL, L., OLSON, R. J., SOSIK, H. M., ABRAHAM, A., HENRICH, D. W., HYATT, C. J. & BUSKEY, E. J. 2010. First Harmful Dinophysis (Dinophyceae, Dinophysiales) Bloom in the U.S. is Revealed by Automated Imaging Flow Cytometry. *Journal of Phycology*, 46.
- CANNY, J. 1986. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8, 679-698.
- CARMICHAEL, W. W. 2001. Health Effects of Toxin-Producing Cyanobacteria: "The CyanoHABS". *Human and Ecological Risk Assessment: An International Journal*, 7, 1393-1407.
- CAVALIER-SMITH, T. 1993. Kingdom protozoa and its 18 phyla. *Microbiol Rev*, 57, 953-94.
- CETINIC, I., PERRY, M., D'ASARO, E., BRIGGS, N., POULTON, N., SIERACKI, M. & LEE, C. 2014. Optical community index to assess spatial patchiness during the 2008 North Atlantic Bloom. *Biogeosciences Discussion*, 11, 12833-12870.

- CHOU, W.-L. 2022. *YoloX* [Online]. Available: <https://hackmd.io/@willy541222/YOLOX> [Accessed].
- CLOERN, J. E. 2018. Why large cells dominate estuarine phytoplankton. *Limnology and Oceanography*, 63, S392-S409.
- COLAS, F., TARDIVEL, M., EVRARD, J., FOREST, B., CRASSOUS, M. P., LUNVEN, M. & DANIELOU, M. M. 2016. Progression of the digitalization system: the FastCAM prototype and its outlook. *Optimisation de l'identification et du dénombrement du micro-phytoplancton avec le système couplé de numérisation et d'analyse d'images FlowCAM – Zoo/PhytoImage (système innovant)* IFREMER.
- COLLIER, J. L. 2000. Flow cytometry and the single cell in phycology. *Journal of Phycology*, 36, 628-644.
- COLLINS, S., ROST, B. & RYNEARSON, T. A. 2014. Evolutionary potential of marine phytoplankton under ocean acidification. *Evol Appl*, 7, 140-55.
- CULVERHOUSE, P. F., WILLIAMS, R., BENFIELD, M., FLOOD, P. R., SELL, A., F., MAZZOCCHI, M. G., BUTTINO, I. & SIERACKI, M. 2006. Automatic image analysis of plankton: future perspectives. *Marine Ecology Progress Series*, 312, 297-309.
- DAPENA, C., BRAVO, I., CUADRADO, A. & FIGUEROA, R. I. 2015. Nuclear and cell morphological changes during the cell cycle and growth of the toxic dinoflagellate *Alexandrium minutum*. *Protist*, 166, 146-60.
- DASHKOVA, V., MALASHENKOV, D., POULTON, N., VOROBEV, I. & BARTENEVA, N. S. 2017. Imaging flow cytometry for phytoplankton analysis. *Methods*, 112, 188-200.
- DELONG, E. F. 2009. The microbial ocean from genomes to biomes. *Nature*, 459, 200-206.
- DENG, J., DONG, W., SOCHER, R., LI, L.-J., LI, K. & FEI-FEI, L. Imagenet: A large-scale hierarchical image database. 2009 IEEE conference on computer vision and pattern recognition, 2009. Ieee, 248-255.
- DI CARLO, D., IRIMIA, D., TOMPKINS, R. G. & TONER, M. 2007. Continuous inertial focusing, ordering, and separation of particles in microchannels. *Proceedings of the National Academy of Sciences*, 104, 18892-18897.
- DOBBINS, J. T. 2000. Chapter 3: Image Quality Metrics for Digital Systems *In*: KUNDEL, R. L. V. M. J. B. H. L. (ed.) *Handbook of Medical Imaging, Volume 1. Physics and Psychophysics*. SPIE Press.
- DOĞAN, N. Ö. 2018. Bland-Altman analysis: A paradigm to understand correlation and agreement. *Turkish Journal of Emergency Medicine*, 18, 139-141.
- DOLLAR, P., WOJEK, C., SCHIELE, B. & PERONA, P. 2011. Pedestrian detection: An evaluation of the state of the art. *IEEE transactions on pattern analysis and machine intelligence*, 34, 743-761.
- DOUGLAS, D. H. & PEUCKER, T. K. 1973. Algorithms for the Reduction of the Number of Points Required to Represent a Digitized Line or Its Caricature. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 10, 112-122.
- DUBELAAR, G. B., GERRITZEN, P. L., BEEKER, A. E., JONKER, R. R. & TANGEN, K. 1999. Design and first results of CytoBuoy: a wireless flow cytometer for in situ analysis of marine and fresh waters. *Cytometry*, 37, 247-54.
- DUBELAAR, G. B. J. & GERRITZEN, P. L. 2000. CytoBuoy: a step forward towards using flow cytometry in operational oceanography. *2000*, 64, 11.
- DUDA, R. O. & HART, P. E. Pattern classification and scene analysis. A Wiley-Interscience publication, 1974.
- DUNKER, S., BOHO, D., WALDCHEN, J. & MADER, P. 2018. Combining high-throughput imaging flow cytometry and deep learning for efficient species and life-cycle stage identification of phytoplankton. *BMC Ecol*, 18, 51.
- ELSWORTH, G. W., LOVENDUSKI, N. S., MCKINNON, K. A., KRUMHARDT, K. M. & BRADY, R. X. 2020. Finding the Fingerprint of Anthropogenic Climate Change in Marine Phytoplankton Abundance. *Current Climate Change Reports*, 6, 37-46.
- EPPLEY, R. W., HOLM-HARISEN, O. & STRICKLAND, J. D. 1968. Some Observations on the Vertical Migration of dinoflagellates. *J Phycol*, 4, 333-40.

- ESTRIBEAU, M. & MAGNAN, P. 2004. *Fast MTF measurement of CMOS imagers using ISO 12333 slanted-edge methodology*, SPIE.
- EVANDER, M., JOHANSSON, L., LILLIEHORN, T., PISKUR, J., LINDVALL, M., JOHANSSON, S., ALMQVIST, M., LAURELL, T. & NILSSON, J. 2007. Noninvasive acoustic cell trapping in a microfluidic perfusion system for online bioassays. *Analytical chemistry*, 79, 2984-2991.
- FALKOWSKI, P. G., BARBER, R. T. & SMETACEK, V. 1998. Biogeochemical Controls and Feedbacks on Ocean Primary Production. *Science*, 281, 200-206.
- FENCHEL, T. & JUEL HANSEN, P. 2006. Motile behaviour of the bloom-forming ciliate *Mesodinium rubrum*. *Marine Biology Research*, 2, 33-40.
