

Durham E-Theses

Mathematical Morphology for Quantification in Biological & Medical Image Analysis

CARL JOHN NELSON

How to cite:

NELSON, CARL JOHN (2017) *Mathematical Morphology for Quantification in Biological & Medical Image Analysis*. Doctoral thesis, Durham University.

Use policy



This work is licensed under a [Creative Commons Attribution 3.0 \(CC BY\)](https://creativecommons.org/licenses/by/3.0/)

Mathematical Morphology for Quantification in Biological & Medical Image Analysis

Carl J. Nelson

A thesis presented for the degree of
Doctor of Philosophy at Durham University



School of Engineering and Computing Sciences
Durham University
United Kingdom

30th May 2017

Mathematical Morphology for Quantification in Biological & Medical Image Analysis

Carl J. Nelson

Submitted for the degree of Doctor of Philosophy
March 2017

Mathematical morphology is an established field of image processing first introduced as an application of set and lattice theories. Originally used to characterise particle distributions, mathematical morphology has gone on to be a core tool required for such important analysis methods as skeletonisation and the watershed transform. In this thesis, I introduce a selection of new image analysis techniques based on mathematical morphology.

Utilising assumptions of shape, I propose a new approach for the enhancement of vessel-like objects in images: the bowler-hat transform. Built upon morphological operations, this approach is successful at challenges such as junctions and robust against noise. The bowler-hat transform is shown to give better results than competitor methods on challenging data such as retinal/fundus imagery.

Building further on morphological operations, I introduce two novel methods for particle and blob detection. The first of which is developed in the context of colocalisation, a standard biological assay, and the second, which is based on Hilbert-Edge Detection And Ranging (HEDAR), with regard to nuclei detection and counting in fluorescent microscopy. These methods are shown to produce accurate and informative results for sub-pixel and supra-pixel object counting in complex and noisy biological scenarios.

I propose a new approach for the automated extraction and measurement of object thickness for intricate and complicated vessels, such as brain vascular in medical images. This pipeline depends on two key technologies: semi-automated segmentation by advanced level-set methods and automatic thickness calculation based on morphological operations. This approach is validated

and results demonstrating the broad range of challenges posed by these images and the possible limitations of this pipeline are shown.

This thesis represents a significant contribution to the field of image processing using mathematical morphology and the methods within are transferable to a range of complex challenges present across biomedical image analysis.

Declaration

The work in this thesis is based on research carried out within the Bioimage Informatics research laboratory, part of the Innovative Computing Group at the School of Engineering and Computing Sciences, Durham University, UK. No part of this thesis has been submitted elsewhere for any other degree or qualification, and it is all the author's work unless referenced to the contrary below.

Note on Publications Included in this Thesis

At the time of submission, three chapters of this thesis had been submitted for publication or published in journals.

Chapter 3 Sazak, Ç., Nelson, C. J. & Obara, B. The Multiscale Bowler-Hat Transform for Vessel Enhancement in Biomedical Images. *IEEE Transactions in Image Processing*. Submitted (2017)

Chapter 4 Nelson, C. J., Duckney, P., Hawkins, T. J., Deeks, M. J., Laissue, P. P., Hussey, P. J. & Obara, B. Blobs and curves: object-based colocalisation for plant cells. *Functional Plant Biology* **42**, 471–485 (2015)

These chapters are presented as submitted, although referencing and notation has been altered and cross-referencing added for consistency throughout this thesis. Some stylistic changes have been made for consistency. The majority of the text is verbatim; however, additions to the body of text are included along with footnotes and, where appropriate, added appendices.

For consistency additional results chapters (chapters 5–6) have been written in the style of papers/proceedings and can be considered self-contained (*c.f.* chapters 3–4). Chapter 5 is planned to be modified and submitted as a research article to an appropriate journal.

The implementation, literature review and writing for chapter 3 (*p.* 20) were completed in partnership with Çiğdem Sazak. The biologically related works and conclusions in chapter 4 (*p.* 48) were contributed by Mike Deeks and Tim Hawkins. The micrographs used in chapter 4 were collected by Mike Deeks, Tim Hawkins and Patrick Duckney. All other aspects of this research, including implementation, data analysis and writing, were carried out by the author.

Note on Publications Not Included in this Thesis

As well as the above papers, the following works have been published during the period of research for this thesis; however, these publications do not fit into the narrative of this thesis and have not been included in the text.

- Nelson, C. J., Dixon, M., Laissue, P. P. & Obara, B. *Speeding up active mesh segmentation by local termination of nodes* in *Medical Image Understanding and Analysis* (July 2014)
- Jackson, P. T. G., Nelson, C. J., Schiefele, J. & Obara, B. *Runway detection in High Resolution remote sensing data* in *Image and Signal Processing and Analysis (ISPA), 2015 9th International Symposium on* (Sept. 2015), 170–175
- Willcocks, C., Jackson, P., Nelson, C. & Obara, B. *Extracting 3D Parametric Curves from 2D Images of Helical Objects*. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PP**, 1–1 (2016)
- Willcocks, C. G., Jackson, P. T. G., Nelson, C. J., Nasrulloh, A. V. & Obara, B. *Interactive GPU Active Contours for Segmenting Inhomogenous Objects*. *Journal of Real-Time Image Processing*. Submitted (2017)

Copyright © 2017 by Carl J. Nelson.

“The copyright of this thesis rests with the author. No quotations from it should be published without the author’s prior written consent and information derived from it should be acknowledged”.

Acknowledgements

Carl J. Nelson was sponsored by the Engineering and Physical Sciences Research Council, UK (Reference 1314229) for all the research presented in this thesis. Further, all research was supported by an academic grant from The Royal Society, UK (RF080232).

Chapter 3 Çiğdem Sazak and Carl J. Nelson contributed equally to this work. Çiğdem Sazak, Carl J. Nelson and Boguslaw Obara are with the School of Engineering and Computing Sciences, Durham University, Durham, UK. Çiğdem Sazak is funded by the Turkey Ministry of National Education.

Chapter 4 The NET1A dataset used for method validation was previously generated as part of the projects supported by the Biotechnology and Biological Sciences Research Council (BB/G006334/1). The authors offer warm thanks to Martin R. Dixon, Durham University, for the Lat52-RFP-FABD2 *Arabidopsis thaliana* lines used in mitochondrial tracking along actin filaments.

Chapter 5 Images for figure 5.11 (p. 104) (b–c) are from the internet and are in the public domain. Images for figure 5.12 (p. 105) (a–c) are from [7], and figure 5.12 (p. 105) (d) is used under Creative Commons Attribution-Share Alike 3.0 Unported license, which can be found here: <https://creativecommons.org/licenses/by-sa/3.0/deed.en>.

Chapter 6 The MRA images used in this paper were kindly provided by Mr Nitin Mukerji of South Tees Hospitals NHS Foundation Trust, UK.



Contents

Abstract	i
Declaration	iii
Note on Publications Included in this Thesis	iii
Note on Publications Not Included in this Thesis	iv
Acknowledgements	v
Contents	vii
List of Figures	xii
List of Tables	xv
List of Algorithms	xvi
List of Acronyms	xvii
1 Introduction	1
1.1 Biomedical Image Analysis	2
1.2 Image Enhancement	2
1.3 Colocalisation	3
1.4 Nuclei Detection	6
1.5 Measuring Vasculature Thickness	7
1.6 Notation and Data in this Thesis	7
1.7 Overview of Thesis and Research Contributions	7
2 An Introduction to Mathematical Morphology	10

2.1	Morphological Operators in Binary Images	11
2.1.1	Dilation & Erosion	11
2.1.2	Opening & Closing	14
2.2	Morphological Image Operators in Greyscale	15
2.3	Some Specific Morphological Operations	16
2.3.1	The Top-Hat Transform	16
2.3.2	Morphological Reconstruction	18
2.3.3	Granulometry	19
3	The Multiscale Bowler-Hat Transform for Vessel Enhancement in Biomedical Images	20
	Prologue	20
3.1	Introduction	21
3.1.1	Vessel Enhancement in Retinal Imaging	21
3.1.2	Contribution and Organisation	23
3.2	Related Work	24
3.2.1	Hessian-based Enhancement Methods	24
3.2.2	Phase Congruency-based Enhancement Methods	26
3.2.3	Contrast Limited Adaptive Histogram Equalisation	27
3.2.4	Matched and Steerable Filter Enhancement Methods	28
3.2.5	Enhancement with Mathematical Morphology	28
3.2.6	Limitations and Challenges	29
3.3	Methods	30
3.3.1	Proposed Method	30
3.3.2	Implementation and Data Sets	34
3.4	Results	34
3.4.1	Synthetic Data	35
3.4.2	Real Data	38
3.4.3	Limitations	45
3.5	Conclusion	45
	Epilogue	45
4	Blobs and Curves: Object-based Colocalisation for Plant Cells	48
	Prologue	48
4.1	Introduction	49

4.1.1	Bioimage Informatics in the Plant Sciences	50
4.1.2	Challenges in Bioimage Informatics	50
4.1.3	Colocalisation of NET1A with Plasmodesmata	51
4.1.4	Colocalisation of Mitochondria with Actin Filaments	52
4.2	Materials and Methods	54
4.2.1	Plant Growth and Sample Preparation	54
4.2.2	Imaging	55
4.2.3	Image Processing	56
4.2.4	Software Packages	63
4.3	Results	63
4.3.1	Colocalisation of NET1A with Plasmodesmata	63
4.3.2	Colocalisation of Mitochondria with Actin Filaments across Time	68
4.4	Discussion	72
4.4.1	Extracting Sub-resolution Information of NET1A	72
4.4.2	Co-imaging of Motile Mitochondria and Cytoskeletal Tracks	75
4.4.3	Prospective Pitfalls in Plant Bioimage Informatics	76
4.4.4	Future Work	77
	Epilogue	78
5	Ellipse Detection by Hilbert-Edge Detection And Ranging (HEDAR)	79
	Prologue	79
5.1	Introduction	80
5.1.1	Ellipse Fitting Methods	81
5.1.2	Voting-based Methods	81
5.1.3	Edge-following Methods	82
5.1.4	Contributions and Overview	82
5.2	Granulometry of Circular Objects in 2D Images	83
5.2.1	From Circles to Ellipses	84
5.2.2	Extension to Greyscale and Higher Dimensions	86
5.2.3	Granulometry to Edge Detection and Ranging	86
5.3	Methods	88
5.3.1	The Algorithm	88
5.3.2	Finding an Ellipse	89
5.3.3	Dealing with Noise	89
5.4	Experiments and Discussion	92

5.4.1	Comparator Methods	92
5.4.2	Running Time and Memory Use	94
5.4.3	Detection Accuracy	95
5.4.4	Robustness Against Noise	100
5.4.5	Experiments on Real Data	101
5.5	Automatic Nuclei Counting in Fluorescent Microscopy	103
5.5.1	Background	103
5.5.2	Experiments & Results	106
5.6	Conclusion	110
	Epilogue	111
6	Granulometry for Thickness Measurement of Cerebral Vasculature	112
	Prologue	112
6.1	Introduction	113
6.2	Methods	114
6.2.1	Automated and Interactive Vasculature Level-Set Segmentation	116
6.2.2	Measuring Vessel Thickness with Mathematical Morphology	116
6.2.3	Visualisation	118
6.2.4	Implementation	118
6.3	Results	118
6.4	Discussion & Conclusions	120
	Epilogue	124
7	Concluding Remarks	125
7.1	Contributions to the Field	126
7.2	Conclusions	126
7.3	Future Work	127
A	The Hessian Matrix for Image Processing	143
A.1	The Scaled 2D Hessian Matrix	143
A.2	Eigenvalues of the 2D Hessian Matrix	144
B	The Phase Congruency Tensor for Image Processing	145
B.1	Phase Congruency	145
B.2	Phase Congruency Tensor	147

C	Using the Hungarian Algorithm for Matching Particles	148
C.1	The Matching Problem.....	148
C.2	Summary of the Hungarian Algorithm	149
D	The Jaccard Similarity Index	150
E	The Hilbert Transform	151
F	Interactive GPU Level-Sets	152

List of Figures

1.1	Vessel enhancement	4
1.2	Colocalisation assay	5
1.3	Nuclei counting and segmentation	6
2.1	Examples structuring elements	11
2.2	Binary dilation	12
2.3	Dilation on a pixel-by-pixel basis	12
2.4	Erosion on a pixel-by-pixel basis	13
2.5	Binary erosion	13
2.6	Dilation with different shape structuring elements (SEs)	14
2.7	Binary opening and closing	15
2.8	Greyscale morphological operations	17
2.9	The top-hat transform	18
2.10	An example of morphological reconstruction	19
3.1	Challenges in vessel enhancement	22
3.2	An example of fundus imaging and blood vessel extraction	23
3.3	Flow chart of the bowler-hat transform	31
3.4	Example openings with disk SEs of various sizes	32
3.5	Example openings with line SEs of various lengths and rotations	32
3.6	Example differences between disk openings and maximum line openings	33
3.7	Cross-sectional profile of original and enhanced vessel	36
3.8	A comparison of vessel enhancement on synthetic images	37
3.9	Mean area under the curve against noise for vessel enhancement methods	39
3.10	Vessel enhancement in the presence of uneven illumination	40

3.11	Examples of true/false positive/negative pixels used for ROC generation	41
3.12	Vessel enhancement of retinal images	43
3.13	Receiver operating characteristic curves for vessel enhancement on real data	44
3.14	Vessel enhancement on real data	46
4.1	A schematic diagram of a plant cell	53
4.2	Retrospective background illumination correction	57
4.3	Schematic explanation of colocalisation	60
4.4	Centroid bias of mass calculations	63
4.5	NET1A-GFP colocalises with aniline blue signals	66
4.6	Euclidean distance map of NET1A-GFP colocalisation with aniline blue signals	67
4.7	RFP-FABD2 and MitoTracker Green segmentation	70
4.8	Bleaching artefacts affecting object colocalisation	71
4.9	Colocalisation of RFP-FABD2 and MitoTracker Green signals	73
4.10	Relative positioning of molecular components of root meristem plasmodesmata	74
5.1	Classical granulometry for circles in 2D binary images	85
5.2	Example detection of an ellipse	90
5.3	Detecting steps with the Hilbert transform	91
5.4	Comparison of computational running time for methods	96
5.5	Algorithm response with respect to major axis length and axes ratio	97
5.6	Algorithm response with respect to rotation and axes ratio	98
5.7	Examples of the results for each algorithm	99
5.8	HEDAR can detect the number of ellipses in an image	100
5.9	Comparing the robustness of HEDAR over different noise types	101
5.10	Examples of HEDAR run on synthetic noisy images	102
5.11	Examples of HEDAR and comparator methods run on real images	104
5.12	Examples of HEDAR on real greyscale images with bright & dark objects	105
5.13	HEDAR on images with bright and dark objects	105
5.14	Examples of images from the Broad Bioimaging Benchmark Collection	107
5.15	HEDAR results on benchmark dataset BBBC001	108
5.16	HEDAR results on benchmark dataset BBBC002	109
5.17	HEDAR results on benchmark dataset BBBC004	110
6.1	Maximum intensity projection of MRA image	113
6.2	Flow chart of the proposed global approach	115

6.3	Automatically identified vascular material for ϕ initialisation.....	117
6.4	Semi-automated segmentation of brain vasculature	119
6.5	Severe artefacts in the image can prevent accurate segmentation.....	121
6.6	Results of thickness calculation on a full segmentation of brain vasculature	122
6.7	Results of thickness calculation on a vasculature network with an aneurysm.....	123
A.1	Visualising the Hessian matrix eigenvalues	144
B.1	Phase congruency at image features.....	146

List of Tables

3.1	Mean area under the curve values for vessel enhancement on real data	42
5.1	HEDAR mean percentage counting errors on benchmark datasets.....	108
C.1	Protein colocalisation framed as an assignment problem	148
C.2	Matching pairs of proteins based on separation distance	149

List of Algorithms

2.1	Morphological reconstruction	18
4.1	Background Illumination Correction	58
4.2	Mask-based Normalisation	59
4.3	Colocalisation of Blobs	61
4.4	Segmentation of cables	62
5.1	Edge Detection And Ranging	88
5.2	Hilbert-Edge Detection and Ranging	93
6.1	Measuring segmented vessel thickness with sequential openings	118

List of Acronyms

ND <i>N</i> -dimensional	81
1D one-dimensional	85
2D two-dimensional	1
3D three-dimensional	1
4D four-dimensional	86
AUC Area Under the Curve	39
CLAHE contrast limited adaptive histogram equalisation	27
DRIVE Digital Retinal Images for Vessel Extraction	23
EliFit ElliFit	93
EMCCD electron-multiplying charge-coupled device	55
FABD2 fimbrin actin-binding domain 2	54
GFP green fluorescent protein	52
GPU graphical processing unit	7
HEDAR Hilbert-Edge Detection And Ranging	9
Hough elliptical hough transform	93
HyperRN hyper renormalisation	93
IGAC interactive GPU active contours	116
LoG Laplacian of Gaussian	59
MRA magnetic resonance angiography	1
NET1A NETworked protein 1A	52
PCT phase congruency tensor	27
RANSAC random samle consensus	82
RFP ref fluorescent protein	54
ROC Receiver Operating Characteristic	40
RVR regularised volume ratio	26

SE structuring element	xii
SF structuring function	15
STARE STructured Analysis of the REtina	23
TOF time-of-flight	7
VMTK Vascular Modelling Toolkit	114
YAED fast and effective ellipse detector	93

Chapter 1

Introduction

In this thesis, I explore a set of mathematical techniques, those of mathematical morphology, in relation to image processing and analysis. The power of these methods for image enhancement is demonstrated in the context of medical images and a technique that was derived from the ideas behind granulometry, a particular morphological approach, is developed and analysed in the context of a bioimage analysis task. The methods developed in this thesis apply to a range of imaging scenarios but particularly to problems found in the medical and biological imaging fields.

First, I introduce the use of mathematical morphology for the enhancement of vessel-like structures in two-dimensional (2D) medical imaging, specifically retinal imaging. Second, I show the use of a mathematical morphology-based approach to blob/particle detection and quantification in a standard biological assay: colocalisation of proteins. Third, I develop a method that has been derived from granulometric approaches for the detection, counting and quantification of ellipse-like objects in general images. This approach is then applied to another common task in bioimage analysis, that of nuclei detection and counting. Finally, I introduce a mathematical morphology approach to measuring the thickness, in three-dimensional (3D), of vasculature networks, specifically those captured by time-of-flight magnetic resonance angiography (MRA).

This chapter will introduce key background concepts in biomedical image analysis including a brief overview of image enhancement, colocalisation, nuclei detection and vasculature thickness. I will also outline the structure of this thesis.

1.1 Biomedical Image Analysis

Imaging has been an essential tool in the life sciences for around 350 years [8]. Much of our current biological knowledge and many medical diagnostic protocols have been built upon imaging technologies. However, in the last decade or so, the size, detail and quantity of biological and medical imaging data has risen to colossal amounts. In modern biological experiments, using new super-resolution and lightsheet approaches, it is not unusual to accumulate terabytes of data per experiment, *e.g.* [9]. Thus, it has become increasingly important for computational methods to fully and partly automate analysis, quantification and annotation of biological and medical images [10].

The key challenges in all biological and medical image analysis lie in the intra- and inter-image variability of features and objects as well as in the complex nature of biological objects themselves [11]. As such, newly developed approaches to any stage of analysis must be either robust to large portions of this variability or be supported by clear and practical assumptions on the choice of modality, preparation, assay and features of interest.