- FIGLEY, W., PYLE, B. & HALGREN, B. 1979. Chapter 14: Socioeconomic Impacts. In: SINDERMAN, C. J. & SWANSON, R. L. (eds.) *Oxygen depletion and associated benthic mortalities in New York Bight, 1976*. Rockville, Md.: National Oceanic and Atmospheric Administration.
- FLEMINGER, A. & CLUTTER, R. I. 1965. Avoidance of Towed Nets by Zooplankton. *Limnology and Oceanography*, 10, 96-104.
- FOFONOFF, N. P. & MILLARD JR, R. 1983. Algorithms for the computation of fundamental properties of seawater.
- FOGG, G. E. 1990. Our perceptions of phytoplankton: an historical sketch the first Founders' Lecture. *British Phycological Journal*, 25, 103-115.
- FORNELL, A., JOHANNESSEN, C., SEARLE, S. S., HAPPSTADIUS, A., NILSSON, J. & TENJE, M. 2019. An acoustofluidic platform for non-contact trapping of cell-laden hydrogel droplets compatible with optical microscopy. *Biomicrofluidics*, 13, 044101.
- GARETH JAMES, DANIELA WITTEN, TREVOR HASTIE & TIBSHIRANI, R. 2013. *An introduction to statistical learning : with applications in R*, New York : Springer, [2013] ©2013.
- GAUTAM, N., SANKARAN, S., YASON, J. A., TAN, K. S. W. & GASCOIGNE, N. R. J. 2018. A high content imaging flow cytometry approach to study mitochondria in T cells: MitoTracker Green FM dye concentration optimization. *Methods*, 134-135, 11-19.
- GE, Z., LIU, S., WANG, F., LI, Z. & SUN, J. 2021. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*.
- GEORGE, T. C., BASIJI, D. A., HALL, B. E., LYNCH, D. H., ORTYN, W. E., PERRY, D. J., SEO, M. J., ZIMMERMAN, C. A. & MORRISSEY, P. J. 2004. Distinguishing modes of cell death using the ImageStream® multispectral imaging flow cytometer. *Cytometry Part A*, 59A, 237-245.
- GIRSHICK, R., DONAHUE, J., DARRELL, T. & MALIK, J. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. 2014 IEEE Conference on Computer Vision and Pattern Recognition, 23-28 June 2014 2014. 580-587.
- GLIBERT, P. M., ICARUS ALLEN, J., ARTIOLI, Y., BEUSEN, A., BOUWMAN, L., HARLE, J., HOLMES, R. & HOLT, J. 2014. Vulnerability of coastal ecosystems to changes in harmful algal bloom distribution in response to climate change: projections based on model analysis. *Global Change Biology*, 20, 3845-3858.
- GLYNNE-JONES, P., BOLTRYK, R. J., HARRIS, N. R., CRANNY, A. W. & HILL, M. 2010a. Mode-switching: a new technique for electronically varying the agglomeration position in an acoustic particle manipulator. *Ultrasonics*, 50, 68-75.
- GLYNNE-JONES, P., BOLTRYK, R. J., HILL, M., ZHANG, F., DONG, L., WILKINSON, J. S., BROWN, T., MELVIN, T. & HARRIS, N. R. 2010b. Multi-modal particle manipulator to enhance bead-based bioassays. *Ultrasonics*, 50, 235-239.
- GOBLER, C. J., DOHERTY, O. M., HATTENRATH-LEHMANN, T. K., GRIFFITH, A. W., KANG, Y. & LITAKER, R. W. 2017. Ocean warming since 1982 has expanded the niche of toxic algal blooms in the North Atlantic and North Pacific oceans. *Proc Natl Acad Sci U S A*, 114, 4975-4980.
- GODDARD, G. R., SANDERS, C. K., MARTIN, J. C., KADUCHAK, G. & GRAVES, S. W. 2007. Analytical Performance of an Ultrasonic Particle Focusing Flow Cytometer. *Analytical Chemistry*, 79, 8740-8746.
- GREENBAUM, A., AKBARI, N., FEIZI, A., LUO, W. & OZCAN, A. 2013. Field-Portable Pixel Super-Resolution Colour Microscope. *Plos One*, 8.

- GUTIÉRREZ-RODRÍGUEZ, A., SELPH, K. E. & LANDRY, M. R. 2016. Phytoplankton growth and microzooplankton grazing dynamics across vertical environmental gradients determined by transplant in situ dilution experiments. *Journal of Plankton Research*, 38, 271-289.
- GÖRÖCS, Z., TAMAMITSU, M., BIANCO, V., WOLF, P., ROY, S., SHINDO, K., YANNY, K., WU, Y., KOYDEMIR, H. C., RIVENSON, Y. & OZCAN, A. 2018. A deep learning-enabled portable imaging flow cytometer for cost-effective, high-throughput, and label-free analysis of natural water samples. *Light Sci Appl*, 7, 66.
- HÄDER, D.-P. 1995. Novel method to determine vertical distributions of phytoplankton in marine water columns. *Environmental and Experimental Botany*, 35, 547-555.
- HAIN, R., KÄHLER, C. J. & TROPEA, C. 2007. Comparison of CCD, CMOS and intensified cameras. *Experiments in Fluids*, 42, 403-411.
- HALLEGRAEFF, G. M. 1993. A review of harmful algal blooms and their apparent global increase. *Phycologia*, 32, 79-99.
- HALLEGRAEFF, G. M. 2010. OCEAN CLIMATE CHANGE, PHYTOPLANKTON COMMUNITY RESPONSES, AND HARMFUL ALGAL BLOOMS: A FORMIDABLE PREDICTIVE CHALLENGE1. *Journal of Phycology*, 46, 220-235.
- HAMMER, A., SCHUMANN, R. & SCHUBERT, H. 2002. Light and temperature acclimation of *Rhodomonas salina* (Cryptophyceae): Photosynthetic performance. *Aquatic Microbial Ecology - AQUAT MICROB ECOL*, 29, 287-296.
- HAN, Y., GU, Y., ZHANG, A. C. & LO, Y. H. 2016. Review: imaging technologies for flow cytometry. *Lab Chip*, 16, 4639-4647.
- HARAGUCHI, L., JAKOBSEN, H. H., LUNDHOLM, N. & CARSTENSEN, J. 2017. Monitoring natural phytoplankton communities: a comparison between traditional methods and pulse-shape recording flow cytometry. *Aquatic Microbial Ecology*, 80, 77-92.
- HARDY, A. C. 1939. Ecological investigations with the Continuous Plankton Recorder: Object, plan and methods. *Hull Bulletins of Marine Ecology*, 1, 1-57.
- HASINOFF, S. W. 2014. Saturation (imaging). In: IKEUCHI, K. (ed.) *Computer Vision: A Reference Guide*. Springer US.
- HAYES, N. M., VANNI, M. J., HORGAN, M. J. & RENWICK, W. H. 2015. Climate and land use interactively affect lake phytoplankton nutrient limitation status. *Ecology*, 96, 392-402.
- HE, J., ERFANI, S., MA, X., BAILEY, J., CHI, Y. & HUA, X.-S. 2021. Alpha-IoU: A Family of Power Intersection over Union Losses for Bounding Box Regression. *Advances in Neural Information Processing Systems*, 34, 20230-20242.
- HEADLAND, S. E., JONES, H. R., D'SA, A. S. V., PERRETTI, M. & NORLING, L. V. 2014. Cutting-Edge Analysis of Extracellular Microparticles using ImageStream(X) Imaging Flow Cytometry. *Scientific Reports*, 4, 10.
- HENSON, S. A., COLE, H. S., HOPKINS, J., MARTIN, A. P. & YOOL, A. 2018. Detection of climate change-driven trends in phytoplankton phenology. *Glob Chang Biol*, 24, e101-e111.
- HENSON, S. A., SARMIENTO, J. L., DUNNE, J. P., BOPP, L., LIMA, I., DONEY, S. C., JOHN, J. & BEAULIEU, C. 2010. Detection of anthropogenic climate change in satellite records of ocean chlorophyll and productivity. *Biogeosciences*, 7, 621-640.
- HEO, Y. J., LEE, D., KANG, J., LEE, K. & CHUNG, W. K. 2017. Real-time Image Processing for Microscopy-based Label-free Imaging Flow Cytometry in a Microfluidic Chip. *Sci Rep*, 7, 11651.
- HESS, D., RANE, A., DEMELLO, A. J. & STAVRAKIS, S. 2015. High-throughput, quantitative enzyme kinetic analysis in microdroplets using stroboscopic epifluorescence imaging. *Anal Chem*, 87, 4965-72.
- HILL, M., GLYNNE-JONES, P., HARRIS, N. R., BOLTRYK, R. J., STANLEY, C. & BOND, D. Trapping and micromanipulation using ultrasonic fields and dual ultrasonic/magnetic forces. Proc.SPIE, 2010.
- HILL, M., SHEN, Y. & HAWKES, J. J. 2002. Modelling of layered resonators for ultrasonic separation. *Ultrasonics*, 40, 385-392.

- HINCAPIÉ GÓMEZ, E., TRYNER, J., ALIGATA, A. J., QUINN, J. C. & MARCHESE, A. J. 2018. Measurement of acoustic properties of microalgae and implications for the performance of ultrasonic harvesting systems. *Algal Research*, 31, 77-86.
- HOLZNER, G., DU, Y., CAO, X., CHOO, J., A, J. D. & STAVRAKIS, S. 2018. An optofluidic system with integrated microlens arrays for parallel imaging flow cytometry. *Lab Chip*, 18, 3631-3637.
- HU, Y., LIU, Y. & LIU, Z. A Survey on Convolutional Neural Network Accelerators: GPU, FPGA and ASIC. 2022 14th International Conference on Computer Research and Development (ICCRD), 7-9 Jan. 2022 2022. 100-107.
- HUANG, T. J. 2019. Acoustofluidics: Merging acoustics and microfluidics for biomedical applications. *The Journal of the Acoustical Society of America*, 145, 1786-1786.
- HUSSAIN, M., CHEN, D., CHENG, A., WEI, H. & STANLEY, D. 2013. Change detection from remotely sensed images: From pixel-based to object-based approaches. *ISPRS Journal of Photogrammetry and Remote Sensing*, 80, 91-106.
- IOFFE, S. & SZEGEDY, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. International conference on machine learning, 2015. pmlr, 448-456.
- ISO 2017. Photography - Electronic still picture imaging - Resolution and spatial frequency responses. *ISO 12233:2017* International Organization for Standardization.
- JAKOBSEN, H. 2001. Escape response of planktonic protists to fluid mechanical signals. *Marine Ecology-progress Series - MAR ECOL-PROGR SER*, 214, 67-78.
- JAKOBSEN, H. H., EVERETT, L. M. & STROM, S. L. 2006. Hydromechanical signaling between the ciliate *Mesodinium pulex* and motile protist prey. *Aquatic Microbial Ecology*, 44, 197-206.
- JAYASINGHE, S. N. 2020. Reimagining Flow Cytometric Cell Sorting. *Adv Biosyst*, 4, e2000019.
- JÖHNK, K. D., HUISMAN, J. E. F., SHARPLES, J., SOMMEIJER, B. E. N., VISSER, P. M. & STROOM, J. M. 2008. Summer heatwaves promote blooms of harmful cyanobacteria. *Global Change Biology*, 14, 495-512.
- KAMYKOWSKI, D., PRIDGEN, K. G., MORRISON, J. M., MCCULLOCH, A. A., NYADJRO, E. S., THOMAS, C. A. & SINCLAIR, G. A. 2013. Cold front induced changes on the Florida panhandle shelf during October 2008. *Continental Shelf Research*, 54, 52-66.
- KAŇA, R., KOTABOVÁ, E., ŠEDIVÁ, B. & KUTHANOVÁ TRSKOVÁ, E. 2019. Photoprotective strategies in the motile cryptophyte alga *Rhodomonas salina*—role of non-photochemical quenching, ions, photoinhibition, and cell motility. *Folia Microbiologica*, 64, 691-703.
- KARAMAN, A., PACAL, I., BASTURK, A., AKAY, B., NALBANTOGLU, U., COSKUN, S., SAHIN, O. & KARABOGA, D. 2023. Robust real-time polyp detection system design based on YOLO algorithms by optimizing activation functions and hyper-parameters with artificial bee colony (ABC). *Expert Systems with Applications*, 221, 119741.
- KARLSON, B., GODHE, A., CUSACK, C. & BRESNAN, E. 2010. Introduction to methods for quantitative phytoplankton analysis. *Microscopic and molecular methods for quantitative phytoplankton analysis*, 5.
- KASS, M., WITKIN, A. & TERZOPOULOS, D. 1988. Snakes: Active contour models. *International Journal of Computer Vision*, 1, 321-331.
- KATSIAPI, M., MOUSTAKA-GOUNI, M., MICHALOUDI, E. & KORMAS, K. A. 2011. Phytoplankton and water quality in a Mediterranean drinking-water reservoir (Marathonas Reservoir, Greece). *Environ Monit Assess*, 181, 563-75.
- KHEIREDDINE, S., SUDALAIYADUM PERUMAL, A., SMITH, Z. J., NICOLAU, D. V. & WACHSMANN-HOGIU, S. 2019. Dual-phone illumination-imaging system for high resolution and large field of view multi-modal microscopy. *Lab on a Chip*, 19, 825-836.