Despite the broad range of modalities, assays, preparations and contexts, much of biological and medical image analysis fall into several key categories: image registration, image enhancement, feature and object detection, object segmentation, measurement, and feature or object classification. In this thesis I have developed tools for image enhancement, object detection and measurement; these tools can be integrated with other techniques to fulfil more complex requirements and pipelines.

1.2 Image Enhancement

Many images, particularly in biological and medical imaging, suffer from problems that obscure objects or features of interest. For example, low contrast between an object of interest and background is a common problem in many forms of medical imaging where tissues of interest may not be distinguishable from neighbouring tissues. As such, image enhancement is often an early step in an analysis pathway [12].

Image enhancement aims to improve the visual quality of an image to aid human qualitative analyses and automated quantitative analyses of features/objects in the image. For example, in our tissue example, image enhancement tools for enhancing the contrast at boundaries between

tissues may be employed before semi-automated segmentation of the tissues of interest. Alternatively, a denoising process may be applied to images with large amounts of noise due to the low efficiency of fluorophores and thus the high exposure needed for fluorescent imaging with particular labels.

In this thesis, I focus on one particular area of image enhancement: the strengthening of vessel-like objects in greyscale images (see figure 1.1). Tools for this problem are designed to enhance (often by brightening) vessel-like objects relative to the image background and lessen non-vessel-like objects (*i.e.* blobs) [13]. However, many of the state-of-the-art approaches still have significant problems with the loss of signals at junctions, the mistaken enhancement of non-vessel-like objects and difficulties in dealing with the intrinsically multi-scale nature of biological systems (see chapter 3 (*p.* 20)).

1.3 Colocalisation

In biology, colocalisation assays are used to investigate if two molecules meet in space and time. If there is significant correlation in the localisation of these two molecules, then there is a high chance that the two proteins interact either directly or indirectly. Colocalisation is usually one of several assays carried out to probe interactions between molecules.

The classical approach to colocalisation is to fluorescently label the two proteins of interest and image them, one protein per channel, with fluorescence microscopy. Often, these images are analysed qualitatively by overlapping the two images (see figure 1.2 (*p.* 5)); however, it is standard to take a simple quantitative measure by calculating the Pearson correlation coefficient between the two images [14].

However, this approach to colocalisation is pixel-dependent and intensity-dependent and object-based approaches to colocalisation have been introduced over recent years [15]. In object-based approaches, the signals of interest are segmented first, and their ‘colocalisation’ is measured against a statistical model. Additional information can be gained from this approach, such as distance or orientation.

In this thesis I extend an object-based colocalisation method (chapter 4 (*p.* 48)). This approach identifies putative particulate objects using a blob enhancement step and Gaussian kernel fitting; identifies putative filamentous objects with local thresholding and pairs signals based on the Hungarian algorithm [16]. Importantly, I highlight the areas of uncertainty in such a method and such concerns extend to other existing object-based colocalisation algorithms.




Figure 1.1: Top: an ideal vessel enhancement: the original image (a) contains a bright blob-like object and a low-contrast vessel-like object, after enhancement (b) the vessel is bright and the blob has been suppressed. Bottom: a real example of vessel enhancement, the original image (c) has been processed using the the bowler-hat transform (described in this thesis) to provide the enhanced image (d). Note that, in the real images, the enhancement is not perfect (*c.f.* synthetic images) but the vessels are enhanced and the blobs dulled in this sample




Figure 1.2: Visual colocalisation of two molecules A and B. Schematically, particles of molecule A (a) and particles of molecule B (b) correlate in position (c). Fluorescence micrographs captured of each protein (d,e) can be overlaid to show areas of colocalisation (f).

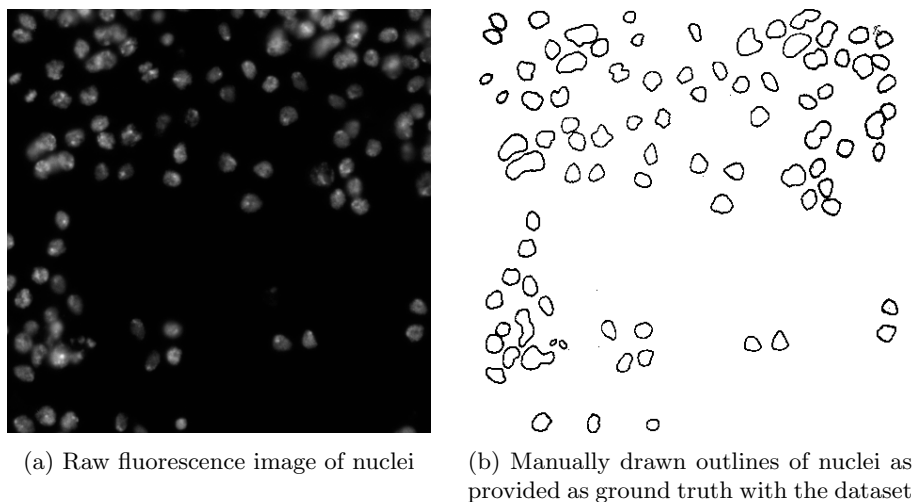


Figure 1.3: Nuclei and cell counting in fluorescent images (a) is an important step in its own right but can also be a starting step in full cell segmentation (b). Image from the Broad Bioimage Benchmark Collection [21, 22].

1.4 Nuclei Detection

Nuclei detection is required for a large amount of quantitative biology: nuclei counting may be the analysis needed for high-throughput microscopy (*e.g.* in histopathology [17]), or detection may be used to seed a segmentation algorithm for detailed analysis of the morphodynamics of a cancerous cell culture [18]. There are many approaches to this problem as different modalities, various cell types and different preparations each provides unique challenges. A large number of these approaches are based on general ellipse detection algorithms, such as elliptical Hough transforms [19], kernel fitting [20] and beyond (see section 5.5 (*p.* 103) for a brief summary of the field).

Many current approaches to nuclei detection or segmentation lack the ability to count nuclei or cells accurately if cells are clustered or overlapping; struggle with inhomogeneity or inter-nuclei variation; may be sensitive to parameters or indeed rely on a large number of parameters; or rely on *a priori* knowledge, such as a predetermined number of expected nuclei (see [23] for a review).

In this thesis, I introduce a new approach to general ellipse detection derived from granulometry and built on the idea of Hilbert-edge detection and ranging and the assumptions of ellipse

symmetry and signal shape (chapter 5 (p. 79)). This method is demonstrated and validated on a broad range of synthetic and real images as well as being applied to the more specific problem of nuclei detection (section 5.5 (p. 103)).

1.5 Measuring Vasculature Thickness

The thickness of vasculature is often seen as an important metric for automated diagnosis and quantification of vascular abnormalities, *e.g.* [24–26]. However, many approaches to measuring the thickness of vascular networks and abnormalities rely on either 2D approximations of the thickness [27, 28], which may be automated, and/or manual measurement or annotation of these images [29, 30].

In this thesis, I describe an approach to 3D thickness calculation for brain vasculature captured using time-of-flight (TOF) MRA. This pipeline includes a semi-automated segmentation stage using state-of-the-art graphical processing unit (GPU) level-set technologies, followed by thickness calculation using sequential mathematical openings. We show the results of this approach to be good at measuring thickness and sensitive enough to highlight abnormalities, such as aneurysms (chapter 6 (p. 112)).

1.6 Notation and Data in this Thesis

Throughout this thesis I will refer to images, structuring elements and operations on these. For the purpose of this text an image is a 2D or 3D array of pixels with floating points values in the range $[0, 1]$ (greyscale) or integer values 0 or 1 (binary). The same applies to structuring elements/functions, introduced in the next chapter, and kernels.

1.7 Overview of Thesis and Research Contributions

This thesis is divided as follows: first, I introduce the basic concepts behind mathematical morphology before moving on to discuss image enhancement, introducing a newly developed approach for vessel enhancement based on mathematical morphology. Then I describe two methods for the detection and quantification of blob-like objects at different scales: one, for small particles, based on Gaussian kernel fitting and the other, for larger objects, based on granulometry. Finally, I introduce a thickness measurement approach built on sequential morphological openings for the analysis of vascular networks in 3D. The chapters are laid out as follows:

Chapter 2 In this chapter I will introduce the basic principles of mathematical morphology, specifically within the context of greyscale image processing. I demonstrate erosion, dilation, opening and closing in 2D binary and greyscale images before briefly introducing other morphological transforms used in this thesis.

Chapter 3 Vessel enhancement, particularly before the extraction of vascular structures, is a key stage in many automated diagnostic and clinical pipelines. Whilst a number of image processing techniques have been proposed for this, many still have issues, particularly with respect to junctions and the broad range of scales involved in a biomedical network. In this chapter, I introduce a new approach to vessel enhancement based on mathematical morphology. Mathematical morphology is widely used in image processing and has been used in image enhancement previously. However, the approach proposed here, the bowler-hat transform, combines different structuring functions together to identify innate features of vessel-like objects. This new approach is robust at junctions and able to cope with variations in thickness and contrast throughout vessel-like structures. I evaluate the proposed approach with both synthetic and real data and compare the results quantitatively with several existing vessel enhancement methods. My results show that this approach is promising with high-quality enhanced images, accuracy at junctions and robustness to noise.

Chapter 4 Blobs and curves occur everywhere in plant bioimaging: from signals of fluorescence-labelled proteins, through cytoskeletal structures, nuclei staining and cell extensions such as root hairs. In this chapter, I look at the problem of colocalisation of blobs with blobs (protein-protein colocalisation) and blobs with curves (organelle-cytoskeleton colocalisation). This chapter demonstrates a clear quantitative alternative to pixel-based colocalisation methods and, using object-based methods, can quantify not only the level of colocalisation but also the distance between objects. Included in this chapter are computational algorithms, biological experiments and guidance for those looking to increase their use of computationally-based and quantified analysis of bioimages. This chapter describes an extension of an original method in a way that aims to explain to non-computing experts how to understand the limitations and pitfalls in computational approaches to solving image analysis problems. The key contributions of this chapter are the improved local pruning of putative particles.

Chapter 5 Detection of ellipses is an issue that continues to arise in many areas of image processing: from identifying objects in complex visual scenes to detecting nuclei in large-scale bioimages. In this chapter, I introduce a new method for the detection of ellipse-like

objects that uses simple morphological operations, building on aspects of granulometry, and common signal processing. With this approach, Hilbert-Edge Detection And Ranging (HEDAR), it is possible to easily extract the position, size and rotation of an unknown number of ellipse-like objects in an image. HEDAR is low parameter and robust to noise in greyscale images. Demonstrations of the effectiveness and robustness of HEDAR are shown and compared to multiple state-of-the-art methods. These demonstrations show the advantages and disadvantages of this approach over other ellipse detection methods. The key contribution of this chapter is the new approach to ellipse detection based on HEDAR). This chapter also includes a case-study in nuclei counting and detection in fluorescence microscopy (section 5.5). Accurate and reliable nuclei counting is still an important part of quantitative and high-throughput microscopy. Many methods exist that attempt to fulfil this task; however, all methods have limitations. Such limitations include sensitivity to user parameters, a need for pre-processing, modality-specific features, and more. Here, I demonstrate the application of HEDAR to nuclei detection. This method exploits the symmetric properties of nuclei and uses Hilbert-edge ranging to identify their positions and sizes. I validate this method and measure its performance by utilising datasets from the Broad Bioimage Benchmark Collection. I have also been careful to show the limitations of this approach by including examples of cases where this method fails to detect nuclei. I show this new approach to have a low mean error and be robust against clustering of nuclei. I believe this approach to be a competitive algorithm for nuclei counting and detection in microscopy. The key contribution of this case-study is the application of HEDAR to fluorescence microscopy and nuclei detection.

Chapter 6 Accurate 3D measurement of vascular thickness is a key step in the automated, quantitative analysis of medical images and is essential for accurate, automatic, differential diagnosis. Few methods exist that can precisely measure thickness in 3D, and many existing approaches are non-automatic, relying on human segmentation and/or measurement, often in 2D rather than 3D images. Here I propose the use of sequential morphological operations to measure the thickness of objects in 3D accurately. In this chapter, I introduce a pipeline for the segmentation and measurement of brain vasculature from time-of-flight magnetic resonance images for a semi-automated solution to this complex problem.

Chapter 7 Finally, I summarise this thesis, highlighting the key contributions to the field and suggest possible future work for extending the use of mathematical morphology to other image processing and analysis problems in the biological and medical imaging domains.

Chapter 2

An Introduction to Mathematical Morphology

Mathematical morphology is an image processing field based on ideas derived from set theory and lattice theory. First introduced in 1964 by Georges Matheron and Jean Serra, mathematical morphology was originally used to characterise the distribution profiles of grains and particles in binary images [31].

Later generalised to greyscale, mathematical morphology has now been utilised for a wide variety of problems including, for example, noise suppression, *e.g.* [32], and edge detection, such as [33]. Techniques based on mathematical morphology now include such key methods as skeletonisation [34], granulometry [35] and the watershed transform [36].

In this chapter, I will introduce the underlying concepts behind mathematical morphology. First, I will introduce the two basic operations of dilation and erosion in a binary scenario, followed by opening and closing, each an ordered combination of dilation and erosion. I will discuss the extension of these four operators to greyscale before going on to briefly introduce three more complex methods built on mathematical morphology: the top-hat transform, reconstruction and granulometry for image quantification and segmentation.




Figure 2.1: (a-d) Example 3×3 binary structuring elements.

2.1 Morphological Operators in Binary Images

For the sake of simplicity, I will first introduce mathematical morphology in terms of binary images and flat*, binary structuring elements. Here I will focus on the inputs and outputs of morphological operations; however, it is worth bearing in mind that, in binary, each image is being considered as a set of points embedded in \mathbb{Z}^2 . With these simple operations, we can begin to build more complex binary image transformations.

First, we must define our structuring element or SE. In mathematical morphology, the SE is equivalent to the kernel used for a convolution. For binary morphology, our structuring element will always be a) binary, b) smaller than the image of interest, and c) of a shape that differentiates objects or features of interest in the image (*e.g.* figure 2.1). For example, if we are interested in circular particles, we might choose to use a SE containing a binary disk of the same size as the particles we are interested in. Obviously, the choice of SE may also be impacted by the operations or transforms being undertaken (see below).

2.1.1 Dilation & Erosion

The majority of early mathematical morphology is built upon two basic operators: dilation and erosion. Dilation will, in general, expand or dilate an object in a manner that matches the SE used. Mathematically, binary dilation, \ominus , on image I with SE b is defined as,

$$I \ominus b = \{ \vec{p} \in \mathbb{Z}^2 \mid (b^s)_{\vec{p}} \cap I \neq \emptyset \}, \quad (2.1)$$

where b^s is the symmetric of b , $b^s = \{x \mid -x \in b\}$, and $(b^s)_{\vec{p}}$ is a translation of b^s by vector \vec{p} , *i.e.* all points in space.

* A flat structuring element is 1, or true, at points of interest and 0, or false, otherwise. A non-flat structuring element may have any finite value at pixels.




Figure 2.2: Dilation of a binary image with a binary structuring element (SE).

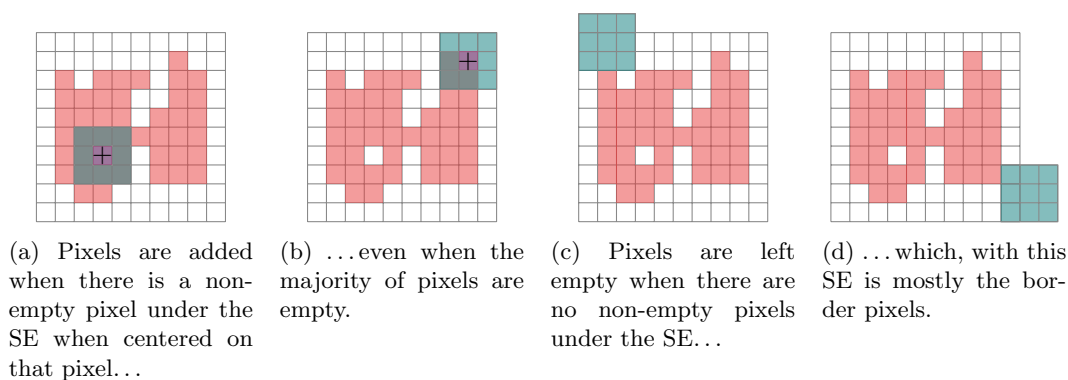


Figure 2.3: Dilation on a pixel-by-pixel basis where SEs cover non-empty (a-b) and empty (c-d) pixels.

Binary dilation for 2D images is more easily described visually. Figure 2.2 shows an example of dilation with a binary image and SE. Note how the original object is dilated based on the square shape of the SE. It is important to remember that dilation (and erosion) occurs simultaneously (see figure 2.3 and figure 2.4) at all pixels, so the kernel always sees the original image only; as such, the ordering in figures 2.2 and 2.5 is arbitrary.

The dual[†] of the dilation operation is erosion. Mathematically, binary erosion, \ominus , on image I with SE b is defined as,

$$I \ominus b = \{ \vec{p} \in \mathbb{Z}^2 | b_{\vec{p}} \subseteq I \}, \quad (2.2)$$

[†] $I \ominus b = (I^c \ominus b^c)^c$, where x^c is the complement of x .




Figure 2.4: Erosion on a pixel-by-pixel basis where SEs cover empty (a-b) and non-empty (c-d) pixels.




Figure 2.5: Erosion of a binary image with a binary structuring element (SE).

where $b_{\vec{p}}$ is a translation of b by vector \vec{p} , *i.e.* all points in the image.

Again, erosion is more easily described visually. Figure 2.5 illustrates erosion of a binary image and SE (*c.f.* figure 2.2 (*p.* 12)). Note how the boundaries of the original object are eroded based on the square shape of the SE.

It is important to consider the shape of the SE used, for example, in figure 2.6 the same image is dilated with two different SE. Note how the results are different; although this is minimal for such a simple example, when using mathematical morphology for feature detection and extraction the shape of the SE becomes increasingly important.

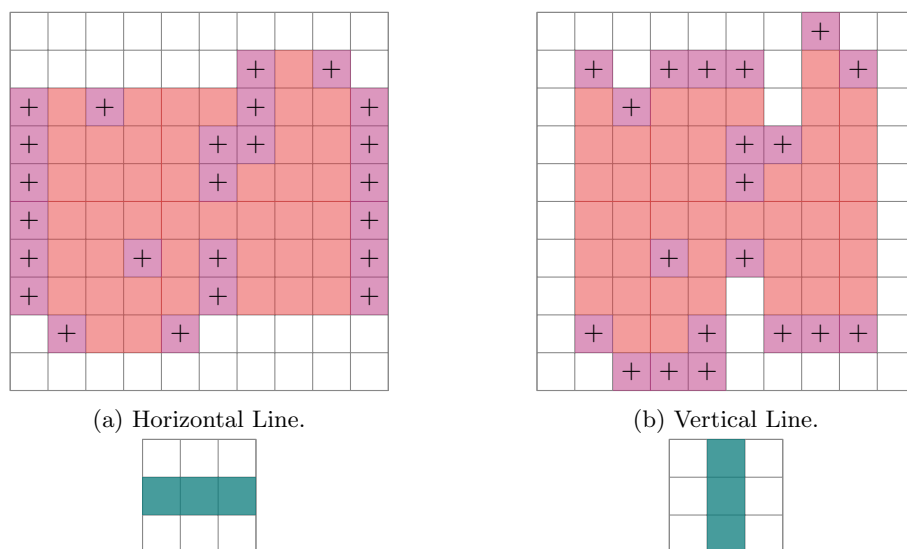


Figure 2.6: Dilation of a binary image with a horizontal (a) and vertical (b) line SEs.