- KIM, H. 2010. An Overview on the Occurrences of Harmful Algal Blooms (HABs) and Mitigation Strategies in Korean Coastal Waters. *Coastal Environmental and Ecosystem Issues of the East China Sea*.
- KIM, M., BAYLY, P. V. & MEACHAM, J. M. 2021. Motile cells as probes for characterizing acoustofluidic devices. *Lab on a Chip*, 21, 521-533.
- KINGMA, D. P. & BA, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

- KINSLER, L. E., FREY, A. R., COPPENS, A. B. & SANDERS, J. V. 1999. *Fundamentals of acoustics*.
- KÖHLER, A. 1893. Zeitschrift für wissenschaftliche. *Mikroskopie X*, 433–440
- KOUAKOU, C. R. C. & PODER, T. G. 2019. Economic impact of harmful algal blooms on human health: a systematic review. *J Water Health*, 17, 499-516.
- KRAFT, K., VELHONOJA, O., EEROLA, T., SUIKKANEN, S., TAMMINEN, T., HARAGUCHI, L., YLÖSTALO, P., KIELOSTO, S., JOHANSSON, M., LENSU, L., KÄLVIÄINEN, H., HAARIO, H. & SEPPÄLÄ, J. 2022. Towards operational phytoplankton recognition with automated high-throughput imaging, near-real-time data processing, and convolutional neural networks. *Frontiers in Marine Science*, 9.
- KREMER, C. T., THOMAS, M. K. & LITCHMAN, E. 2017. Temperature- and size-scaling of phytoplankton population growth rates: Reconciling the Eppley curve and the metabolic theory of ecology. *Limnology and Oceanography*, 62, 1658-1670.
- KRIMHOLTZ, R., LEEDOM, D. A. & MATTHAEI, G. L. 1970. New equivalent circuits for elementary piezoelectric transducers. *Electronics Letters*, 6, 398-399.
- KRIZHEVSKY, A., SUTSKEVER, I. & HINTON, G. E. 2017. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60, 84-90.
- KUDELA, R. M., LANE, J. Q. & COCHLAN, W. P. 2008. The potential role of anthropogenically derived nitrogen in the growth of harmful algae in California, USA. *Harmful Algae*, 8, 103-110.
- LANEY, S. R. & SOSIK, H. M. 2014. Phytoplankton assemblage structure in and around a massive under-ice bloom in the Chukchi Sea. *Deep Sea Research Part II: Topical Studies in Oceanography*, 105, 30-41.
- LATSOS, C., BAKRATSAS, G., MOERDIJK, T., VAN HOUCKE, J. & TIMMERMANS, K. R. 2021. Effect of salinity and pH on growth, phycoerythrin, and non-volatile umami taste active compound concentration of *Rhodomonas salina* using a D-optimal design approach. *Journal of Applied Phycology*, 33, 3591-3602.
- LEGENDRE, L., COURTIES, C. & TROUSSELLIER, M. 2001. Flow cytometry in oceanography 1989–1999: Environmental challenges and research trends. *Cytometry*, 44, 164-172.
- LEI, C., ITO, T., UGAWA, M., NOZAWA, T., IWATA, O., MAKI, M., OKADA, G., KOBAYASHI, H., SUN, X. L., TIAMSAK, P., TSUMURA, N., SUZUKI, K., DI CARLO, D., OZEKI, Y. & GODA, K. 2016. High-throughput label-free image cytometry and image-based classification of live *Euglena gracilis*. *Biomedical Optics Express*, 7, 2703-2708.
- LEIBACHER, I., HAHN, P. & DUAL, J. 2015. Acoustophoretic cell and particle trapping on microfluidic sharp edges. *Microfluidics and Nanofluidics*, 19, 923-933.
- LENSHOF, A., MAGNUSSON, C. & LAURELL, T. 2012. Acoustofluidics 8: Applications of acoustophoresis in continuous flow microsystems. *Lab on a Chip*, 12, 1210-1223.
- LEVASSEUR, M., THOMPSON, P. A. & HARRISON, P. J. 1993. Physiological Acclimation of Marine Phytoplankton to Different Nitrogen Sources. *Journal of Phycology*, 29, 587-595.
- LEVINSEN, H., NIELSEN, T. G. & HANSEN, B. W. 1999. Plankton community structure and carbon cycling on the western coast of Greenland during the stratified summer situation. II. Heterotrophic dinoflagellates and ciliates. *Aquatic Microbial Ecology*, 16, 217-232.
- LI, Q., SUN, X., DONG, J., SONG, S., ZHANG, T., LIU, D., ZHANG, H., HAN, S. & BEYAN, C. 2020. Developing a microscopic image dataset in support of intelligent phytoplankton detection using deep learning. *ICES Journal of Marine Science*, 77, 1427-1439.
- LI, S., GLYNNE-JONES, P., ANDRIOTIS, O. G., CHING, K. Y., JONNALAGADDA, U. S., OREFFO, R. O., HILL, M. & TARE, R. S. 2014. Application of an acoustofluidic perfusion bioreactor for cartilage tissue engineering. *Lab on a Chip*, 14, 4475-4485.
- LIANG, J., CHEN, X., LIANG, C., LONG, T., TANG, X., SHI, Z., ZHOU, M., ZHAO, J., LAN, Y. & LONG, Y. 2023. A detection approach for late-autumn shoots of litchi based on unmanned aerial vehicle (UAV) remote sensing. *Computers and Electronics in Agriculture*, 204, 107535.
- LIM, H. C., LEAW, C. P., SU, S. N., TENG, S. T., USUP, G., MOHAMMAD-NOOR, N., LUNDHOLM, N., KOTAKI, Y. & LIM, P. T. 2012. Morphology and Molecular Characterization of Pseudo-Nitzschia (Bacillariophyceae) from Malaysian Borneo, Including the New Species Pseudo-Nitzschia Circumpora Sp. Nov. *J Phycol*, 48, 1232-47.

- LIN, W., HASENSTAB, K., MOURA CUNHA, G. & SCHWARTZMAN, A. 2020. Comparison of handcrafted features and convolutional neural networks for liver MR image adequacy assessment. *Scientific Reports*, 10, 20336.
- LIU, S. & HUA, H. 2011. Extended depth-of-field microscopic imaging with a variable focus microscope objective. *Optics Express*, 19, 353-362.
- LIU, S., QI, L., QIN, H., SHI, J. & JIA, J. Path aggregation network for instance segmentation. Proceedings of the IEEE conference on computer vision and pattern recognition, 2018. 8759-8768.
- LOMBARD, F., BOSS, E., WAITE, A. M., VOGT, M., UITZ, J., STEMMANN, L., SOSIK, H. M., SCHULZ, J., ROMAGNAN, J. B., PICHERAL, M., PEARLMAN, J., OHMAN, M. D., NIEHOFF, B., MOLLER, K. M., MILOSLAVICH, P., LARA-LPEZ, A., KUDELA, R., LOPES, R. M., KIKO, R., KARP-BOSS, L., JAFFE, J. S., IVERSEN, M. H., FRISSON, J. O., FENNEL, K., HAUSS, H., GUIDI, L., GORSKY, G., GIERING, S. L. C., GAUBE, P., GALLAGER, S., DUBELAAR, G., COWEN, R. K., CARLOTTI, F., BRISENO-AVENA, C., BERLINE, L., BENOIT-BIRD, K., BAX, N., BATTEN, S., AYATA, S. D., ARTIGAS, L. F. & APPELTANS, W. 2019. Globally Consistent Quantitative Observations of Planktonic Ecosystems. *Frontiers in Marine Science*, 6.
- LOPEZ-RIQUELME, N., MINGUELA, A., VILLAR-PERMUY, F., CIPRIAN, D., CASTILLEJO, A., ALVAREZ-LOPEZ, M. R. & SOTO, J. L. 2013. Imaging cytometry for counting circulating tumor cells: comparative analysis of the CellSearch vs ImageStream systems. *Apmis*, 121, 1139-1143.
- LUND, R. D. 2016. Analyzing the Effect of Pipette Tip Geometries on Fluid Velocity and Shear Strain Rate: Biomek Wide Bore vs. Standard Pipette Tips. *Discovery In Motion*. Indianapolis, IN: Beckman Coulter Life Sciences.
- M. ART, N. CHANDELIER, U. GAST & MILLER, D. 2023. A "Shear" Mystery – Uncovered. *APPLICATION NOTE I No. 442*. Hamburg · Germany: Eppendorf SE.
- MADRID-WOLFF, J. & FORERO, M. 2019. *Simple and open 4f Koehler transmitted illumination system for low-cost microscopic imaging and teaching*.
- MAHSERECI, M., BALLE, L., LASSNER, C. & HENNIG, P. 2017. Early stopping without a validation set. *arXiv preprint arXiv:1703.09580*.
- MANN, H. B. & WHITNEY, D. R. 1947. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, 50-60.
- MARANON, E. 2015. Cell size as a key determinant of phytoplankton metabolism and community structure. *Ann Rev Mar Sci*, 7, 241-64.
- MARIE, D., RIGAUT-JALABERT, F. & VAULOT, D. 2014. An improved protocol for flow cytometry analysis of phytoplankton cultures and natural samples. *Cytometry A*, 85, 962-8.
- MARTIN BLAND, J. & ALTMAN, D. 1986. Statistical Methods for Assessing Agreement between Two Methods of Clinical Measurement. *The Lancet*, 327, 307-310.
- MATHWORKS. 2023. *Multilabel Image Classification Using Deep Learning* [Online]. MathWorks Available: <https://www.mathworks.com/help/deeplearning/ug/multilabel-image-classification-using-deep-learning.html> [Accessed 26 March 2023].
- MEDWIN, H. 1975. Speed of sound in water: A simple equation for realistic parameters. Available: <http://hdl.handle.net/10945/40176> [Accessed 2020-07].
- MENDEN-DEUER, S., MORISON, F., MONTALBANO, A. L., FRANZÈ, G., STROCK, J., RUBIN, E., MCNAIR, H., MOUW, C. & MARREC, P. 2020. Multi-Instrument Assessment of Phytoplankton Abundance and Cell Sizes in Mono-Specific Laboratory Cultures and Whole Plankton Community Composition in the North Atlantic. *Frontiers in Marine Science*, 7.
- MITJA, C., ESCOFET, J., TACHO, A. & REVUELTA, R. 2011. *Slanted Edge MTF* [Online]. Available: <https://imagej.nih.gov/ij/plugins/se-mtf/index.html> [Accessed 11th June 2019].
- MIURA, T., MIKAMI, H., ISOZAKI, A., ITO, T., OZEKI, Y. & GODA, K. 2018. On-chip light-sheet fluorescence imaging flow cytometry at a high flow speed of 1 m/s. *Biomed Opt Express*, 9, 3424-3433.
- MOLINA-NAVARRO, E., MARTINEZ-PEREZ, S., SASTRE-MERLIN, A., VERDUGO-ALTHOFER, M. & PADISAK, J. 2014. Phytoplankton and suitability of derived metrics for assessing the ecological status in a limno-reservoir, a Water Framework Directive nondefined type of Mediterranean waterbody. *Lake and Reservoir Management*, 30, 46-62.

- MOORE, J. K. & ABBOTT, M. R. 2000. Phytoplankton chlorophyll distributions and primary production in the Southern Ocean. *Journal of Geophysical Research-Oceans*, 105, 28709-28722.
- MORAN, X. A. G., LOPEZ-URRUTIA, A., CALVO-DIAZ, A. & LI, W. K. W. 2010. Increasing importance of small phytoplankton in a warmer ocean. *Global Change Biology*, 16, 1137-1144.
- MULLER-KARGER, F. & KUDELA, R. 2016. Essential Ocean Variables (EOV) for Biology and Ecosystems: Phytoplankton biomass and diversity. In: BATTEN, S. & CHECKLEY, D. (eds.) *Essential Ocean Variables*. IOC.
- NGUYEN, H.-V., BAE, J.-H., LEE, Y.-E., LEE, H.-S. & KWON, K.-R. 2022. Comparison of Pre-Trained YOLO Models on Steel Surface Defects Detector Based on Transfer Learning with GPU-Based Embedded Devices. *Sensors*, 22, 9926.
- NIEMI, A. 1973. Ecology of phytoplankton in the Tvarinne area, SW coast of Finland. I. Dynamics of hydrography, nutrients, chlorophyll a and phytoplankton. *Acta Botanica Fennica*.
- OLLI, K. 1999. Diel vertical migration of phytoplankton and heterotrophic flagellates in the Gulf of Riga. *Journal of Marine Systems*, 23, 145-163.
- OLSON, R. J., SHALAPYONOK, A., KALB, D. J., GRAVES, S. W. & SOSIK, H. M. 2017. Imaging FlowCytobot modified for high throughput by in-line acoustic focusing of sample particles. *Limnology and Oceanography-Methods*, 15, 867-874.
- OLSON, R. J., SHALAPYONOK, A. & SOSIK, H. M. 2003. An automated submersible flow cytometer for analyzing pico- and nanophytoplankton: FlowCytobot. *Deep Sea Research Part I: Oceanographic Research Papers*, 50, 301-315.