2.1.2 Opening & Closing

Using these two basic operations, we can begin to build more complex operations. Two commonly used operations are opening and closing. The opening is composed of erosion with image I and SE b followed by dilation of the result and SE b :

$$I \circ b = (I \ominus b) \oplus b. \quad (2.3)$$

Contrariwise, closing consists of a dilation followed by an erosion:

$$I \bullet b = (I \oplus b) \ominus b. \quad (2.4)$$

Figure 2.7 shows an example of opening and closing using the same starting image and SE used previously.

Given a regular structuring element, *e.g.* a small disk, then the opening and closing operations are non-linear transformations that smooth the contours of the image. Opening with such a SE will suppress sharp protrusions and will cut narrow bridges between larger components (figure 2.7a); imagine your image as an island, opening raises the sea level in a way based on the SE shape and size. On the other hand, closing with such a SE will fill thin breaks between and small holes

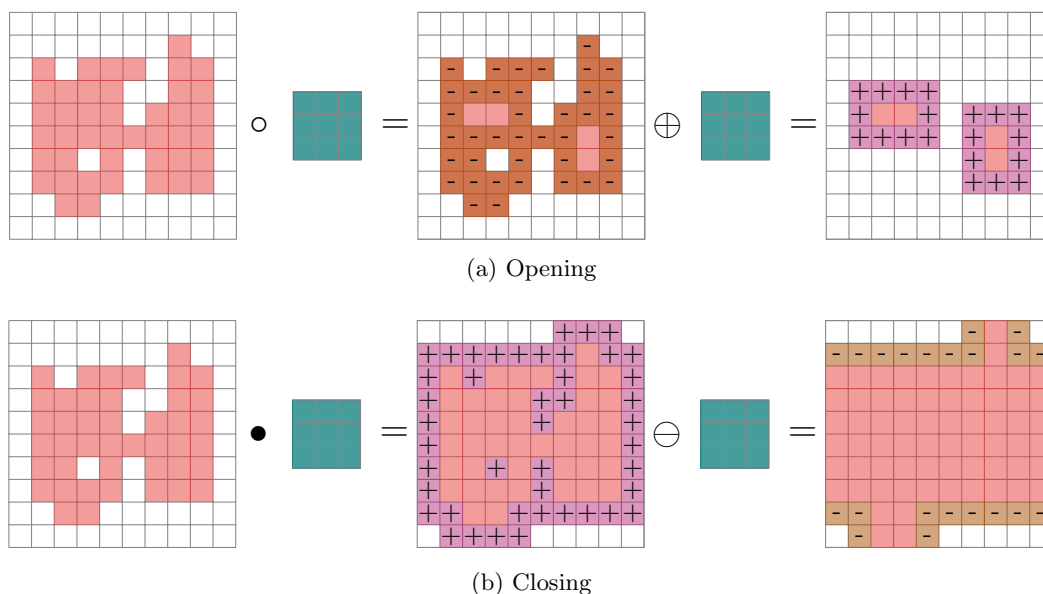


Figure 2.7: Opening (a) and closing (b) a binary image with a binary SE.

within larger components (figure 2.7b); here our island is suffering from drought, and the lakes and rivers in the island have dried up.

2.2 Morphological Image Operators in Greyscale

At this point, it is easier to leave the mathematics of set theory and lattice theory behind and focus on the processes and results of greyscale morphology purely in terms of images and structuring functions (*c.f.* kernels for convolution).

First, when we move to greyscale images, we can now consider a generalisation of structuring elements: a structuring function. These structuring functions (SFs) can be seen as analogous to a weighted kernel and will be a) of ‘pixels’ with values in the greyscale domain, *e.g.* between 0 and 1, b) smaller than the image of interest, and c) of a shape and weighting that differentiates objects or features of interest in the image. For example, whilst in binary images we may have used an approximation to a disk, which will contain various discretisation errors, for greyscale morphology we might use an aliased disk to help decrease these errors.

We must also redefine dilation and erosion to account for the embedding of the image in \mathbb{R} as opposed to \mathbb{Z} as for the binary case. As such, we generalise dilation to,

$$(I \oplus b)(\vec{p}) = \sup_{\vec{q} \in \mathbb{R}^2} (I(\vec{q}) + b(\vec{p} - \vec{q})). \quad (2.5)$$

Moreover, erosion generalises to,

$$(I \ominus b)(\vec{p}) = \inf_{\vec{q} \in \mathbb{R}^2} (I(\vec{q}) - b(\vec{q} - \vec{p})). \quad (2.6)$$

Where $\sup_{\vec{q}}$ is the supremum over all \vec{q} , for image processing this is the equivalent to the weighted maximum over the SF centred at \vec{q} , and $\inf_{\vec{q}}$ is the infimum over all \vec{q} , equivalent to the weighted minimum over the SF centred at \vec{q} .

Figure 2.8 a–b shows examples of greyscale dilation and erosion. Dilation expands maxima within the image whilst erosion decreases the area of maxima, *i.e.* expands minima. Note that, due to the disk-shaped SF, corners and features of bright (dark) areas are rounded off in the dilation (erosion). Greyscale opening and closing are demonstrated in figure 2.8 c–d. Here, opening smooths the contour of the image and removes spurs and bridges between components. Closing, on the other hand, smooths the contour by filling valleys and ‘coves’ of the main component.


2.3 Some Specific Morphological Operations

Here I briefly introduce a small number of the myriad of morphological operations and applications. These are transforms that have been used in the development of the methods in this thesis. Readers who wish to extend their knowledge of mathematical morphology within the field of image processing are directed to [31, 37].


2.3.1 The Top-Hat Transform

The top-hat transform, $TH(I, b)$, is a commonly used, simple morphological operation that extracts small details from images, $I(\vec{p})$, given a specific SF, $b(\vec{p})$. The top-hat transform is the difference between an image and its opening, *i.e.*


$$TH(I(\vec{p}), b(\vec{p})) = I(\vec{p}) - (I \circ b)(\vec{p}). \quad (2.7)$$




(a) Dilation.



(b) Erosion.



(c) Opening.



(d) Closing.

Figure 2.8: Morphological operations on a greyscale image with a greyscale SF: dilation (a), erosion (b), opening (c), and closing (d).




Figure 2.9: The top-hat transform removes objects greater in size than the kernel used.

Where $TH(I, b)$ is an image that contains only those bright objects smaller than the SF, b (see figure 2.9 for example). The top-hat transform has been used, amongst other purposes, as the basis for vessel enhancement and segmentation approaches [38].

2.3.2 Morphological Reconstruction


Morphological reconstruction is a general class of mathematical-morphology-based tools that can be used to extract features from an image (figure 2.10). A generalisation of flood-filling, morphological reconstruction starts with two images 1) a ‘marker’ image, ϕ , which describes where the flooding will begin, and 2) a ‘mask’ image, I_{mask} , which describes the boundaries where the flooding must stop.

During reconstruction, flooding begins at high peaks and ridges in the marker image and spreads or dilates whilst being forced to fit within the mask image until no more changes can be made. Algorithmically, reconstruction is an iterative procedure, as shown in algorithm 2.1 [39].

Input : ϕ, I_{mask}
Output: I_{recon}

- 1 $I_1 \leftarrow \phi$
- 2 $b = ones(3)$
- 3 **repeat**
- 4 | $I_{k+1} = (I_k \oplus b) \cap I_{mask}$
- 5 **until** $I_{k+1} = I_k$
- 6 $I_{recon} \leftarrow I_k$

Algorithm 2.1: Morphological reconstruction



(a) Example mask image with objects of various intensities
(b) Example seed image with seeds only for those objects brighter than 0.5
(c) Resulting morphological reconstruction captures only the bright objects from the mask image

Figure 2.10: A simple example case of morphological reconstruction of a seed image with a given mask image to capture only certain features of the mask image.

2.3.3 Granulometry

Originally introduced in the 1960's, granulometry is one of the earlier techniques developed using mathematical morphology and is a method for determining the distribution of granule/particle size within an image [35]. Granulometry relies on the difference of morphological openings of sequentially larger SEs to determine the distribution of sizes in an image. Granulometry is described in more detail in section 5.3 (*p.* 88).

Chapter 3

The Multiscale Bowler-Hat Transform for Vessel Enhancement in Biomedical Images

Prologue

In chapter 2 (p. 10), I have provided the basic background to mathematical morphology for greyscale image processing and analysis. This knowledge will be key to understanding the methods proposed in this and other chapters of this thesis. In this chapter, I will introduce a mathematical-morphology-based approach to image enhancement, specifically the enhancement of vessels in biological and medical images.

Vessel enhancement in biomedical images is an important pre-processing step as, in an ideal case, the enhancement allows segmentation methods to segment networks of vessels, *e.g.* vasculature, correctly and, particularly, maintain the connectivity of the system.

In this chapter, I will briefly summarise the field of vessel enhancement for 2D images followed by a description of a novel vessel enhancement approach, the bowler-hat transform. This new method is qualitatively and quantitatively compared to a wide range of comparator methods on both synthetic data and real datasets. These results show that the bowler-hat transform is as good or better a vessel enhancer than these comparator methods for the datasets used.

Declaration: This chapter is based on the following publication: Sazak, Ç., Nelson, C. J. & Obara, B. The Multiscale Bowler-Hat Transform for Vessel Enhancement in Biomedical Images. *IEEE Transactions in Image Processing*. Submitted (2017). This chapter is presented as submitted, although referencing and notation has been altered and cross-referencing added for consistency throughout this thesis. Some stylistic changes have been made for consistency. The majority of the text is verbatim; however, additions to the body of text are included along with footnotes and, where appropriate, added appendices.

3.1 Introduction

Many biological and medical images contain vessel-like structures, such as vasculature or cytoskeletal networks. Extraction of these vessels and their network is often a key step in the quantified or automated analysis of such images and, for medical imaging, such quantification may be used for automated diagnoses. However, most biological and medical modalities may suffer from poor quality due to many reasons, including, but not limited to, noisy image capture, sample/patient variability and low contrast scenarios. As such, standard segmentation or network extraction methods may not work on the original, unenhanced image and some form of vessel enhancement is required (see [40] for a recent review).


There are many approaches to vessel enhancement, such as Hessian matrix approaches [41, 42], phase congruency tensor methods [43], contrast limited adaptive histogram equalisation [44] and more [45]. These methods all show increased segmentation performance compared to the unenhanced, original images in, for example, vasculature extraction of fundus images. However, many existing methods have practical issues, such as the loss of signals at junctions (figure 3.1a) [41], false introduction of vessel-like structures (figure 3.1b) [46] and complex parameter spaces [43].

3.1.1 Vessel Enhancement in Retinal Imaging

Retinal vessel segmentation and measurement of blood vessels are used for diagnosis, screening and evaluation for a range of diseases including diabetes and arteriosclerosis [47]. Manual




(a) Junction artefacts such as the loss of signal at the centre of this cross (left: original; right: enhanced with vesselness method)



(b) False vessels created by the accidental enhancement of noise (left: original; right: enhanced with RVR method)

Figure 3.1: Some of the challenges faced when using many current and state-of-the-art vessel enhancement methods include loss of signal at junctions (a) and false vessels (b).



(a) A fundus image of a retina showing the optic disc, where vessels appear to branch from, the macula and, importantly, many blood vessels.

(b) A manual segmentation of the blood vessels in (a).

Figure 3.2: Fundus or retinal imaging is an important tool for ophthalmologists and used for diagnosing and monitoring many retinal diseases and their treatment.

segmentation of the vasculature is both time consuming and requires expertise and patience [48]. However, automated segmentation methods often struggle with unenhanced retinal images and require an image enhancement preprocessing step [48].

The vasculature in retinal images appear as elongated, branching structures with a wide variety of vessel widths and tortuosities and usually a wide range of intensities and contrast. Such images also include various non-vasculature objects, such as the fovea. These challenges make retinal imaging an ideal case study for the development of new vessel enhancement measurements. In this paper, we compare our method to various other approaches using publicly available retinal image databases for ground truths and quantitative analyses.

3.1.2 Contribution and Organisation

In this paper, we introduce a new vessel enhancement approach based on mathematical morphology and the elongated shape of vessel structures. The proposed method, the bowler-hat transform, shows visually and quantitatively improved results on a range of test, synthetic images and also on two public databases of fundus images: the Digital Retinal Images for Vessel Extraction (DRIVE) and STructured Analysis of the REtina (STARE) collections.

This rest of this paper is organised as follows. In section 3.2, we introduce key papers and approaches from the vessel enhancement literature. Section 3.3 introduces and explains our approach. Section 3.4 presents validation experiments and results on synthetic and real data. Finally, we conclude this paper and discuss the future of this approach.

3.2 Related Work

3.2.1 Hessian-based Enhancement Methods

Hessian matrix-based methods have been very popular in image enhancement methods. These methods stem from the early work by Frangi and colleagues [41] who used the eigenvalues of Hessian matrices at various scales to enhance curvilinear/tubular objects in 2D/3D respectively [41, 42]. Here we describe and compare three Hessian matrix-based enhancement methods: 1) the original ‘vesselness’ measurement as introduced in [41], the ‘neuriteness’ measurement described in [46] and a recently proposed measurement based on the regularised volume ratio in [13].

Vesselness

The vesselness measurement describes images in terms of the eigenvalues of the Hessian matrix. For vessel-like structures, *i.e.* elongated, smooth structures, we define the following eigenvalue relationship:

$$\lambda_2 \gg \lambda_1, \quad (3.1)$$

where the eigenvalues are defined by the elements of the Hessian matrix^{*}, H_{ij} and $(i, j) \in [1, 2]$, such that

$$\lambda_{1,2} = \frac{1}{2} \left(H_{11} + H_{22} \pm \sqrt{(H_{11} - H_{22})^2 + 4H_{12}^2} \right). \quad (3.2)$$

In general, if the magnitude of both eigenvalues is small, then the pixel is likely to be background; if λ_1 is small and λ_2 high then the feature is elongated, *e.g.* vessel-like; and if both are high then the feature is blob-like. This leads to a failure at the intersection of vessels as the difference between eigenvalues tends towards zero and thus that feature can be lost [49][†].

^{*} Actually, these methods use a scaled Hessian that, for completeness, has been described in appendix A.1 (p. 143).

[†] See appendix A.2 (p. 144) for a visualisation of this problem in 2D.

In the original paper by Frangi and colleagues [41], the vesselness measurement is then defined based on the ratio of eigenvalues as

$$\mathcal{V}_o = \begin{cases} 0 & \text{if } \lambda_2 > 0, \\ \left(e^{-\frac{\mathcal{R}_B^2}{2\beta^2}} \right) \left(1 - e^{-\frac{\mathcal{S}^2}{2c^2}} \right) & \end{cases}, \quad (3.3)$$

where

$$\mathcal{R}_B = \lambda_1/\lambda_2, \quad \mathcal{S} = \sqrt{\lambda_1^2 + \lambda_2^2}, \quad (3.4)$$

where generally β is fixed to 0.5 and c is equal to half of the maximum Frobenius norm (\mathcal{S}).

Neuriteness

An alternative to the vesselness measure was developing by Meijering and colleagues in [46] to enhance low contrast and highly inhomogeneous images of neurites, small extensions of neurons including axons and denrites. This neuriteness measurement consists of a detection stage [50], which infers a putative neurite in every pixel of the image that has a non-zero value; and a tracing stage, which determines the pixels that are most likely to reflect the centre-lines of true neurites based on the original image [51].

In [46], Meijering and colleagues use a modified Hessian matrix, H' , based on second-order derivatives, with a tunable parameter, α , to calculate a measure of neuriteness such that,

$$N = \begin{cases} \frac{\lambda_{\max}}{\lambda_{\min}} & \text{if } \lambda_{\max} < 0, \\ 0 & \text{if } \lambda_{\max} \geq 0, \end{cases} \quad (3.5)$$

where

$$\lambda_{\max} = \max(|\lambda_1'|, |\lambda_2'|), \quad (3.6)$$

$$\lambda_{\min} = \min(\lambda_{\max}), \quad (3.7)$$

$$\lambda_1' = \lambda_1 + \alpha\lambda_2, \quad (3.8)$$

$$\lambda_2' = \lambda_2 + \alpha\lambda_1, \quad (3.9)$$

where λ_i' are the normalized eigenvalues of H' ; λ_{\max} is the largest eigenvalue at each pixel; and λ_{\min} the smallest value of all λ_{\max} [46].

The neuriteness measurement enhances bright, neurite-like structures whilst ignoring dark, line-like structures. Further, background intensity discontinuities that are immune to first order derivatives are suppressed by the use of second order derivatives.

Regularized Volume Ratio

Recently, a new Hessian-based vessel enhancer was introduced in [13]. This approach attempts to resolve the drawbacks found in most previous Hessian-based approaches, which are directly proportional to λ_2 (2D and 3D) or λ_3 (3D only). These drawbacks are a) the response of $\lambda_{2,3}$ is non-uniform within an ideal elongated or rounded structure of uniform intensity, b) the magnitude of $\lambda_{2,3}$ varies with image intensity, and c) enhancement is not uniform across scales. Jerman et al. [13] attempt to solve this by modifying the volume ratio to indicate elongated structures, such as vessels, and regularising λ_3 (where $\lambda_3 = \lambda_2$ in 2D) to ensure robustness to low magnitude changes, *e.g.* noise in regions of uniform intensity. Their final solution in 2D is given by

$$\mathcal{V}_P = \begin{cases} 0 & \text{if } \lambda_2 \leq 0 \vee \lambda_\rho \leq 0, \\ 1 & \text{if } \lambda_2 \geq \lambda_\rho/2 > 0, \\ \lambda_2^2(\lambda_\rho - \lambda_2) \left[\frac{3}{\lambda_2 + \lambda_\rho} \right]^3 & \text{otherwise,} \end{cases} \quad (3.10)$$

where λ_ρ is the regularised form of λ_2 , given by

$$\lambda_\rho(s) = \begin{cases} \lambda_2 & \text{if } \lambda_2 > \tau \max_{\vec{p}} \lambda_2(\vec{p}, s), \\ \tau \max_{\vec{p}} \lambda_2(\vec{p}, s) & \text{if } 0 < \lambda_2 \leq \tau \max_{\vec{p}} \lambda_2(\vec{p}, s), \\ 0 & \text{otherwise,} \end{cases} \quad (3.11)$$

for scale s , where τ is a cut-off threshold between zero and one and for any pixel $\vec{p} = (x, y)$. Throughout this paper we refer to this approach as the regularised volume ratio (RVR).

3.2.2 Phase Congruency-based Enhancement Methods

One problem with many image enhancement methods is that they depend, to some extent, on image intensity and therefore fine, and usually lower contrast, vessels may be missed. The development of a contrast-independent image enhancement measurement has been shown in [43]. This approach builds upon the idea of phase congruency: a measurement that looks for features at points where Fourier phase components align [52].

The original measurement is sensitive to noise, one dimensional and not a robust detector of features; this has been extended by Kovési and colleagues [53–56] into a 2D phase congruency approach that combines both sine and cosines phase elements, improving the localisation information and noise robustness[‡].