- OLSON, R. J. & SOSIK, H. M. 2007a. A submersible imaging-in-flow instrument to analyze nano- and microplankton: Imaging FlowCytobot. *Limnology and Oceanography-Methods*, 5, 195-203.
- OLSON, R. J. & SOSIK, H. M. 2007b. A submersible imaging-in-flow instrument to analyze nano and microplankton: Imaging FlowCytobot. *Limnology and Oceanography: Methods*, 5, 195-203.
- PAERL, H. W. & HUISMAN, J. 2009. Climate change: a catalyst for global expansion of harmful cyanobacterial blooms. *Environmental Microbiology Reports*, 1, 27-37.
- PAERL, H. W. & PAUL, V. J. 2012. Climate change: Links to global expansion of harmful cyanobacteria. *Water Research*, 46, 1349-1363.
- PAPAGEORGIOU, C. P., OREN, M. & POGGIO, T. A general framework for object detection. Sixth International Conference on Computer Vision (IEEE Cat. No. 98CH36271), 1998. IEEE, 555-562.
- PARK, J., LEE, H., PARK, C. Y., HASAN, S., HEO, T. Y. & LEE, W. H. 2019. Algal Morphological Identification in Watersheds for Drinking Water Supply Using Neural Architecture Search for Convolutional Neural Network. *Water*, 11, 1338.
- PASZTALENIC, A. & PONIEWOZIK, M. 2010. Phytoplankton based assessment of the ecological status of four shallow lakes (Eastern Poland) according to Water Framework Directive – a comparison of approaches. *Limnologica - Ecology and Management of Inland Waters*, 40, 251-259.
- POLLINA, T., LARSON, A. G., LOMBARD, F., LI, H., COLIN, S., DE VARGAS, C. & PRAKASH, M. 2020. PlanktonScope: Affordable modular imaging platform for citizen oceanography. *bioRxiv*, 2020.04.23.056978.
- POMATI, F., MATTHEWS, B., SEEHAUSEN, O. & IBELINGS, B. W. 2017. Eutrophication and climate warming alter spatial (depth) co-occurrence patterns of lake phytoplankton assemblages. *Hydrobiologia*, 787, 375-385.
- POSPÍŠIL, J., FLIEGEL, K. & KLÍMA, M. 2017. *Assessing resolution in live cell structured illumination microscopy*, SPIE.
- POULTON, N. J. 2016. FlowCam: Quantification and Classification of Phytoplankton by Imaging Flow Cytometry. In: BARTENEVA, N. S. & VOROBEV, I. A. (eds.) *Imaging Flow Cytometry: Methods and Protocols*. New York, NY: Springer New York.
- PRAKASH, A., MEDCOF, J. C. & TENNANT, A. D. 1971. *Paralytic Shellfish Poisoning in Eastern Canada*, Fisheries Research Board of Canada.

- RANZATO, M. A., BOUREAU, Y. L. & LECUN, Y. 2008. Sparse feature learning for deep belief networks. *Advances in Neural Information Processing Systems*, 20, 1185-1192.
- REDMON, J., DIVVALA, S., GIRSHICK, R. & FARHADI, A. You Only Look Once: Unified, Real-Time Object Detection. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 27-30 June 2016 2016. 779-788.
- REDMON, J. & FARHADI, A. 2017. *YOLO9000: Better, Faster, Stronger*.
- REDMON, J. & FARHADI, A. 2018. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.
- RICHARDS, T. A., JONES, M. D. M., LEONARD, G. & BASS, D. 2012. Marine Fungi: Their Ecology and Molecular Diversity. *Annual Review of Marine Science*, 4, 495-522.
- RICHARDSON, A. J. & SCHOEMAN, D. S. 2004. Climate Impact on Plankton Ecosystems in the Northeast Atlantic. *Science*, 305, 1609.
- RITZMAN, J., BRODBECK, A., BROSTROM, S., MCGREW, S., DREYER, S., KLINGER, T. & MOORE, S. K. 2018. Economic and sociocultural impacts of fisheries closures in two fishing-dependent communities following the massive 2015 U.S. West Coast harmful algal bloom. *Harmful Algae*, 80, 35-45.
- RODENACKER, K., HENSE, B., JÜTTING, U. & GAIS, P. 2006. Automatic analysis of aqueous specimens for phytoplankton structure recognition and population estimation. *Microscopy Research and Technique*, 69, 708-720.
- ROERDINK, J. & MEIJSTER, A. 2003. The Watershed Transform: Definitions, Algorithms and Parallelization Strategies. *Fundam Inf*, 41, 187-228.
- ROMERO-MARTINEZ, L., VAN SLOOTEN, C., NEBOT, E., ACEVEDO-MERINO, A. & PEPERZAK, L. 2017. Assessment of imaging-in-flow system (FlowCAM) for systematic ballast water management. *Sci Total Environ*, 603-604, 550-561.
- SAMIR, S., EMARY, E., EL-SAYED, K. & ONSI, H. 2020. Optimization of a Pre-Trained AlexNet Model for Detecting and Localizing Image Forgeries. *Information* [Online], 11.
- SANTHANAM, P., PACHIAPPAN, P. & BEGUM, A. 2019. Methods of Collection, Preservation and Taxonomic Identification of Marine Phytoplankton. In: SANTHANAM, P., BEGUM, A. & PACHIAPPAN, P. (eds.) *Basic and Applied Phytoplankton Biology*. Singapore: Springer Singapore.
- SCHMIDT, M., GLASSON, J., EMMELIN, L. & HELBRON, H. 2008. *Standards and thresholds for impact assessment*, Springer Science & Business Media.
- SCHNEIDER, C. A., RASBAND, W. S. & ELICEIRI, K. W. 2012. NIH Image to ImageJ: 25 years of image analysis. *Nature Methods*, 9, 671.
- SCHNEPF, E. & KÜHN, S. F. 2000. Food uptake and fine structure of *Cryothecomonas longipes* sp. nov., a marine nanoflagellate incertae sedis feeding phagotrophically on large diatoms. *Helgoland Marine Research*, 54, 18-32.
- SCHOPF, J. W. 1993. Microfossils of the Early Archean Apex Chert: New Evidence of the Antiquity of Life. *Science*, 260, 640-646.
- SEE, J. H., CAMPBELL, L., RICHARDSON, T. L., PINCKNEY, J. L., SHEN, R. & GUINASSO JR, N. L. 2005. COMBINING NEW TECHNOLOGIES FOR DETERMINATION OF PHYTOPLANKTON COMMUNITY STRUCTURE IN THE NORTHERN GULF OF MEXICO 1. *Journal of Phycology*, 41, 305-310.