Phase Congruency Tensor

A contrast-independent, tensor-based measurement of local image behaviour, the phase congruency tensor (PCT), was introduced in [43] as a detector of curvilinear structures regardless of intensity variation and inhomogeneous contrast issues. For a given set of scales s and a given set of phase congruency measures $PC_s(\vec{p}, \theta)$ for each orientation θ , the proposed phase congruency tensor is,

$$T_{PC_s}(\vec{p}) = \sum_{\theta} PC_s(\vec{p}, \theta) (\vec{n}_{\theta} \vec{n}_{\theta}^T), \quad (3.12)$$

where \vec{n}_{θ} is the normalised column vector for each orientation and $\alpha = 1/N - 1$, with N is the image dimensionality.

The eigenvalues of the tensor can then be used in the same way as the Hessian matrix eigenvalues are used in equations (3.3) and (3.5) to define PCT Neuriteness and PCT Vesselness.

3.2.3 Contrast Limited Adaptive Histogram Equalisation

A somewhat simpler approach in image enhancement is to carry out contrast enhancement by histogram equalisation. A popular approach in this category is contrast limited adaptive histogram equalisation (CLAHE) [44].

The image is first divided into small regions, each of which will undergo histogram equalisation (thus making the algorithm adaptive); however, this would also massively enhance noise and, to prevent this, contrast limiting is applied. Contrast limiting identifies any original histogram bin of the region with more than a set number pixels above a threshold and ‘clips’ those pixels, distributing them evenly to the rest of the bins, histogram equalisation can then be carried out on the clipped region. Most implementations will also bilinearly interpolate to remove artefacts at the borders between regions.

[‡] The mathematics of phase congruency are described in more technical detail in appendix B (p. 145).

3.2.4 Matched and Steerable Filter Enhancement Methods

Matched filtering is an enhancement or detection approach where a set of filters matching key features of the objects of interest, such as vessel, are convolved with the image to indicate the presence and location of such filters. Steerable filters dimensionally reduce this approach and are filters with an orientation that can be formed of a linear combination of basis filters. Such directional filters have been used to obtain features such as lines, edges or corners [57–59].

Matched filters require discretization of the orientation for matching and good results require high angular resolution; however, finer resolution comes with increased computational costs. Steerable filters offer an alternative way by reducing the number of orientations required to the basis vectors [60]; however, such a reduction in orientations can lead to slightly reduced accuracy.

3.2.5 Enhancement with Mathematical Morphology

An alternative approach to image enhancement is the use of mathematical morphology. Morphological operations are a set of non-linear filtering methods formed through combination of two basic operators: dilation and erosion[§]. Where we describe dilation, \oplus , for a point in a grayscale image, $I(\vec{p})$, as the maximum of the points in the weighted neighbourhood described by the structuring function $b(\vec{p})$ or, mathematically, as

$$(I \oplus b)(\vec{p}) = \sup_{\vec{q} \in \mathbb{R}^2} (I(\vec{q}) + b(\vec{p} - \vec{q})), \quad (3.13)$$

where sup is the supremum and $\vec{q} \in \mathbb{R}^2$ denotes all points within the image. Likewise, we describe the erosion, \ominus , as the minimum of the points in the neighbourhood described by the structuring function or, mathematically as

$$(I \ominus b)(\vec{p}) = \inf_{\vec{q} \in \mathbb{R}^2} (I(\vec{q}) - b(\vec{q} - \vec{p})), \quad (3.14)$$

where inf is the infimum. Dilation is able to expand bright areas and reduce dark areas, whilst erosion expands dark areas reducing bright areas.

[§] In this chapter we give a brief overview of mathematical morphology, however, a full description is provided in chapter 2 (p. 10)

From these two operators we can define four commonly used morphological filters:

$$\text{opening : } (I \circ b)(\vec{p}) = ((I \ominus b) \oplus b)(\vec{p}) \quad (3.15)$$

$$\text{closing : } (I \bullet b)(\vec{p}) = ((I \oplus b) \ominus b)(\vec{p}) \quad (3.16)$$

$$\text{top-hat : } TH(I, b) = I(\vec{p}) - (I \circ b)(\vec{p}) \quad (3.17)$$

$$\text{bottom-hat : } BH(I, b) = I(\vec{p}) - (I \bullet b)(\vec{p}). \quad (3.18)$$

Where an opening will preserve dark features and patterns, suppressing bright features, and a closing will preserve bright features whilst suppressing dark patterns. Top-hat and bottom-hat transforms will enhance bright and dark regions, respectively, where those patterns correspond to the size and shape of the structuring element used.

Mathematical morphology has long been used for image enhancement including noise suppression, feature detection and contrast sharpening. Traditionally, the top-hat transform is used to identify ‘peak’ like signals. Modifications of the top-hat approach have been used in, for example, pattern classification [61], enhancing depth-of-field [62] and edge detection [63].

The Top-Hat Transform

Mathematical morphology has been used by many researchers to enhance and segment vessel-like structures [38, 45, 48, 64–66]. One popular approach is that put forward by Zana and Klein [38] where they combine morphological transforms and cross-curvature evaluation to segment vessel-like objects. This relies on the assumption that vessels are linear, connected and have smooth variations of curvature along the peak of the line. Their algorithm combines a sum of top-hats, calculated using linear structuring elements at different angles, which reduces noise and improves contrast, and computation of the curvature, using a Laplacian of Gaussian approach, to further reduce noise and enhance vessel-like signals.

3.2.6 Limitations and Challenges

Due to the nature of many vessel enhancement approaches, *i.e.* the suppression of disk-like objects, many existing approaches lose intensity at junctions and some bends which are more locally similar to disk-like elements than extended vessel-like objects. Contrariwise, those methods that have been designed not to lose signal at these features often see vessel-like features everywhere in the image, including in noisy regions, and thus enhance ‘vessels’ that are not there. Our

approach, described in the next section neither suppresses junctions and bends nor enhances noise as vessel-like structures. This helps to maintain the connectivity of networks and key feature needed for visual and automated analyses in both medical and biological images.

3.3 Methods

Our proposed method, summarised in figure 3.3, aims to solve the problems of junction/bend suppression and noise enhancement when attempting to enhance the signal of vessel-like objects in biomedical images. We achieve this by combining the outputs of morphological operations upon the input image carried out with two different banks of structural elements: one bank of disk elements (figure 3.4 (p. 32)), with varying radii, and one bank of line elements (figure 3.5 (p. 32)), with varying radii and rotation. These banks of images are then combined in a multiscale manner before being integrated into a single output image (figure 3.6 (p. 33)). In this section we introduce our proposed method, the bowler-hat transform, highlighting the key concepts that allow this approach to deal with the challenges such as those found in retinal imaging.

3.3.1 Proposed Method

For a given greyscale input image, I , we carry out a series of morphological openings with disk-shaped structuring elements, b_s , of size $s \in [1, s_{\max}]$, where s_{\max} is the manually chosen expected maximum vessel size. The result for each s is an image, I_{disk} , in which vessel segments wider than s remain and those segments smaller than s are attenuated (figure 3.4 (p. 32)). We produce a stack of such images for all s , such that

$$\{I_{disk}\} = \{I \circ b_s : \forall s \in [1, s_{\max}]\}. \quad (3.19)$$

We also produce a similar stack of images using a bank of line structural elements, $b_{s,\theta}$; each line is of length $s \in [1, s_{\max}]$, has a width of 1 pixel, and of orientation $\theta \in [0, 180)$. The resultant image for a line of length s and orientation θ is that all strips on an object that are longer than s along angle θ will remain and those strips shorter than s along angle θ will be attenuated (figure 3.5 (p. 32)). For each s we produce a stack of images for all θ and then take the pixel-wise

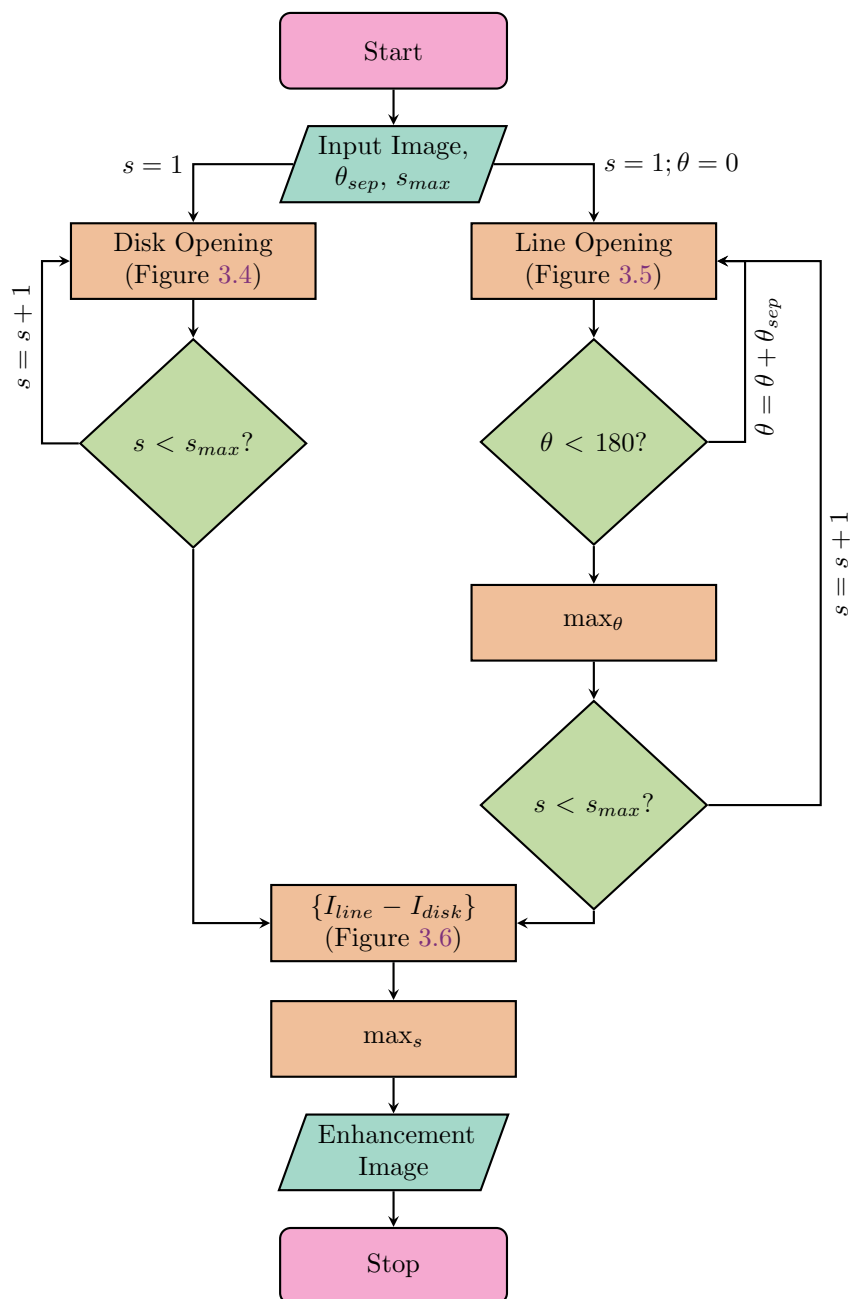


Figure 3.3: Flow chart of the bowler-hat transform showing how two banks of images are created using disk SEs (left route; figure 3.4) and line SEs (right route; figure 3.5) of different scales and, for the line elements, rotations. These banks are then combined in a multiscale fashion (figure 3.6) before forming a single enhancement image.




Figure 3.4: Example openings with disk SEs of various sizes. The set of all openings with disks of differing size represents the left route in figure 3.3

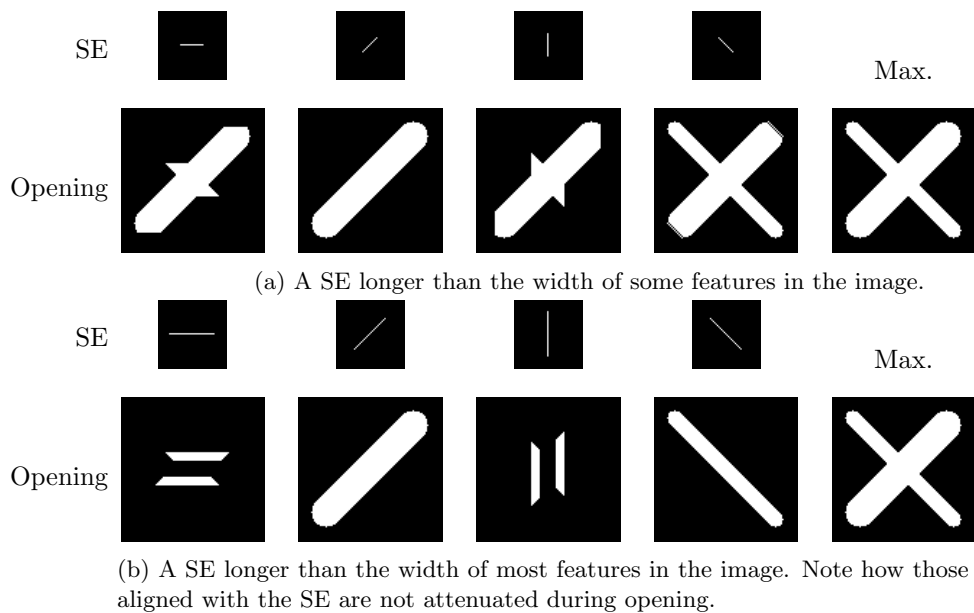


Figure 3.5: Example openings with line SEs of various lengths and rotations. The set of all openings with lines of differing length and rotation represents the left route in figure 3.3. In this extremely simple case the maximum (right hand column) includes values at all features.

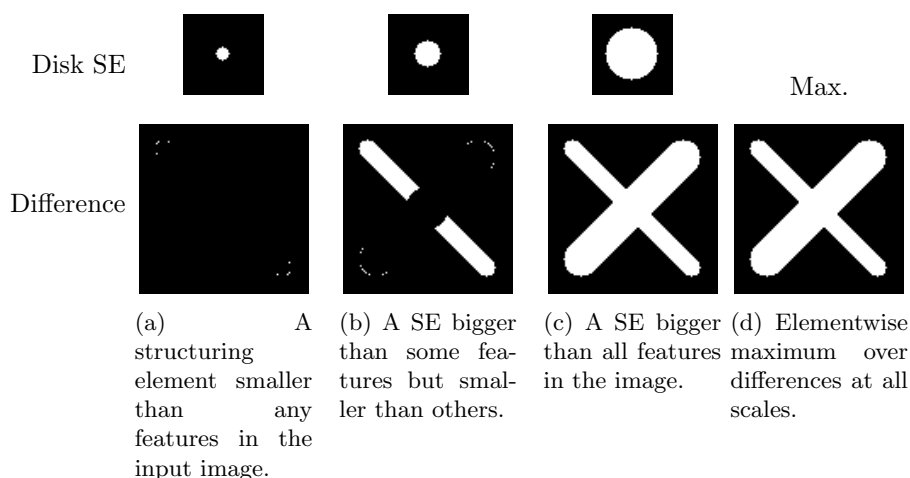


Figure 3.6: Example differences between disk openings and maximum (over orientation) line openings of various sizes. The set of all differences represents the coming together of the two routes in figure 3.3. Disk SEs are shown to give an idea of scale. The maximum (far right) represents the enhanced image.

maximum to give a single image, I_{line} , for each s ; these resultant images then produce a stack with one image for each value of s , such that

$$\{I_{line}\} = \left\{ \max_{\theta} (\{I \circ b_{s,\theta} : \forall \theta\}) : \forall s \in [1, s_{\max}] \right\}. \quad (3.20)$$

These two stacks are then combined by taking the stack-wise different, *i.e.* the difference between the maximum opening with a line of length s across all angles and an image formed of opening with a disk of size s , to form an enhanced image at all scales (figure 3.6). Our enhanced greyscale image is then formed from the maximum difference at each pixel across all stacks,

$$I_{enhanced} = \max_s (\{I_{line} - I_{disk}\}). \quad (3.21)$$

Pixels in the background, *i.e.* dark regions, will have a low value due to the use of openings; pixels in the foreground of blob-like objects will have a low value as the differences will be minimal, *i.e.* similar values for disk-based and line-based openings; and pixels in the foreground of vessel-like objects will have a high value, *i.e.* large differences between longer line-based openings and disk-based openings.

The comparison of line and disk elements gives our proposed method a key advantage over existing methods, given an appropriate s_{\max} , *i.e.* larger than any vessels in the image, a junction should appear bright like those vessels joining that junction, something that many other vessel enhancement methods fail to do. This is due to the ability to fit longer line-based structural elements within the junction area. The result of this improvement is that the network of vessels in an image stays connected when thresholded, if junctions were not enhanced this connectivity is lost.

In section 3.4 we demonstrate the key advantages of the bowler-hat transform over other vessel enhancement approaches.

3.3.2 Implementation and Data Sets

To evaluate our proposed method we show experiments of synthetic and real data sets. All synthetic images have been generated and modified, *i.e.* added noise, in MATLAB. We have shown the quality of the proposed method using two, common retinal imaging data sets: the DRIVE and STARE databases. These datasets have been chosen because of their availability and their ground truth. We have used these ground truth segmentations to quantitatively compare our method with other vessel enhancement approaches.

All codes were implemented and written in MATLAB 2016b [67] on Windows 8.1 pro 64-bit PC running an Intel Core i7-4790 CPU (3.60 GHz) with 16GB RAM. All codes and synthetic data is available in a GitHub repository ([68]).

3.4 Results

In this section we demonstrate and quantify the enhancement results of known synthetic data, including example junctions and noisy data, and real data, with human-annotated ground truths, when enhanced with the proposed method; Hessian-based vesselness [41] & neuriteness [69]; Phase Congruency Tensor-based vesselness & neuriteness [43]; CLAHE [44]; Regularised volume ratio [13]; and a morphological approach using a sum of top-hats [38].

3.4.1 Synthetic Data

As with any method, an understanding of how the parameters involved affect the result is essential. Throughout our experiments we have manually tuned the bowler-hat transform and the comparator methods. In general we have found the bowler-hat transform to be robust against parameter changes above around 10–12 orientations and with a scale a few pixels greater than the largest expected object thickness.

In figure 3.7 we plot the enhanced responses of the proposed and comparator methods on a synthetic vessel-like object. As can be seen, whilst the Hessian-based approaches have an enhanced signal at the center of the vessel, *i.e.* a peak value of one at the vessels centre-line, their value drops off quickly and decreases the perceived thickness of the vessel. Contrariwise, Phase Congruency Tensor methods do not necessarily peak at the vessel centre but their response does not drop off quickly and they maintain a higher response through to the edges of the vessel, *i.e.* the perceived thickness is not decreased. Our proposed method has both these benefits: a maximal peak value at the vessel centre-line and an enhanced response till the edges of the vessel, so reliable vessel thicknesses can be captured.

Comparison to Existing Approaches

Figure 3.8 compares the enhanced results of the proposed method and several existing approaches for comparison. Key issues that occur across the comparator methods include defects at junctions (purple arrows), noise enhancement, tip artefacts (orange arrows) and loss of signal (yellow arrows). These issues are all absent in our proposed approach.

However, a key negative to our approach is shown in figure 3.8 row 4, which shows a vessel like object with an attached ‘blob’ (green arrow), a perfect vessel enhancement method would enhance all of the linear structure and none of the blob. Whilst none of the comparison approaches act in this ideal manner many of them show a clear difference between the blob response and vessel response, our proposed method shows some difference but this difference is much less obvious when, as here, the blob touches a vessel. This is due to a blob touching another object appearing as an elongated blob, *i.e.* a vessel.