- SHAHINFAR, S., MEEK, P. & FALZON, G. 2020. "How many images do I need?" Understanding how sample size per class affects deep learning model performance metrics for balanced designs in autonomous wildlife monitoring. *Ecological Informatics*, 57, 101085.
- SHARMA, S. & GULERIA, K. Deep Learning Models for Image Classification: Comparison and Applications. 2022 2nd International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE), 28-29 April 2022 2022. 1733-1738.
- SHERRIT, S., LEARY, S. P., DOLGIN, B. P. & BAR-COHEN, Y. Comparison of the Mason and KLM equivalent circuits for piezoelectric resonators in the thickness mode. 1999 IEEE Ultrasonics Symposium. Proceedings. International Symposium (Cat. No.99CH37027), 17-20 Oct. 1999 1999. 921-926 vol.2.

- SHI, K., ZHANG, Y., ZHANG, Y., LI, N., QIN, B., ZHU, G. & ZHOU, Y. 2019. Phenology of Phytoplankton Blooms in a Trophic Lake Observed from Long-Term MODIS Data. *Environ Sci Technol*, 53, 2324-2331.
- SHIN, H. C., ORTON, M., COLLINS, D. J., DORAN, S. & LEACH, M. O. 2016a. Chapter 7 - Organ Detection Using Deep Learning. In: ZHOU, S. K. (ed.) *Medical Image Recognition, Segmentation and Parsing*. Academic Press.
- SHIN, H. C., ROTH, H. R., GAO, M., LU, L., XU, Z., NOGUES, I., YAO, J., MOLLURA, D. & SUMMERS, R. M. 2016b. Deep Convolutional Neural Networks for Computer-Aided Detection: CNN Architectures, Dataset Characteristics and Transfer Learning. *IEEE Trans Med Imaging*, 35, 1285-98.
- SIEBURTH, J. M., SMETACEK, V. & LENZ, J. 1978. Pelagic ecosystem structure: Heterotrophic compartments of the plankton and their relationship to plankton size fractions 1. *Limnology and Oceanography*, 23, 1256-1263.
- SIERACKI, M., BENFIELD, M., HANSON, A., DAVIS, C., PILSKALN, C., CHECKLEY, D., SOSIK, H., ASHJIAN, C., CULVERHOUSE, P., COWEN, R., LOPES, R., BALCH, W. & IRIGOIEN, X. 2010. *Optical Plankton Imaging and Analysis Systems for Ocean Observation*.
- SINGARAJAH, K. V. 1969. Escape reactions of zooplankton: The avoidance of a pursuing siphon tube. *Journal of Experimental Marine Biology and Ecology*, 3, 171-178.
- SMAYDA, T. J. 1997. Harmful algal blooms: Their ecophysiology and general relevance to phytoplankton blooms in the sea. *Limnology and Oceanography*, 42, 1137-1153.
- SMETACEK, V. & CLOERN, J. E. 2008. On Phytoplankton Trends. *Science*, 319, 1346-1348.
- SOYDANER, D. 2020. A Comparison of Optimization Algorithms for Deep Learning. *International Journal of Pattern Recognition and Artificial Intelligence*, 34, 2052013.
- STAVRAKIS, S., HOLZNER, G., CHOO, J. & DEMELLO, A. 2019. High-throughput microfluidic imaging flow cytometry. *Curr Opin Biotechnol*, 55, 36-43.
- STRAUß, S. 2018. From Big Data to Deep Learning: A Leap Towards Strong AI or 'Intelligentia Obscura'? *Big Data and Cognitive Computing*, 2, 16.
- SUIKKANEN, S., LAAMANEN, M. & HUTTUNEN, M. 2007. Long-term changes in summer phytoplankton communities of the open northern Baltic Sea. *Estuarine, Coastal and Shelf Science*, 71, 580-592.
- SUN, J. & NING, X. R. 2005. Marine phytoplankton specific growth rate. *J. Adv Earth Sci.*, 20, 939-945.
- SUN, X., SHEN, F., BREWIN, R. J. W., LI, M. & ZHU, Q. 2022. Light absorption spectra of naturally mixed phytoplankton assemblages for retrieval of phytoplankton group composition in coastal oceans. *Limnology and Oceanography*, 67, 946-961.
- TERVEN, J. & CORDOVA-ESPARZA, D. 2023. A Comprehensive Review of YOLO: From YOLOv1 to YOLOv8 and Beyond. *arXiv preprint arXiv:2304.00501*.
- THOISEN, C., VU, M. T. T., CARRON-CABARET, T., JEPSEN, P. M., NIELSEN, S. L. & HANSEN, B. W. 2018. Small-scale experiments aimed at optimization of large-scale production of the microalga *Rhodomonas salina*. *Journal of Applied Phycology*, 30, 2193-2202.
- THYSSEN, M., TARRAN, G. A., ZUBKOV, M. V., HOLLAND, R. J., GRÉGORI, G., BURKILL, P. H. & DENIS, M. 2008. The emergence of automated high-frequency flow cytometry: revealing temporal and spatial phytoplankton variability. *Journal of Plankton Research*, 30, 333-343.
- TRALLER, J. C. & HILDEBRAND, M. 2013. High throughput imaging to the diatom *Cyclotella cryptica* demonstrates substantial cell-to-cell variability in the rate and extent of triacylglycerol accumulation. *Algal Research*, 2, 244-252.
- TROMBETTA, T., VIDUSSI, F., MAS, S., PARIN, D., SIMIER, M. & MOSTAJIR, B. 2019. Water temperature drives phytoplankton blooms in coastal waters. *PLoS One*, 14, e0214933.
- TSALOGLOU, M.-N. 2016. Microfluidics and in situ Sensors for Microalgae. In: TSALOGLOU, M.-N. (ed.) *Microalgae: Current Research and Applications*. Caister Academic Press.
- VAINRUB, A. 2008. Precise measurement of the resolution in light microscopy using Fourier transform. *Review of Scientific Instruments*, 79, 046112.
- VHENGANI, L., GRIFFITH, D. & LYSKO, M. 2012. *Effects of slant angle and illumination angle on MTF estimations*.

- VOELKEL, R., VOGLER, U., BICH, A., PERNET, P., WEIBLE, K. J., HORNING, M., ZOBEBIER, R., CULLMANN, E., STUERZEBECKER, L., HARZENDORF, T. & ZEITNER, U. D. 2010. Advanced mask aligner lithography: new illumination system. *Optics Express*, 18, 20968-20978.
- VONSHAK, A. 1985. Micro-Algae: Laboratory Growth Techniques and Outdoor Biomass Production. In: COOMBS, J., HALL, D. O., LONG, S. P. & SCURLOCK, J. M. O. (eds.) *Techniques in Bioproductivity and Photosynthesis*. Pergamon.