Response to Noise

Figure 3.9 shows the effect of three different noise types on the proposed and comparison methods. Given that the proposed approach has no built-in noise suppression or cancelling, it is unsurprising

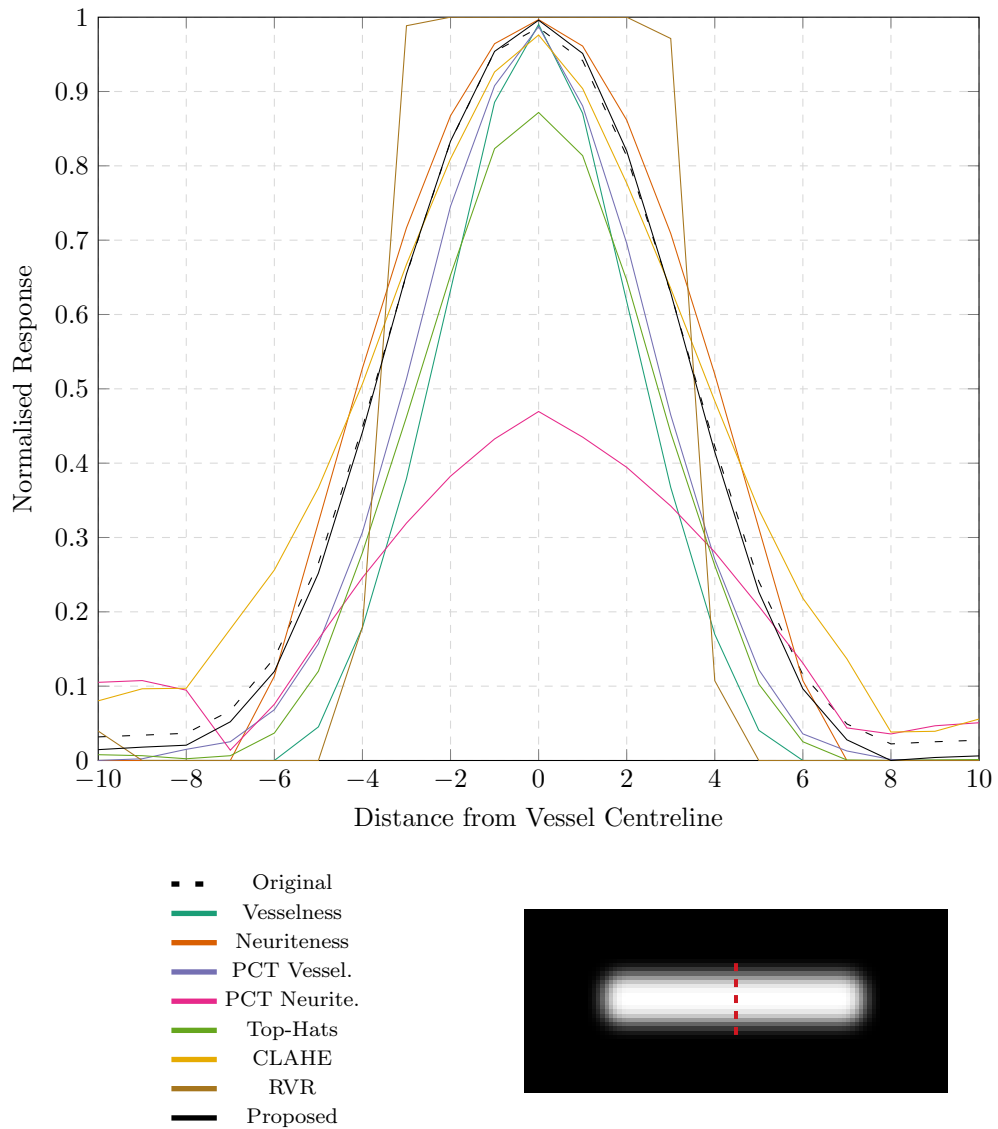




Figure 3.7: Cross-sectional profile of original synthetic vessel (black, dashed line), the enhanced image under the proposed method (grey, solid line) and the enhanced images under the comparator methods (see legend for colours). All images were normalised such that the brightest pixel in the whole image has a value of 1 and the darkest a value of 0.



(a) Original (b) Proposed (c) Vesselness (d) Neurite. (e) PCT V.



(f) PCT N. (g) Top-Hats (h) CLAHE (i) RVR

Figure 3.8: A comparison of the enhancement of vessel-like structures and other features for the proposed and comparator approaches. (a) shows the original images, all vessels have a thickness of nine pixels and the ‘blob’ in 4 has a diameter of 21 pixels. (3.8b) shows the results of the proposed method, whilst (c–i) show the enhanced images for the given comparator methods. Arrows indicate features of interest: vessel structures (yellow arrows), junctions (purple), blob-like features (green) and tips (orange).

that the effect of noise on the enhanced image is in-line with the original image. We note that the method is weakest in response to speckle noise, which follows from the noise-sensitivity in morphological operations and should be taken into consideration when choosing an enhancement approach.

Response to Uneven Background Illumination

In figure 3.10 we show the response of the proposed method under cases of uneven illumination. Key features such as junctions are preserved and appear unaffected by even severe illumination problems. This ability to preserve junctions under uneven illumination is important for many real applications of vessel enhancement and the proposed method is able to do this, unlike the comparator approaches.

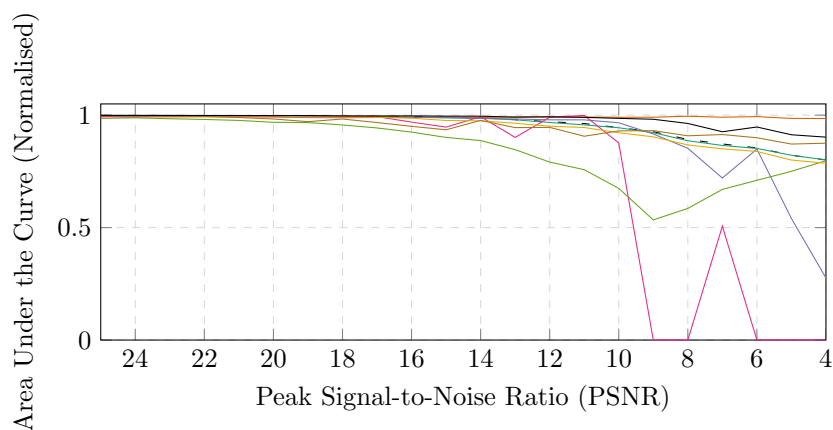
3.4.2 Real Data

In this section we provide quantitative and qualitative comparisons of the proposed method and comparators. We show the quality of the proposed method validated on two retinal imaging data sets: the DRIVE and STARE databases. These datasets have been chosen because of their availability and their ground truth data. We have used these ground truth segmentations to quantitatively compare our method with other vessel enhancement approaches.


Datasets

The DRIVE [70] dataset is a published database of retinal images for research and educational purposes. The database consist of twenty colour images that are JPEG compressed, as for many screening programs. These images were selected randomly from a screening of 400 diabetic subjects between the ages of 25 and 90. The ground truth provided with this dataset consists of a manual segmentation of the vasculature for each image. Ground truths were prepared by trained observers and ‘true’ pixels are those for which observers where $> 70\%$ certain.


The STARE dataset is another publicly available database [47] and within the full dataset contains twenty colour images with human-determined vasculature ground truth. We have compared against the AH labelling for all images.



(a) Additive Gaussian Noise



(b) Multiplicative Gaussian (Speckle) Noise



(c) Salt and Pepper Noise

■ Original ■ Vesselness ■ Neuriteness ■ PCT Vessel. ■ PCT Neurite.
 ■ Top-Hats ■ CLAHE ■ RVR ■ Proposed

Figure 3.9: Mean Area Under the Curve (AUC) for the original image (black, dashed line), the enhanced image under the proposed method (grey, solid line) and the enhanced images under the comparator methods with different peak signal-to-noise ratios (PSNRs) for three different noise types: (a) additive Gaussian noise, (b) multiplicative Gaussian noise (speckle noise), and (c) salt and pepper noise.

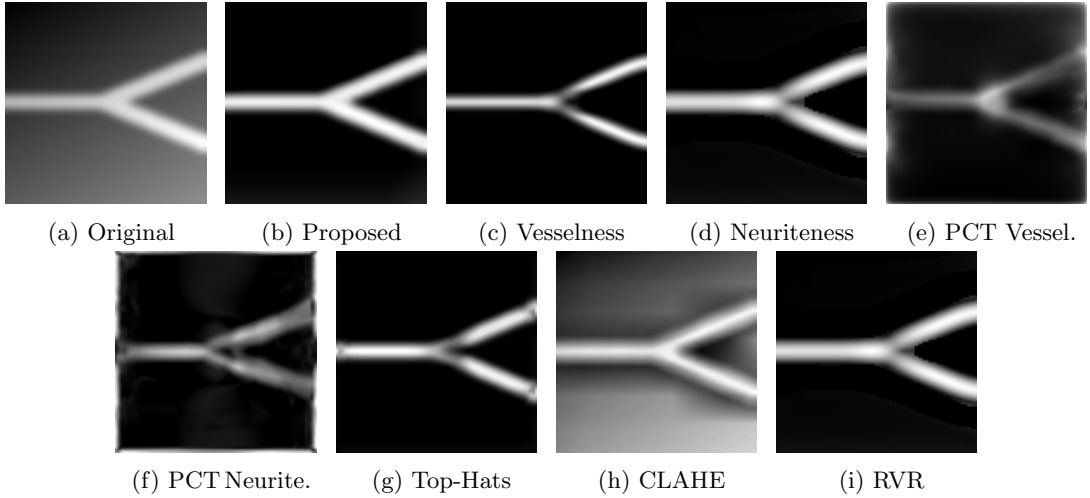


Figure 3.10: Comparison of methods' abilities to deal with uneven background illumination. (a) The original image, (b) the proposed method, (c-i) the comparator methods.

Quantification of Enhancement

Whilst visual inspection can give much information regarding the effectiveness of a vessel enhancement algorithm, a form of comparable quantification allows for a more objective comparison of methods. In order to compare our method with those approaches mentioned above we have chosen to calculate the Receiver Operating Characteristic (ROC) curve, using the sensitivity and specificity of thresholded enhanced images, and the AUC. The sensitivity and specificity are defined as:

$$\text{sensitivity} = \frac{TP}{TP + FN} \quad , \quad (3.22)$$

$$\text{specificity} = \frac{TN}{TN + FP} \quad , \quad (3.23)$$

where TP is the true positive count, FP the false positive count, TN the true negative and FN the false negative counts of the segmented pixels after globally thresholding the enhanced images with different greyscale thresholds values (see figure 3.11 for examples of each). These metrics take in an image, either the original or an enhanced image, and the ground truth annotated data, available with each data set. A higher AUC value indicates a better enhancement, with a value of 1 indicating that the enhancement image is identical to the binary, ground truth data for all thresholds.



(a) Region of image enhanced with the bowler-hat transform and globally thresholded at 0.5.

(b) Same region of human-annotated ground truth.

(c) True positive (red) and true negative (blue) pixels; false positive (green) and false negative pixels (yellow).

Figure 3.11: Examples of true positive (red pixels in (c)); true negative (blue pixels in (c); false positive (green pixels in (c); and false negative (yellow pixels in (c) as described in the main text. (a) shows the region of interest in a retinal image that has been enhanced with the bowler-hat transform and globally thresholded at a value of 0.5. (b) shows the manually annotated ground truth for that region.)

Enhancement Method	AUC (StDev)	
	DRIVE	STARE
Original	0.518 (0.064)	0.481 (0.076)
Proposed	0.921 (0.032)	0.926 (0.034)
Vesselness	0.552 (0.243)	0.722 (0.215)
Neuriteness	0.864 (0.022)	0.889 (0.039)
PCT Vessel.	0.826 (0.037)	0.843 (0.056)
PCT Neurite.	0.674 (0.121)	0.649 (0.165)
Top-Hat	0.915 (0.023)	0.945 (0.037)
CLAHE	0.455 (0.068)	0.639 (0.087)
RVR	0.907 (0.024)	0.907 (0.042)

Table 3.1: Mean AUC values for the proposed and comparator methods across the DRIVE and STARE datasets. Best results for each dataset are shown in bold. Individual examples of ROC curves can be seen in figure 3.13.

Comparison of Results

Figure 3.12 shows the results of the proposed and comparator methods on an example image from the STARE dataset. We can see that the proposed method seems able to enhance finer structures as detected by the human observer but not emphasised by many of the comparator methods (see arrows). We can also see that, whilst the connectivity seems to be maintained (unlike in figure 3.12d), false vessels are not introduced (*c.f.* figure 3.12e). However, we do note that there is no noise suppression in the proposed approach and this has led to the enhancement of a lot of background noise, similar to figure 3.12g.

Whilst one would expect the lack of noise suppression to be a major issue with regard to quantified approaches to vasculature segmentation, we find that the proposed method gives the best enhancement of all approaches on the DRIVE dataset and second best for the STARE dataset (see table 3.1 and figure 3.13). We believe this result to be due to vessels that are not captured by the human annotator, which can be seen in figure 3.12 (row 2, yellow arrows) where the proposed method clearly highlights true vessel that are not captured by either the human or the sum of top-hats approach.






Figure 3.12: An example image from the STARE dataset (see main text). The zoomed in regions (2 and 3) show areas both near the optic disk, with thick vasculature, and far from the disk, where vessels are much fine and more branching. The arrows point to key areas of interest, such as junctions, fine tips and vessels not captured by all approaches.



(a) DRIVE



(b) STARE

- - Original — Vesselness — Neuriteness — PCT Vessel. — PCT Neurite.
— Top-Hats — CLAHE — RVR — Proposed

Figure 3.13: ROC curves for examples images of the DRIVE (a) and STARE (b) datasets. Mean AUC values can be found in table 3.1.

Other Biomedical Examples

Although we've shown the enhancement ability of the proposed method on retinal imaging data, this approach will also work on a variety of other biomedical images. Figure 3.14 shows just some examples on microscopy images.

3.4.3 Limitations

As with any approach, the proposed method has some limitations. Firstly, morphological operations are renowned for their large computational requirements. However, whilst our approach is generally slower than the comparator we have found this to be a minimal difference of the order of a few seconds.

3.5 Conclusion

This paper proposes a new method for vessel enhancement in medical and biological images based on mathematical morphology and assumptions of a vessel's shape. The proposed method is shown to have benefits over existing methods, including no loss of signal at junctions and minimised artefacts at vessel ends. We show and quantify this effectiveness on both synthetic and real data.

Future extensions of this work will include the development of a three dimensional equivalent, exploration of blob-enhancing variants of this approach and the analysis of parameter sensitivity for different modalities. For now, the work presented in this paper introduces an easy to use alternative for vessel enhancement that can be used in a wide range of imaging scenarios.

Epilogue

In this chapter I have introduced a new mathematical morphology transform, the bowler-hat transform, which has shown to be a successful vessel enhancement algorithm with results as good or better than the most common approaches for a variety of synthetic and real images.

Image enhancement is a key task in the image analysis pipeline, often followed by feature or object detection and measurement. In the remaining chapters, I introduce improved and new approaches to the next stages in the pipeline. Starting with chapter 4 (*p.* 48) I introduce a modified small particle detector for 3D and then, in chapter 5, I describe a novel ellipse detector for 2D and

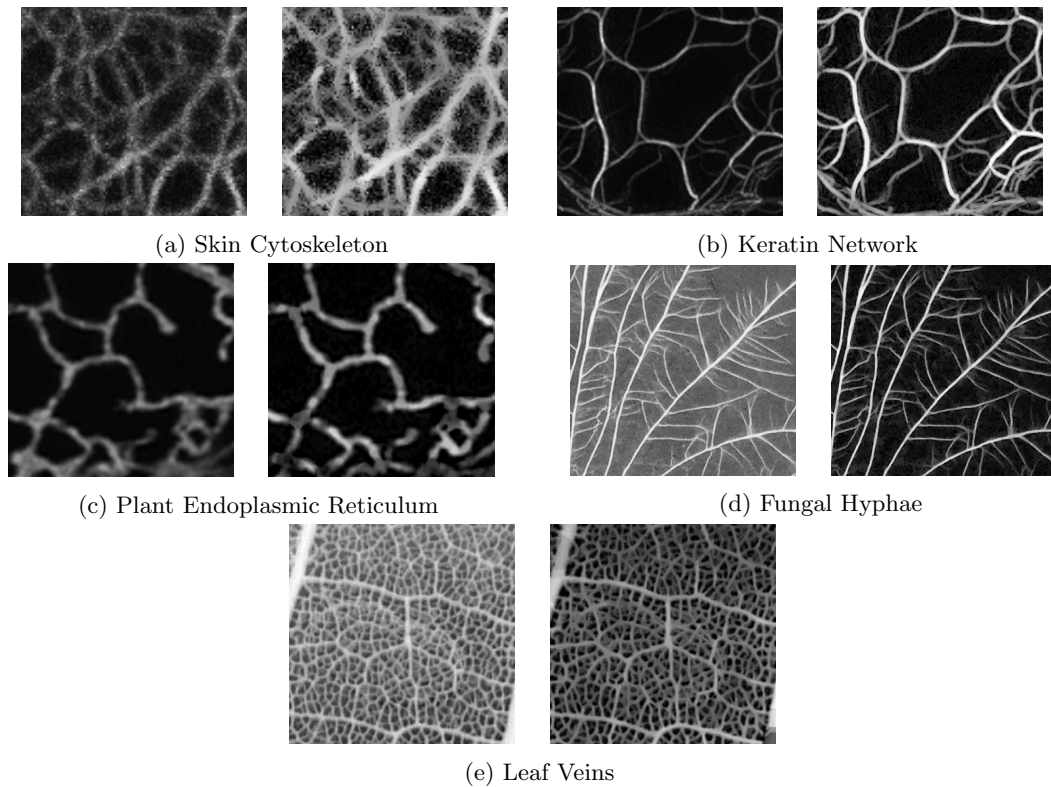


Figure 3.14: As well as medical applications, our proposed enhancement method works in a variety of biological scenarios including cytoskeletal networks (a–b), endoplasmic reticulum (c) and macro-scale networks (d–e). Left image shows original data, right shows data after enhancement by the bowler-hat transform. (a) provided by Prof. Rudolf Leube, RWTH Aachen University, Germany. (b) provided by Dr Tim Hawkins, Durham University, UK. (c–e) provided by Prof. Mark Fricker, Oxford University, UK

its application to nuclei detection and counting in fluorescence microscopy. Finally, chapter 6 (p. 112) introduces the use of mathematical morphology for measurement of segmented objects, in particular, vascular systems and abnormalities.

Chapter 4

Blobs and Curves: Object-based Colocalisation for Plant Cells

Prologue

In chapter 3 (p. 20) I have investigated and developed a novel approach to the first stage of many image analysis pipelines: image enhancement. In the following chapter, I move onto the next task in the image analysis pipeline: that of object detection, specifically ellipse detection. I introduce the problem of small particle detection and colocalisation in fluorescence microscopy, as applied to a case study of plant cell biology.

In this chapter, I introduce the field of bioimage informatics in the context of plant cell biology along with the background biological context of the case study. I then go on to describe, in some detail, all the algorithms used in this analysis pipeline including the introduction of a new local intensity pruning solution for over detection. Here I have focused on understandability for non-expert readers, highlighting key areas of concern surrounding uncertainty and error.