- WANG, J., TANG, C. & LI, J. Towards Real-time Analysis of Marine Phytoplankton Images Sampled at High Frame Rate by a YOLOX-based Object Detection Algorithm. OCEANS 2022 - Chennai, 21-24 Feb. 2022 2022. 1-9.
- WANG, Z. A., MOUSTAHFID, H., MUELLER, A. V., MICHEL, A. P. M., MOWLEM, M., GLAZER, B. T., MOONEY, T. A., MICHAELS, W., MCQUILLAN, J. S., ROBIDART, J. C., CHURCHILL, J., SOURISSEAU, M., DANIEL, A., SCHAAP, A., MONK, S., FRIEDMAN, K. & BREHMER, P. 2019. Advancing Observation of Ocean Biogeochemistry, Biology, and Ecosystems With Cost-Effective in situ Sensing Technologies. *Frontiers in Marine Science*, 6.
- WILLERT, C., STASICKI, B., KLINNER, J. & MOESSNER, S. 2010. Pulsed operation of high-power light emitting diodes for imaging flow velocimetry. *Measurement Science and Technology*, 21, 075402.
- WU, J. & CHAN, R. K. Y. 2013. A fast fluorescence imaging flow cytometer for phytoplankton analysis. *Optics Express*, 21, 23921-23926.
- WU, M., OZCELIK, A., RUFO, J., WANG, Z., FANG, R. & JUN HUANG, T. 2019a. Acoustofluidic separation of cells and particles. *Microsyst Nanoeng*, 5, 32.
- WU, Q., MERCHANT, F. & CASTLEMAN, K. R. 2008. *Microscope Image Processing*, San Diego, UNITED STATES, Elsevier Science & Technology.
- WU, Q. & ZHOU, Y. 2019. *Real-Time Object Detection Based on Unmanned Aerial Vehicle*.
- WU, Z., JIANG, H., ZHANG, L., YI, K., CUI, H., WANG, F., LIU, W., ZHAO, X., ZHOU, F. & GUO, S. 2019b. The acoustofluidic focusing and separation of rare tumor cells using transparent lithium niobate transducers. *Lab Chip*, 19, 3922-3930.
- XIE, W., NOBLE, J. A. & ZISSERMAN, A. 2018a. Microscopy cell counting and detection with fully convolutional regression networks. *Computer methods in biomechanics and biomedical engineering: Imaging & Visualization*, 6, 283-292.
- XIE, X., FAN, H., WANG, A., ZOU, N. & ZHANG, Y. 2018b. Regularized slanted-edge method for measuring the modulation transfer function of imaging systems. *Applied Optics*, 57, 6552-6558.
- XIE, X., FAN, H., WANG, H., WANG, Z. & ZOU, N. 2018c. Error of the slanted edge method for measuring the modulation transfer function of imaging systems. *Applied Optics*, 57, B83-B91.
- XIE, Y., MAO, Z., BACHMAN, H., LI, P., ZHANG, P., REN, L., WU, M. & HUANG, T. J. 2020. Acoustic Cell Separation Based On Biophysical Properties. *Journal of Biomechanical Engineering*.
- XU, P., LI, Q., ZHANG, B., WU, F., ZHAO, K., DU, X., YANG, C. & ZHONG, R. 2021. On-Board Real-Time Ship Detection in HISEA-1 SAR Images Based on CFAR and Lightweight Deep Learning. *Remote Sensing*, 13, 1995.
- XUE, M., CHEN, M., PENG, D., GUO, Y. & CHEN, H. 2021. One Spatio-Temporal Sharpening Attention Mechanism for Light-Weight YOLO Models Based on Sharpening Spatial Attention. *Sensors*, 21, 7949.
- YANG, R.-J., FU, L.-M. & HOU, H.-H. 2018a. Review and perspectives on microfluidic flow cytometers. *Sensors and Actuators B: Chemical*, 266, 26-45.
- YANG, R. J., FU, L. M. & HOU, H. H. 2018b. Review and perspectives on microfluidic flow cytometers. *Sensors and Actuators B-Chemical*, 266, 26-45.
- YANG, T., VITALI, V. & MINZIONI, P. 2018c. Acoustofluidic separation: impact of microfluidic system design and of sample properties. *Microfluidics and Nanofluidics*, 22, 44.
- YANG, Z. B. & HODGKISS, I. J. 2004. Hong Kong's worst "red tide" — causative factors reflected in a phytoplankton study at Port Shelter station in 1998. *Harmful Algae*, 3, 149-161.

- YODER, J. A., DONEY, S. C., SIEGEL, D. A. & WILSON, C. 2010. Study of Marine Ecosystems and Biogeochemistry Now and in the Future Examples of the Unique Contributions from Space. *Oceanography*, 23, 104-117.
- YUAN, R. & HUI, W. Object Identification and Recognition Using Multiple Contours Based Moment Invariants. 2008 International Symposium on Information Science and Engineering, 20-22 Dec. 2008 2008. 140-144.
- ZARAUZ, L., IRIGOIEN, X. & FERNANDES, J. A. 2008. Changes in plankton size structure and composition, during the generation of a phytoplankton bloom, in the central Cantabrian sea. *Journal of Plankton Research*, 31, 193-207.
- ZHANG, J., HARTMAN, J. H., CHEN, C., YANG, S., LI, Q., TIAN, Z., HUANG, P. H., WANG, L., MEYER, J. N. & HUANG, T. J. 2020. Fluorescence-based sorting of *Caenorhabditis elegans* via acoustofluidics. *Lab Chip*, 20, 1729-1739.
- ZHANG, J., LI, C., YIN, Y., ZHANG, J. & GRZEGORZEK, M. 2023. Applications of artificial neural networks in microorganism image analysis: a comprehensive review from conventional multilayer perceptron to popular convolutional neural network and potential visual transformer. *Artificial Intelligence Review*, 56, 1013-1070.
- ZHU, J., JIANG, J., CHEN, X. & TSUI, C.-Y. SparseNN: An energy-efficient neural network accelerator exploiting input and output sparsity. 2018 Design, Automation & Test in Europe Conference & Exhibition (DATE), 2018. IEEE, 241-244.
- ZMIJAN, R. 2016. *Visible light detection in microflow cytometry - advancing miniaturisation and detection methods*. Doctoral, University of Southampton.
- ZMIJAN, R., JONNALAGADDA, U. S., CARUGO, D., KOCHI, Y., LEMM, E., PACKHAM, G., HILL, M. & GLYNNE-JONES, P. 2015. High throughput imaging cytometer with acoustic focussing. *RSC Adv*, 5, 83206-83216.
- ZOHARY, T., FLAIM, G. & SOMMER, U. 2021. Temperature and the size of freshwater phytoplankton. *Hydrobiologia*, 848, 143-155.