Finally, I demonstrate how particle detection and segmentation can be used to quantify a classically qualitative biological assay: colocalisation before drawing biologically relevant conclusions.

Declaration: This chapter is based on the following publication: Nelson, C. J., Duckney, P., Hawkins, T. J., Deeks, M. J., Laissue, P. P., Hussey, P. J. & Obara, B. Blobs and curves:

object-based colocalisation for plant cells. *Functional Plant Biology* **42**, 471–485 (2015). This chapter is presented as submitted, although referencing and notation has been altered and cross-referencing added for consistency throughout this thesis. Some stylistic changes have been made for consistency. The majority of the text is verbatim; however, additions to the body of text are included along with footnotes and, where appropriate, added appendices.

4.1 Introduction

In the past 20 years, a revolution in plant science has occurred: the age of genomics. From the development of genome sequencing technologies to the bioinformatics tools used to store, manage and analyse the data (see [71] for a recent example of the tools available), a whole new interdisciplinary field of research has emerged, supported by accessible tools and dedicated international journals. Unfortunately with this ‘gold rush’ of genomic tools, techniques and understanding, a classic tool of biology has been left behind: phenotyping. Phenotypic analysis has been essential from the days of Mendel but is now considered the bottleneck of high-throughput biological sciences [72]. A lack of automation in phenotype analysis and indeed a lack of investment in the necessary programming skills within the plant science community are slowing an informatics-driven enlightenment.

With the advent of digital imaging and ever more powerful computers, microscopy as a tool for phenotyping has moved from qualitative observation to quantitative analysis of biological samples and assays (see [73] for a review). Exciting developments in the world of microscopy (see [74] for a review) have brought about a dramatic increase in spatial and temporal resolution, producing massive image files that hold a great amount of data. We can now see the entrance of bioimage informatics, an area where biological study, bioimaging methods and computer vision analysis meet to manage and analyse data from a wide variety of experiments [75].

In plant sciences, this alliance is often hidden under the term ‘phenotyping’; the term ‘image analysis’ only mentioned when explicit computational methods are used (*e.g.* [72]). This can lead to problems of reproducibility and comparison of results [76]. In this paper, we will outline the state of bioimage informatics in plant sciences before showcasing two examples of fluorescence colocalisation; see [14] for an early example of quantitative colocalisation. These examples show how bioimage informatics technology can be applied across experiments and systems, and also highlight possible pitfalls in designing and using bioimage informatics tools. We have chosen to demonstrate bioimage informatics tools on colocalisation, as it is a commonly used technique

with a broad impact and is often used to infer biological function with little quantification. Explicitly quantifying such results can increase this impact and provide researchers with more information from a commonly used assay [77]. The cases we have laid out in this paper also highlight a couple of key considerations when quantifying results in this way.

4.1.1 Bioimage Informatics in the Plant Sciences

Although the term ‘bioimage informatics’ is often used in relation to cellular and molecular biology, it is primarily the management, analysis and storage of large-scale datasets from images within the life sciences (see [78] for a review of the computational involvement in imaging and optical technologies). In plant sciences, this can encompass anything from the canopies of forests to the colocalisation of protein signals within leaf cells. Many bioimage informatics tools are, of course, transferable across all of the biological sciences; examples include the Open Microscopy Environment Remote Objects [79] and Bio-Image Semantic Query User Environment [80] platforms, open-source management and access software for a wide range of biological data. Indeed, new features are being developed that add functionality to such systems (*e.g.* [81]).

More specifically to plant sciences, there are now collaborative groups, such as the iPlant Collaborative ([82], aiming to develop ‘cyberinfrastructure and computational tools’ within plant sciences. Such groups aim to help collaboration and the development of a wide range of high-quality tools for analysis, inference, management and prediction among plant scientists. Further, on-line collections of bioimage informatics tools are being created to help plant scientists find the appropriate tools without having to sift through a plentitude of academic journals [83].

[83] show that their database has software that performs bioimage informatics from the canopy level through to the leaf, root and shoot, right through to cellular analysis. Algorithms range from cell wall segmentation [84] and cell tracing [85] to leaf disease identification, species identification, root network mapping and canopy coverage. Network extraction methods are particularly popular, from the extraction of fungal networks [86] to connected plant cell walls [87] and leaf venation [88, 89]. Such network extraction systems can often be adjusted for different scales and applied to macroscopic fungal or root networks, as well as sub-cellular filament networks.

4.1.2 Challenges in Bioimage Informatics

Many available bioimage informatics software packages give similar measurements as their final outcome but it is up to the individual experimenter to discover which algorithm produces the

most reliable results for their own experimental data. It is often the case that one program is more suited to one modality or one scenario. Scientists in biological laboratories require software solutions that are tailored to their specific needs, so collaborations with computer scientists can often be the most efficient way forward in developing new tools. These tools should be well documented, such that the developed algorithm and approach can be replicated and understood by those trying to validate or contest the results of the experiments [90].

From the computer science perspective, there are also the challenges of developing user-friendly software while still maintaining scientific openness within the code. The challenges regarding bioimage informatics software usability have recently been reviewed, identifying many different criteria [91]. These criteria include user-friendliness alongside the ability to adapt the software for new uses, the interoperability of the software (*i.e.* the ability to use different languages, operating systems, file formats and data formats) and the validation of the software at different levels.

Validation of software (*i.e.* the process of proving that software-produced results both accurately and precisely represent the real data) is particularly important. If an experimental scientist takes to using a particular set of computational tools, that scientist must be aware of how the tools have been validated and whether the criteria assumed for that validation hold true for this experiment's data [90]. To aid such validation, new tools are being developed in the bioimage informatics sphere, with repositories of synthetic images (*e.g.* [92]), fully annotated images (*e.g.* [83]) or segmentation benchmark infrastructures (*e.g.* [93]).

4.1.3 Colocalisation of NET1A with Plasmodesmata

Plasmodesmata are the numerous channels that provide routes of intercellular transport and communication between the cells of multicellular plants (see [94] for a review; figure 4.1). Retraction of the plasma membrane caused by osmotic shock can reveal their presence through the pinched strands of the protoplast running through their centre, thus making them visible to early botanical microscopists. In the root meristem of *Arabidopsis thaliana* (L.) Heynh., plasmodesmata are only 30 nm in diameter but accommodate a complex membrane-bound structure of protein scaffolds and a tubular endoplasmic reticulum. Assignment of protein localisation to plasmodesmata using fluorescence microscopy is achieved using co-labelling with either antibodies or fusions to known plasmodesmal components, or the staining of plasmodesmal-associated callose in the surrounding cell wall (*e.g.* [95]). At a standard optical resolution, the

location of callose accumulation is often indistinguishable from the plasmodesmal mouth and channel (*e.g.* [95]). The majority of publications demonstrating positive colocalisation between known pore components and uncharacterised proteins use transient expression assays in plant tissues where plasmodesmata exist as isolated and distinctive entities; for this reason, complex colocalisation analysis is not required. In tissues such as the *A. thaliana* root meristem, the pattern of plasmodesmata is condensed, with neighbouring pores clustering at distances close to the diffraction limit of resolution (*e.g.* [95, 96]). The segregation of root meristem cell boundaries into distinctive nano-domains of protein localisation, which are not necessarily connected to plasmodesmal distribution ([97]), further complicates analysis.

In this study, we re-analyse a published colocalisation study of the actin-binding protein probe, a fusion of NETWORKED protein 1A (NET1A) and green fluorescent protein (GFP), with labelled root meristem plasmodesmata ([95]) using object-based algorithms (*i.e.* algorithms that identify distinct objects within an image (*e.g.* [15])) rather than using the statistical properties of the image as a whole. We show that this method provides a quantified overview of the spatial relationship between NET1A-GFP fusion protein and aniline blue signals in units of distance rather than an abstract coefficient. We also demonstrate that this method is effective at removing the influence of background signals inherent to deep-tissue plant fluorescence microscopy.

4.1.4 Colocalisation of Mitochondria with Actin Filaments

The pollen tube is a model example of anisotropic plant cell expansion, in which growth occurs only at the tip [98]. Polarised growth is dependent on the tightly regulated spatio-temporal coordination of the pollen tube growth machinery by the actin cytoskeleton [99]. The actin cytoskeleton facilitates the transport of biosynthetic materials to the growing tip, including cellulose synthase and callose synthase, both trafficked in Golgi-derived bodies [100]. Actin is also required for the integration of cargoes into the tip and for the recycling of membrane [99, 101]. The study of the spatio-temporal regulation of growth machinery in the growing pollen tube faces the challenges of observing any dynamic system: in order to fully understand a process, it may be necessary to recognise and quantify large numbers of individual cellular components, quantifying their interaction with actin filaments in a highly dynamic manner. Much information is available from biological images of these dynamic systems, yet meaningful quantitative analysis is often difficult to be done accurately by hand and eye, especially when analysing the dynamic interactions of sub-populations of individual entities. Here, we take a computational approach to quantitatively analyse the association of mitochondria with actin filaments in growing pollen




Figure 4.1: A schematic diagram of a plant cell with a stiff, non-permeable cell wall (dark cyan), punctuated with plasmodesmata (gaps) to allow for transport of molecules through the semi-permeable membrane (gold). Inside the cell has many organelles including, but not limited to, the vacuole, mitochondria and chloroplasts; organelles such as chloroplasts and mitochondria are trafficked by the cell cytoskeleton, *e.g.* actin and microtubule filaments. The nucleus and endoplasmic reticulum are the centres of protein production in the cell.

tubes. Rapid transport of mitochondria in pollen tubes has been shown to be dependent on the actin cytoskeleton [102], and mitochondrial dynamics can be used as an indicator of the rate of actin-mediated cytoplasmic streaming and cargo transport. Mitochondrial movement can range from low-speed ($<5 \mu\text{m s}^{-1}$) movements to dashes along actin filaments at rates up to $10 \mu\text{m s}^{-1}$ [102–104]. Motion is believed to be predominantly powered by myosin motor proteins, although a role for actin dynamics in mitochondrial movement has also been suggested [104–106], with myosin alone producing the low-speed movements, and a combination of actin dynamics and myosin action resulting in rapid dashes [104]. In addition, microtubule dynamics also have effects on mitochondrial velocity, trajectory and positioning in pollen tubes via their role in directing the arrangement of actin filaments [102] or by providing tracks for low-speed movement and anchoring [107].

Using the same object-based algorithms as for NET1A-GFP:plasmodesmata colocalisation, we demonstrate a novel approach that can be used to quantify dynamic interactions effectively in a complex system and quantitatively analyse the interactions of mitochondria (discrete objects) with actin filaments (a continuous object) in the growing pollen tube. This algorithm can, with user-defined adjustments, be used to study the colocalisation of static discrete objects and static continuous objects, as well as dynamic interactions between colocalising structures.

4.2 Materials and Methods

4.2.1 Plant Growth and Sample Preparation

Seedlings of *Arabidopsis thaliana* (*L.*) *Heynh.* were grown on Murashige & Skoog Basal Salt Mixture (M5524; Sigma-Aldrich, St Louis, MO, USA) [108] with 0.8% plant agar (Duchefa Biochemia B.V. Haarlem, The Netherlands) under standard growth conditions (16 h light, 22 °C day temperature, 16 °C night temperature). Stable transgenic *A. thaliana* seedlings expressing the NET1A-GFP fusion protein, under the control of the native NET1A promoter, were incubated for 10 min in 0.1% aniline blue solution at pH 9.5 before imaging.

For dynamic studies using *A. thaliana* pollen tubes, we used plants stably expressing the fluorescent actin reporter ref fluorescent protein (RFP)-fimbrin actin-binding domain 2 (FABD2), which is controlled by the pollen-specific Lat52 promoter. Pollen was germinated as described by [109], with incubation periods of 3 h to 6 h at 22 °C. Prior to imaging, samples were incubated for 10 min with 2 μm MitoTracker Green FM (Life Technologies, Paisley, UK; henceforth referred to as MitoTracker) in liquid pollen germination media.

4.2.2 Imaging

Three-dimensional microscopy images of NET1A-GFP and aniline blue staining were acquired using a Revolution XD spinning disk confocal microscope (Andor, Belfast, UK) with an iXon DU897E back-illuminated electron-multiplying charge-coupled device (EMCCD) with a pixel size of $0.16\ \mu\text{m} \times 0.16\ \mu\text{m}$ (Andor). Z-stacks were collected with an PlanApoN 60×1.42 NA oil lens (Olympus, Tokyo, Japan) at optimum z-spacing of $0.156\ \mu\text{m}$, according to the Nyquist-Shannon reconstruction theorem [110, 111], with each channel acquired sequentially at each step with laser and filter switching. NET1A-GFP was excited with a 488 nm laser and fluorescence was collected via a band pass emission filter (510 nm to 540 nm) at 500 ms of EMCCD exposure. Aniline blue was excited with a 405 nm laser and fluorescence was collected via a band pass emission filter (417 nm to 477 nm) at 500 ms of EMCCD exposure.

Four-dimensional time series of RFP-FABD2 and mitochondrial dynamics (MitoTracker) were acquired again with the Revolution XD spinning disk confocal microscope (Andor), with a pixel size $0.13\ \mu\text{m} \times 0.13\ \mu\text{m}$. Z-stacks were collected with an UPlanSApo 100×1.4 NA oil lens (Olympus) at a z-spacing of $0.252\ \mu\text{m}$. Stacks were repeated 48 times, once every 20 s. Channels were acquired sequentially. MitoTracker was excited with a 488 nm laser and RFP-FABD2 was excited with a 561 nm laser with fluorescence collected via a dual band pass emission filter (512 nm and 630 nm). Excitation lasers are alternated by rapid acousto-optical tunable filter (AOTF) switching. Both images were collected at 32 ms of EMCCD exposure with an average of 32 exposures.

Focussing different wavelengths through the same lens led to slight differences in the focus position; this shift is known as chromatic aberration. Chromatic aberration can be largely compensated by the use of apochromatic lenses, as used in this paper; however, non-lens-based sources of chromatic aberration (*e.g.* differences between plane of coverslip and the plane of slide, optical path misalignments, etc.) may also affect the final images. We used carboxylate-modified polystyrene latex beads $0.03\ \mu\text{m}$ in mean diameter (L5155, Sigma-Aldrich) that have the excitation and emission spectra most closely suited for the two-wavelength set-up used in this paper. Beads were diluted to a concentration of 1:1000, sonicated, air-dried on a coverslip, embedded in Vectashield (Vector Laboratories, Burlingame, CA, USA), sealed on a slide and imaged using the same parameters as for the NET1A-GFP study. The images were then run through the same colocalisation program as the cellular images. For the field of view used, 75 beads were identified at both wavelengths. Vector shifts between a bead's centroid in the red channel and the green channel were analysed, and the average values in three dimensions were

used as a measure of translational shift. The average shift in each dimension was smaller than the uncertainty due to centroid bias (below); as such, no chromatic aberration correction was applied to the final data.

4.2.3 Image Processing

Illumination Correction and Normalisation

In fluorescence microscopy, uneven background illumination and vignetting due to the sensor used for image capture can cause serious problems during analysis: bright regions of background can obscure the desired foreground signals (*i.e.* data). In order to decrease the obscuring of data by the background, many different techniques can be applied to estimate the background illumination and thus remove the background signal (*e.g.* [112]).

Common background correction techniques include entropy-minimisation, envelope estimation, fitting and filtering. In entropy-minimisation, first introduced by [113], illumination and noise artefacts are modelled as a combinations of additive and multiplicative terms. By inverting this image forming model and minimising the entropy uneven illumination and shading effects can be estimated and thus removed. Alternatively, one can model the image shading as a surface, say quadratic, and fit said surface to the raw image, this will be affected by both background and foreground signals but will produce an estimate of the background signal fitting the surface shape of choice [112].

For the images presented in this paper, the simpler approach of filtering out foreground objects can be used. This approach is valid for this data as the objects are very small compared to the scale of background variation (figure 4.2).

For the 3D colocalisation of NET1A and plasmodesmata labelling, we used a relatively simple background approximation technique: low-pass filtering [114]. Low-pass filtering selects for low-frequency signals and attenuates high-frequency signals in the images. Signals of very high-frequency in an image are often noise and this technique is used to ‘smooth’ an image, removing such noise. Using a low-pass filter with a very low cut-off removes not just noise but also foreground signals; what remains is a good approximation of the background illumination.

Although it is possible to run such a filter on a 3D volume, due to the spinning disk’s imaging modality, the uneven illumination for these images was restricted to each 2D x-y plane. As such,

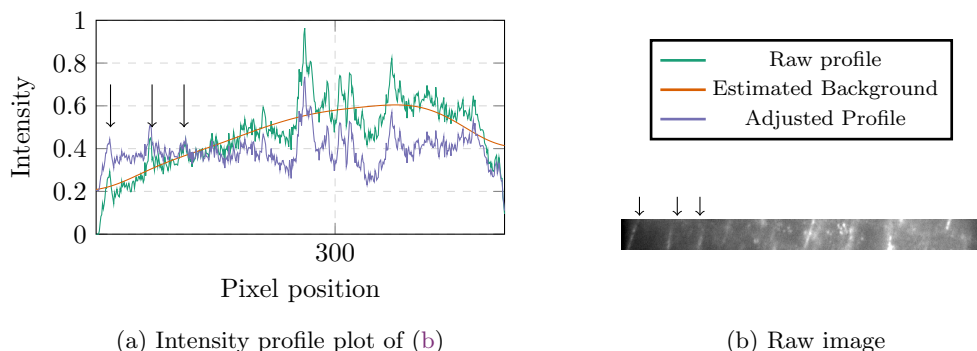


Figure 4.2: Retrospective background illumination correction helps to correct for low frequency changes in illumination due to the imaging modality. The low frequency changes in the raw image, a strip of which is shown in (b) can be estimated and removed to create a more even intensity profile where less intense features (black arrows in (a)) are brought into line with brighter features.

we ran 2D background illumination correction on each frame using the algorithm in algorithm 4.1. The value of standard deviation, σ , used for each sequential Gaussian filter is an integer from 1 to a user-defined maximum σ_{max} , which should be selected depending on the size of foreground objects. The early filters, with a low value of σ , will remove small signals such as noise, whereas the later filters, with higher values of σ , will remove larger signals (*i.e.* foreground objects). For these data, we used rotationally symmetric Gaussian low-pass filters with a kernel size of 3σ and an s.d. of $\sigma_{max} = 20$.

For colocalisation of MitoTracker and RFP-FABD2 labelling, we did not use full-image background illumination correction techniques; instead, it was necessary to create a mask of the pollen tube region. The mask was created using mean local thresholding techniques [115] to create binary image of the tube. For the MitoTracker channel, the minimum and maximum intensity values within the mask were used to normalise the image, giving the image a maximum dynamic range for the signal within the pollen tube (algorithm 4.2 (*p.* 59)). This gives the signal within the pollen tube the maximum dynamic range possible while preventing signals from outside the pollen tube from interfering with the colocalisation algorithm. Hot pixels (*i.e.* noise-related pixels with a high intensity value) were eliminated across the image series to ensure that the dynamic range after normalisation was biologically meaningful.

```
Input :  $I, \sigma_{max}$   
Output:  $I_c$   
1 for  $z \in [1, depth(I)]$  do /* for each z-slice in the image */  
2    $I_{temp} \leftarrow I(z)$   
3   for  $\sigma \in [1, \sigma_{max}]$  do /* for sequential kernel sizes */  
4      $G \leftarrow \text{gaussian}(\sigma)$   
5      $I_{temp} \leftarrow I_{temp} \star G$  /* convolution with  $G$  */  
6   end  
7    $I_{bgd}(z) \leftarrow I_{temp}$   
8 end  
9  $I_c \leftarrow I - I_{bgd}$ 
```

Algorithm 4.1: Algorithm used for uneven background illumination correction of images. For this image series, correction was run on each 2D frame in turn. The algorithm takes a 3D image (I) and a sigma value. For each sequential filtering, a Gaussian low-pass filter of size 3σ and s.d. σ is created. The filter is applied to the image through a convolution (\star) operation, σ is increased by one and the process is repeated until $\sigma = \sigma_{max}$. Once all sequential filtering steps have taken place, the over-smoothed image is the background approximation; this is removed from the original image. The output is a background-corrected 3D image, I_c . Note: the \leftarrow symbol denotes ‘becomes’ (e.g. in Line 2, I_{temp} becomes that z-slice of I); the \in symbol denotes ‘an element of’ (e.g. in Line 3 for each element σ of the interval $[1, \sigma_{max}]$ (i.e. all integers between 1 and σ_{max})).

```

Input :  $I, mask$ 
Output:  $I$ 
1  $M \leftarrow \max(\text{intensity}(I(mask)))$ 
2  $m \leftarrow \min(\text{intensity}(I(mask)))$ 
3 for  $\forall p \in I$  do
4   if  $\text{intensity}(p) < m$  then
5      $p \leftarrow m$ 
6   else if  $\text{intensity}(p) > M$  then
7      $p \leftarrow M$ 
8 end
9  $I \leftarrow \{i \in I : (\text{intensity}(i) < m) \times m\}$ 
10  $I \leftarrow \{i \in I : (\text{intensity}(i) > M) \times M\}$ 
11  $I \leftarrow (I - m)/(M - m)$ 

```

Algorithm 4.2: Algorithm used for pollen tube masking and normalisation of the dynamic range of MitoTracker (Life Technologies) within the mask. The algorithm takes the image (I) and a pollen tube mask ($mask$), and calculates the maximum (M) and minimum (m) intensity values within the mask. Any pixels in the image with an intensity below the minimum are set to the minimum; likewise for pixels above the maximum. The whole image is then normalised based on the pollen tube maximum and minimum intensities.

Object-based Colocalisation

Object-based colocalisation in this paper is based on the algorithm proposed by [77], which is based on 3D blob-like detection, as described in [116]. This is a two-stage detection algorithm: detection of all blob-like signals followed by matching and pruning. All steps are carried out in 3D.

The local blob detection algorithm (algorithm 4.3 (p. 61)) uses an expected blob size to identify likely candidates; for NET1A-GFP and aniline blue signals, this size was a sphere with a 260 nm radius. The expected blob size is used in the Laplacian of Gaussian (LoG) convolution of the image and the statistical pruning of blob candidates. Blob candidates are selected from local maxima* in the image with intensity above a user-defined threshold; the user-defined threshold is selected to prevent noise being selected as blobs. All blobs are treated as ellipsoids for descriptor and centroid calculation before undergoing statistical pruning of the candidates[†]. Weighted

* Local maxima where detected by dilating the image with an element equal to the expected particle size; maxima where then set to locations where $I = I \circ b$

[†] Pruning of particles is done based on distance between itself and all other particles within a certain distance. If two particles are deemed to be the same particle, the one with a stronger intensity is kept under the assumption that the optical centroid of a particle of this type will be the brightest point.

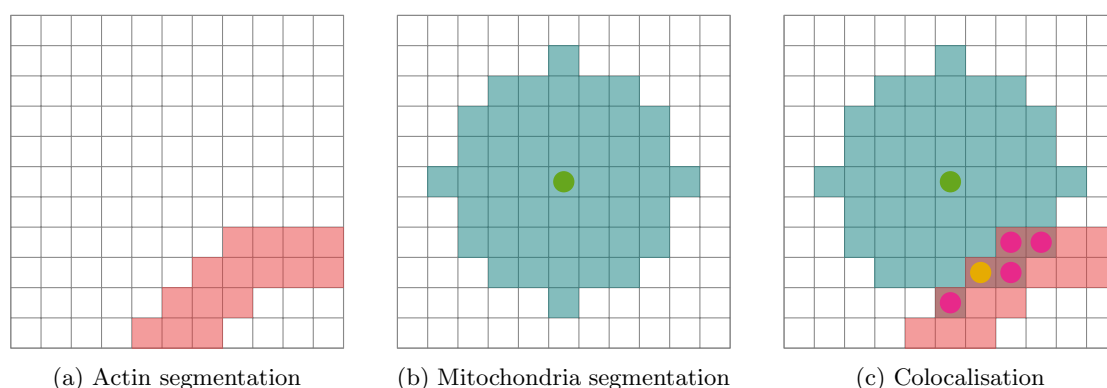


Figure 4.3: A schematic explanation of actin-mitochondria colocalisation in *Arabidopsis thaliana*. (a) A pixel-by-pixel example of actin segmentation. Each square on the grid counts as a single pixel and each red square is a pixel that has been segmented as actin. (b) A mitochondrial ‘blob’ in cyan. The blob is centred on the mitochondrial centroid, labelled with a green dot, as detected by the algorithm above. (c) Each actin pixel within the mitochondrial ellipsoid is considered a candidate for colocalisation; candidates are labelled with a pink dot. Of these candidates, only one is considered as a colocalising match, labelled yellow. This figure is a considerable simplification, as many actin pixels may be candidates for matching to multiple mitochondrial signals. The technique used by this algorithm is designed to optimise the matching results.

centroids are then calculated before the candidates in both images are paired using the Hungarian matching algorithm [16] (see appendix C). As the blobs are matched, they are removed from the image to prevent mismatching.

For MitoTracker-actin colocalisation, MitoTracker signals were identified as above, with an expected blob radius of $0.75\ \mu\text{m}$. Actin cables were then segmented using a mean thresholding technique (algorithm 4.4 (p. 62)); this technique identified all pixels from the actin channel images that were within actin cable signals. The ellipsoids calculated from the MitoTracker signals were then used to identify putative colocalising actin pixels (figure 4.3). The putative matches were then run through a Hungarian matching algorithm [16] to identify all actin pixels that colocalise with mitochondrial signals.

Due to the nature of the centroid calculation algorithm used in our and many other programs, there is a centroid bias in the program. The centroid, which was refined with an intensity-weighted centre of mass algorithm, was biased towards the centre of pixels for small signals (*i.e.* one or two pixels). Figure 4.4 (p. 63) shows how sub-pixel signals emanating from the corners and sides of a pixel’s region of view are drawn in towards the centre of the pixel. This

Input : $I_1, I_2, r_x, r_y, r_z, s, th_1, th_2$
Output: $CoLoc$

```

1  $S \leftarrow [s/r_x, s/r_y, s/r_z]$ 
2  $L \leftarrow \text{LoG}(S)$ 
3 for  $i \in [1, 2]$  do /* for each 3D channel */
4    $I_i \leftarrow I_i \star L$ 
5    $Candidates_i \leftarrow \text{localMaxima}(I_i)$ 
6   for  $C \in Candidates_i$  do /* for each candidate */
7      $S_C \leftarrow \text{meanIntensity}(\text{shell}(C))$ 
8      $E_C \leftarrow \text{meanIntensity}(\text{ellipsoid}(C))$ 
9   end
10   $Blobs_i \leftarrow \{C \in Candidates_i : S_C < E_C\}$ 
11   $Blobs_i \leftarrow \{B \in Blobs_i : \text{intensity}(B) > th_i\}$ 
12  for  $B \in Blobs_i$  do /* for each blob */
13     $E_B \leftarrow \text{ellipsoid}(B)$ 
14     $D_B \leftarrow \text{descriptor}(E_B)$ 
15  end
16   $Blobs_i \leftarrow \{B \in Blobs_i : \text{prune}(B, D_B) = \text{FALSE}\}$ 
17   $W_i \leftarrow \text{weightedCentroid}(\text{ellipsoid}(Blobs_i))$ 
18 end
19  $CoLoc \leftarrow \text{hungarianMatching}(W1, W2)$ 

```

Algorithm 4.3: Algorithm used for locally based colocalisation. Each step is carried out in 3D and each preprocessed and normalised channel (I_1, I_2) is analysed separately until the centroids are matched in the final process. The LoG convolution (\star) of each image enhances blob-like structures. After convolution, the images are renormalised. The filter size for the LoG convolution is determined by the resolution of the image in each direction (r_x, r_y, r_z) and the expected size of the blob (s). Local maxima are detected in each channel through morphological opening and identifying the locations of particles that are brighter than their immediate surroundings. Each candidate is given a centroid position and the ellipsoid around the candidate is calculated; a thin shell around this ellipsoid is also calculated. The comparison of the mean image intensity in the ellipsoid to the mean image intensity in the shell allows candidates to be checked and selected for. Blobs with an intensity above the user-set threshold (th_i) for each channel are then selected for. For this experiment, identified blobs are pruned using statistical measures. The remaining blobs are then given a weighted centroid and the blobs from both images are matched using the Hungarian matching algorithm [16].

Input : I, r_x, r_y, r_z, s
Output: $Cables$

- 1 $S \leftarrow [s/r_x, s/r_y, s/r_z]$
- 2 $E \leftarrow \text{ellipsoid}(S)$
- 3 $h \leftarrow I \star E$
- 4 $mean \leftarrow h / \text{sum}(h)$
- 5 $Cables \leftarrow I > mean$
- 6 $Cables \leftarrow \text{reconstruct}(Cables, \text{imOpen}(Cables))$

Algorithm 4.4: Algorithm used for mean-based thresholding of actin cables. For a 3D image, an ellipsoid size (S), scaled by the image resolution (R_x, R_y, R_z), is used as a kernel. The mean of the surrounding neighbourhood, as defined by the kernel, is calculated for each pixel. The value for each pixel is then compared with this mean; pixels greater than the mean of their neighbourhood are counted as actin pixels. The thresholded image is then reconstructed by morphological opening. Morphological opening is the erosion followed by the dilation of a binary image; reconstruction by morphological opening is a process by which the shapes of the actin cables can be extracted more accurately.

adds an unknown uncertainty to all measurements of distance between two centroids. For the purpose of our results, we have assumed a minimum uncertainty of 0.5 pixels in each direction (*i.e.* 170 nm in 3D Euclidean distance); this the maximum uncertainty due to centroid bias between two points. For signals larger than one pixel, this centroid bias becomes negligible, as the bias from each pixel helps to cancel bias of neighbouring pixels.

Pixel-based Colocalisation

Colocalisation analysis as described in [95] was performed by selecting square regions of interest (150×150 pixels in x-y) within the dataset where punctae in both the red and green channels were of superior contrast. Background fluorescence was removed from the colocalisation analysis using masks generated by the ‘threshold’ tool in ImageJ (National Institutes of Health, Bethesda, MD, USA). The manually adjusted threshold selected for each region of interest depended upon the local level of background fluorescence, as this level was not consistent throughout the dataset. Eight images encompassing 10 cells were combined onto one canvas in ImageJ and analysed using ColocalizerPro (Colocalisation Research Software) to generate a Pearson’s correlation value. The negative control value was obtained by displacing the red and green channel data by 11 pixels along the y-axis; this was chosen to ensure that the cell end walls were compared with background levels in the opposing channel. Regions of interest and masks were duplicated and displaced accordingly.




Figure 4.4: The centroid bias of intensity-weighted centre of mass calculations draws the calculated position of objects towards the centre of pixels. A point of known position was placed in one of nine virtual sub-pixel regions; the figure shows a single pixel divided into these nine uniform regions. Each point was convoluted with a disc one pixel in diameter. This was repeated for all nine regions. Each artificial image then underwent centroid detection and weighted centroid refinement, as per the main algorithms. The figure shows the true centroid in cyan and an approximation of the calculated centroid in red; the arrows show the displacement of centroids towards the pixel's centre. The signal centred on the pixel centre had negligible displacement (<0.01 of a pixel).

4.2.4 Software Packages

Object-based colocalisation algorithms were implemented off-line using the software package MATLAB and Image Processing Toolbox Release 2012b (The MathWorks Inc., Natick, MA, USA). The full algorithm, as used in this paper, is available from the corresponding author.

Pixel-based colocalisation used ImageJ (National Institutes of Health) [117] to generate masks, which were analysed with CoLocalizer Pro (CoLocalisation Research Software).

4.3 Results

4.3.1 Colocalisation of NET1A with Plasmodesmata

To demonstrate the performance of our technique we have re-analysed a dataset first published by [95]. This z-stack of images of *A. thaliana* root meristems was collected using spinning disk confocal microscopy. Nano-domains along plasma membranes were labelled using two probes. The first was a GFP fusion to the actin binding protein NET1A (NET1A-GFP), a protein that localises to the plasma membrane and concentrates at plasmodesmata and has been suggested

as being involved in the link between the actin cytoskeleton and the plasma membrane. NET1A-GFP localises to distinct punctae along the plasma membrane. The second probe was aniline blue, a dye excitable at 405 nm that labels callose-rich cell-to-cell junctions called plasmodesmata. The signal along the plasma membrane domains was isolated using a threshold approach and the correlation between pixel intensities was measured using Pearson's coefficient (pixel-based colocalisation; Materials and methods). This gave a positive value of 0.859, suggesting a strong coincidence between the NET1A-GFP structures and plasmodesmata. A maximum value of 1 denotes an absolute correlation of pixel intensities but should be impossible to obtain due to random noise and optical constraints.

In this particular case, the Pearson's correlation value does not give a quantity that can be related to the distance distribution between the nanopunctae of the two channels. Variability throughout the image of fluorescence from out-of-focus planes, instrument noise and autofluorescence prevents the effective masking of background signal from the correlation calculation. To understand whether green (NET1A) nanopunctae perfectly coincide with red ones (plasmodesmata), it would be necessary to compare Pearson's correlation values between several control experiments to test alternative hypotheses. Ideally, these would use two alternative fluorescent labels of equivalent intensity to aniline blue. The first label would localise homogeneously across the cell wall, whereas the second would highlight the zones between plasmodesmata. These ideal probes are not available. To estimate the element of background fluorescence, the two channels were displaced by 11 pixels, giving a Pearson's correlation value of 0.132. Confidence that the difference of 0.727 represents colocalisation at the resolution of nano-domains is derived from high-resolution transmission electron microscopy imaging combined with an immunoaffinity label [95]. These techniques showed that NET1A protein specifically localises to within nanometres of the membrane surrounding plasmodesmata. Such tools and techniques are expensive, time-consuming and rely upon specialist skills and equipment. More highly resolved optical solutions to image nano-structures are often limited by the brightness and specific characteristics of the probes. There is a requirement for an alternative method to extract higher-quality quantitative colocalisation results from fluorescence imaging data.

The probe signals for each label were isolated using a local maxima detection approach and candidates were colocalised between channels using the Hungarian matching algorithm [16] (figure 4.5 (*p.* 66)). Of 954 NET1A-GFP signals identified, 682 colocalised with signals from aniline blue; there were 1162 individual aniline blue signals. As a control, we shifted the image three pixels horizontally to keep plasma membrane with plasma membrane and cytoplasm with

cytoplasm while shifting the nano-domains enough to alternate (rather than superimpose) red and green punctae. The resulting image showed that only 351 of 954 NET1A signals now colocalised with aniline blue signals; there were 1162 aniline blue signals.

The algorithm allows us to more than just look at whether signals colocalise or not. Because of the object-based nature, the program can also find the Euclidean distance of any colocalising signals (*e.g.* the distance between a NET1A-GFP and an aniline blue signal). Figures 4.6a and 4.6b (*p.* 67) show this information in two formats: figure 4.6a (*p.* 67) shows a 3D representation of the colocalising pairs, colour coded based upon their distances, whereas figure 4.6b (*p.* 67) shows a maximum z-projection with the data overlaid. Figure 4.6c (*p.* 67) shows the distance histograms for the data (blue) and the control (red; see above). The mean colocalisation distance of the population is 330 nm. The distance at which 50% of objects colocalised is \sim 250 nm. The mean value for the control is 685 nm with the distance at which 50% of objects colocalised being \sim 700 nm. It is clear from the superimposed histograms (figure 4.6c (*p.* 67)) that a 660 nm spatial shift has resulted in an almost equal displacement of the colocalisation distances. For a comparison, we repeated the three-pixel control using the Pearson's coefficient workflow. The modest shift in red and green channels generated a correlation value of 0.66. Although this was reduced from the original value of 0.86, it still indicated a positive correlation despite the 660 nm shift. The object-based method allowed intuitive quantification of colocalisation patterns, whereas the Pearson's coefficient was difficult to interpret. Pixel-perfect co-alignment of the channels is required for maximum correlation, as would be expected from two components of plasmodesmata, but this correlation was not completely broken by the shift in channels.

Figure 4.6d (*p.* 67) shows the intensity difference between channels of colocalised pairs. For each pair, the absolute difference between the normalised intensity of the signal on each channel is plotted. The 3D data have been overlaid onto a maximum z-axis projection of the aniline blue channel, which shows that the difference increases where the aniline blue signal is brighter; this acts as a good control to show that bleed-through of signal is not the cause of colocalisation. If bleed-through were a major issue then one would expect the difference between channels to remain similar for all colocalising pairs or to vary only due to uneven illumination of the brighter channel.

It is also possible to extract a biological context from this control. Assuming that aniline blue stains proportionally to callose content, the intensity difference is a quantification of the ratio of NET1A accumulation to callose content. It can be seen from figure 4.6d (*p.* 67) that at the local




Figure 4.5: NET1A-green fluorescent protein (GFP) in *Arabidopsis thaliana* colocalises with plasmodesmal signals as labelled with aniline blue. (a) Identified plasmodesmal labelling (aniline blue). Yellow dots ($n = 682$) represent object-based colocalisation of plasmodesmal signals with NET1A signals; red dots ($n = 480$) represent plasmodesmal signals that are not colocalised. (b) Identified NET1A-GFP signals. Yellow dots ($n = 682$) represent object-based colocalisation of NET1A-GFP signals with plasmodesmal signals; red dots ($n = 272$) represent NET1A-GFP signals that are not colocalised. (c) Merged channels show colocalisation of NET1A-GFP signals at plasmodesmal cell junctions as labelled with aniline blue. All images show a maximum projection through the z-axis.

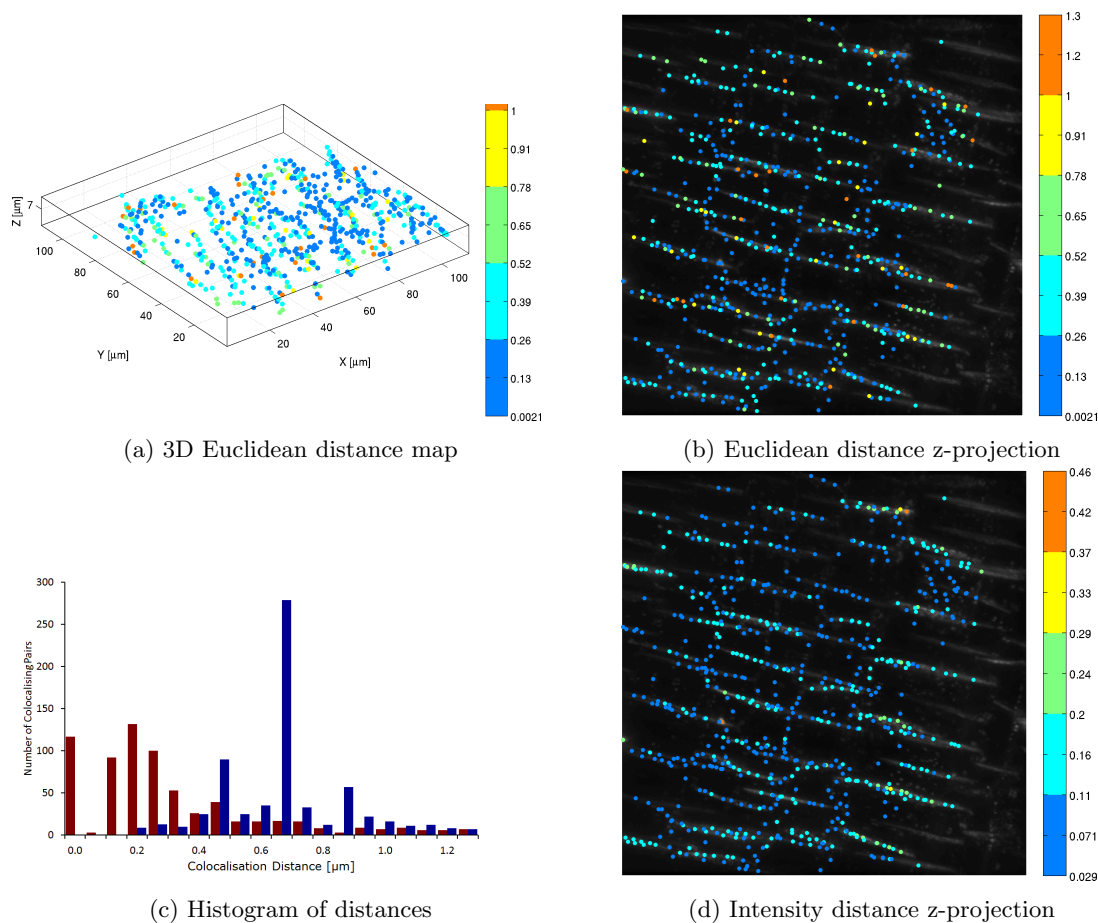


Figure 4.6: (a) Euclidean distance map of colocalised object pairs of NET1A-green fluorescent protein (GFP) and plasmodesmal cell junctions of *Arabidopsis thaliana*, labelled with aniline blue. The scale runs from 0 nm to 260 nm (cobalt blue) to 1 μm to 1.3 μm (orange), as indicated by the bar on the right of the figure. Note that this Euclidean distance carries a known uncertainty of ± 170 nm due to centroid bias (above). The histogram of the Euclidean distance to the number of colocalisations is in the lower left of the scatterplot. (b) Distance map overlay with a maximum z-axis projection of the cells. (c) Distance histogram for the colocalisation of NET1A-GFP and aniline blue signals. The red histogram shows the distance distribution for the standard image (20 bins); note the drop in the second (0.07 μm to 0.13 μm) bin. The blue histogram shows the distribution for the control image (see main text). (4.6d) Intensity difference map of colocalised object pairs of NET1A-GFP and plasmodesmal cell junctions labelled with aniline blue. The scale runs from 0.0 to 0.11 (cobalt blue) to 0.37 to 0.46 (orange), as indicated by the bar on the right of the figure.

level, NET1A-GFP and callose content are closely linked, as the variation of intensity difference between neighbouring plasmodesmata was much less when compared with the variation between neighbouring cell files and distant regions.

NET1A-GFP signal within the dataset is not distributed evenly at the tissue scale. The cell files at the root cap epidermis are devoid of GFP signal. Neighbouring cell files show an asymmetric signal, due to differential expression of NET1A. The outcomes are two distinct zones of ‘orphan’ aniline blue punctae to the right and far left of the field of view. This could be compensated for by either masking the zones lacking a NET1A-GFP signal, or by discounting the ratio of paired to unpaired aniline blue punctae and measuring only the ratio of paired to unpaired NET1A-GFP (since NET1A-GFP punctae always appear in zones with a strong aniline blue signal). We have opted for the latter, since the boundaries of any mask proved to be subjective.

Our results confirm the colocalisation of plasmodesmata (aniline blue) and NET1A (GFP) as originally reported in [95]. Furthermore, we have produced quantified results that show the distances between colocalising pairs and quantified results showing the direction of displacement and the intensity of signals, which can be used in further experiments and to gain imaging and biologically relevant information.

4.3.2 Colocalisation of Mitochondria with Actin Filaments across Time

To demonstrate the extent of our technique we produced a dataset showing the association of mitochondria, labelled with MitoTracker, and actin filaments (RFP-FABD2); this shows how our algorithms can be used in a novel way to examine the colocalisation of discrete objects, *i.e.* mitochondria, with a continuous object, *i.e.* actin filamentous cables. Furthermore, these data were extended into four dimensions: a 3D time series. The time series of images of *A. thaliana* pollen tubes was collected using spinning disk confocal microscopy. The MitoTracker labelling of mitochondria shows these as distinct punctae throughout the pollen tube, whereas the actin-binding RFP fusion protein labels filamentous actin within the pollen tube. Quantifying the association of these two probes is a unique situation: usually, colocalisation examines whether or not two sets of ‘blobs’ are associated; here, we must look at whether one set of ‘blobs’ is associated with any sub-domain of a ‘cable’ or ‘network’.

Figure 4.7 (p. 70) shows an example of segmentation of the actin network within the pollen tube and the associated mitochondrial signals. In areas of the pollen tube with indistinct actin signals


(*e.g.* the pollen tube tip) there is little to no segmentation of actin cables or colocalisation of the mitochondria with actin. This region of the pollen tube contains a more motile and unbundled network of filamentous actin, which is required to integrate vesicles into the pollen tube tip.

Figure 4.8 (p. 71) shows how bleaching[‡] of the image affects the segmentation of actin cables and thus the colocalisation program. Figure 4.8a (p. 71) shows that the overall identifiable volume of the pollen tube, as used in the masking process, did not change much with time; however, it should be noted that the length of the pollen tube increased during the imaging time but the segmentable actin cable network did not invade the new volume. However, it is clear from the raw data (not shown) that the actin segmentation volume decreased, with an apparent exponential drop with time. Figure 4.8b (p. 71) shows exponential curve fitting for the volume of actin segmented at each time point (red; $y = 0.94 \times e^{-0.012x}$) and the mean pixel intensity within the pollen tube mask (blue; $y = 0.93 \times e^{-0.016x}$; both graphs have been normalised so that the maximum mean pixel intensity or the maximum actin volume segmented equals 1).


The segmentation of actin by the program depends upon the amount of RFP-FABD2 signal above a certain manually chosen threshold limit; if the intensity of the image decreases, then either less filamentous actin will be detected or the cables detected will be ‘seen’ as thinner (*i.e.* have a lower segmented volume). Having observed these images manually, we find that the code detected clear actin filaments in all time frames; however, the apparent width of actin filaments detected decreased with time (see figure 4.8d (p. 71) for a diagrammatic explanation). This decrease in apparent width seems to match the decrease in brightness of the RFP-FABD2 signal within the pollen tube. Figure 4.8c (p. 71) shows how the average distance between segmented actin cables increased with time; this increase suggests that the program is identifying thinner cables within the image.

A further example of bleaching is in the first frame of the data (not shown). The data in this frame were very hard to extract and resulted in quantified data that were noticeably different from the subsequent frames. It is plausible that the pre-imaging accumulation of fluorescent molecules within mitochondria and non-specific labelling are responsible for this difference. Continuous replenishment of mitochondrial fluorescence due to oxidation of mitochondrial MitoTracker may make these objects more resistant to photobleaching effects than non-specifically labelled artefacts. Many of the correction techniques applied (*e.g.* masking and normalisation) helped to


[‡] Bleaching is the phenomena where fluorescent molecules become unable to fluoresce overtime leading to a decrease in fluorescent image intensity.



(a) 3D filament and blob segmentation



(b) Maximum Z-projection



(c) Zoomed in view

Figure 4.7: RFP-FABD2 and MitoTracker Green (Life Technologies) label segmentation and colocalisation. (a) An example of segmentation from a single time point; the image has been rotated to give the best view and the *Arabidopsis thaliana* pollen tube tip is at the top of the image. The image shows actin cables (red) and mitochondria (green) as detected by their respective probes. Mitochondria that have been calculated as colocalising are shown in yellow. Mitochondrial size is artificial to make them easier to see. (b) A maximum z-projection of the RFP-FABD2 label at the same time point. The projection has been rotated so that it matches the segmentation rotation. (c) Close-up of the region highlighted by the white box in (a).

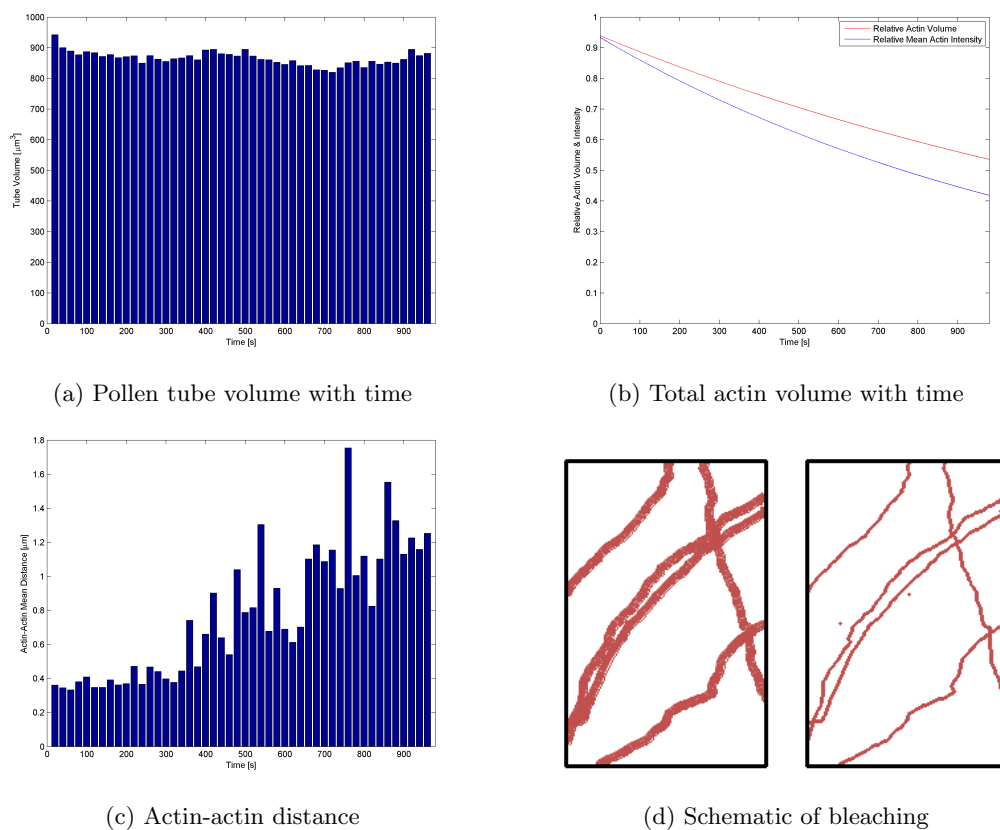


Figure 4.8: Bleaching artefacts affect the amount of actin that can be identified by the program. (a) The identifiable *Arabidopsis thaliana* pollen tube volume remains constant for all 48 frames; the expectation is that the volume of actin within the tube would stay the same or increase due to more filaments forming thick cables. (b) The volume of actin cables within the pollen tube decreases with time (red fitted exponential curve; root mean squared error: 0.0312); this drop is exponential in nature and follows a similar curve to the drop in average intensity within the pollen tube (blue fitted exponential curve; root mean squared error: 0.0736). The less rapid decrease in the volume of actin compared with the mean intensity within the pollen tube is probably due to the Gaussian spread of light during the imaging process and the cubic nature of the volume measurement. (c) A decrease in segmented actin volume is balanced by an increased mean actin-actin distance; this implies that detected actin cables appear thinner and that the effect of bleaching may be the detection of thinner actin cables over the loss of actin cable number. (d) A schematic of detected actin cables at early and late times within the same region of the pollen tube, showing how bleaching can lead to a decrease in thickness of the actin and the Gaussian blur around the actin signals.

increase the number of segmentable MitoTracker signals and helped give a better colocalisation result when compared manually to the data.


Figure 4.9 shows the number of MitoTracker signals (mitochondria) and the number of colocalising mitochondria for each time point of the series. The number of identified mitochondria signals increased slightly with time, whereas the number of signals ‘colocalising’ with actin decreased with time. The increase in mitochondria signals was due to an overall increase of mitochondria across the pollen tube (data not shown) rather than in a particular region of the tube. Unfortunately, the decrease in the number of colocalising mitochondria signals appeared to be due to the bleaching artefacts of the actin channel and followed a mild exponential decrease. The numbers of mitochondria in the pollen tube tip region, where actin cables are diffraction-limited and do not segment clearly, did not increase dramatically (data not shown); further human viewing of the time series showed that the decrease in colocalising mitochondria was apparently evenly distributed across the rest of the tube.

Our results show that the quantification of organelles as they are transported along cytoskeletal systems is a highly achievable goal. Our data also show that bleaching through time is a major factor influencing the results from such an experiment. The data above have been shown to demonstrate both these sides of bioimage informatics: novel techniques allowing access to new measurements and quantifications, and the possible pitfalls in striving for these new tools.


4.4 Discussion

4.4.1 Extracting Sub-resolution Information of NET1A


We have used algorithms to detect individual plasmodesmata automatically within the crowded 3D environment of the *A. thaliana* root apex. Co-imaging of two markers for plasmodesmata (aniline blue for cell wall callose and NET1A-GFP for plasmodesmal membrane) has allowed us to demonstrate the utility of this technique in quantifying 3D colocalisation relationships within populations. Seventy-one percent of plasmodesmata labelled with aniline blue can be matched with foci of NET1A-GFP. Three-quarters ($\sim 75\%$) of these partners have weighted centroids that lie less than 400 nm apart. Artificially displacing the two channels by just three pixels shows obvious disruption of this colocalisation pattern, in line with expectations. By contrast, the alternative method of assessing colocalisation using Pearson’s correlation testing provides no spatial quantities. Figure 4.10 (*p.* 74) describes the potential ultra-structure of the observed plasmodesmata based on transmission electron microscopy data [95]. Individual



(a) Colocalisation of mitochondria particles with actin filaments



(b) Mean distance between colocating mitochondria and actin



(c) Percentage of colocalising mitochondria with time

Figure 4.9: Colocalisation of *Arabidopsis thaliana* mitochondria with actin is successfully assessed despite bleaching. (a) Number of mitochondria ‘particles’ detected at each time point; the red portion of the bar shows the colocalised mitochondria; the blue portion shows the non-colocalised ones. (b) The mean distance between colocalising mitochondria and actin signals vary between frames; however, if bleaching was creating false positives, we would expect to see an exponential increase that mirrors the exponential decrease in intensity due to brightness. Note that the percentage of colocalising mitochondria decreases with time; this could be due to an increase in total number of recognised MitoTracker Green (Life Technologies) signals or an increase in false positives due to bleaching of the actin channel, which fits the data (not shown). (c) The percentage of colocalising mitochondria decreases with time. The pattern of decrease may match the bleaching issues identified for actin segmentation; it is likely that some mitochondria are considered to be too far from segmentable actin due to the effects of bleaching on the program. This decrease in identified MitoTracker Green signals is not due to an increase in mitochondria in regions where actin filaments are below resolution and unsegmentable.




Figure 4.10: Diagrams showing the relative positioning of some molecular components of *Arabidopsis thaliana* root meristem plasmodesmata. The cell wall and channel width are drawn to approximate scale. (Left) The arrangement of plasmodesmata as suggested by transmission electron microscopy and light microscopy (*e.g.* [95]). (Right) In this perfectly symmetrical scenario, the centroid positions for callose and NET1A-GFP would be perfectly aligned. The signal for NET1A-GFP (green ring) would overlap almost entirely the signal for aniline blue (callose deposit) in the cell wall.

pores will differ in their orientations but the pore presented in figure 4.10 is drawn to show the organisation of components in channels along apical-basal cell walls (horizontal cell walls as viewed within figures 4.3 and 4.6 (*p.* 60) and (*p.* 67)). The cell wall interrupts the callose sphincters at either end of the pore. In theory, with sufficient magnification and numerical aperture, it should be possible to differentiate between symmetrical and asymmetrical deposition of components across the pore (figure 4.10). Our dataset, however, was not optimised for high-resolution analysis. Ideally, we would require data at the Nyquist resolution ($\sim 3 \times$ the magnification of the current dataset). Such a resolution would decrease the uncertainty associated with the centroid bias, along with implementation of possible bias-reducing techniques, and would allow information to be gleaned at a resolution below that of the cell wall thickness.

Along with this higher resolution comes further information that can be extracted from object-based colocalisation techniques. We were also able to produce maps, such as those in figure 4.6

(*p.* 67), which show the direction between colocalising pairs (data not shown). With greater resolution, we could identify if there was any consistent direction of displacement between colocalising signals (*i.e.* any asymmetry within the system).

Furthermore, the use of an intensity difference map (figure 4.6d (*p.* 67)) may lead to a novel quantifiable feature that can be used to discriminate between different plasmodesmal fluorescent markers. It would be of great interest to compare maps of various live-cell plasmodesmal markers, as it might be possible, using this program and the intensity difference maps, to discriminate between markers that localise proportionally to the number of channels and markers that specifically respond to the local callose content of the cell wall. This could potentially differentiate between ‘open’ versus ‘closed’ plasmodesmata in living tissue responding to stimuli in real time.

4.4.2 Co-imaging of Motile Mitochondria and Cytoskeletal Tracks

Co-imaging and population analysis of mitochondria and actin filaments within the pollen tube have shown that not all mitochondria colocalised with actin cables. This is, perhaps, not surprising, as the association of mitochondria with microtubules has been documented as being transient in the pollen tube [107]. Romagnoli and colleagues [107] proposed a model of antagonistic action of microtubules and actin filaments on the movement of mitochondria at the pollen tube cell cortex. Cortical microtubules provide a means of anchoring mitochondria during their motion towards the pollen tube tip. Observation of the movement of purified pollen tube mitochondria showed that their movements on actin are salutatory rather than processive. Also, removing microtubules *in vivo* using drug treatments helped to smooth the flow of mitochondria [107]. This is in contrast to the situation in root hairs where the depolymerisation of microtubules increased the frequency of interrupted motion [104]. It would be interesting to repeat this experiment using our own dual-fluorescent pollen lines to test the hypothesis that the association of mitochondria with the actin cytoskeleton increases upon such treatments. Ideally, triple labelling to include the microtubule cytoskeleton would allow testing of the hypothesis that the organelles must be engaged with at least one cytoskeletal system at any one time. One limitation of the current experimental setup is that the fine cortical cables used for mitochondrial movement are difficult to image with the RFP-FABD2 probe. Longer exposure times are ideally required, which would be incompatible with mitochondria movement. An extra fluorescent probe would compound this problem.

4.4.3 Prospective Pitfalls in Plant Bioimage Informatics

Showing the object-based colocalisation of mitochondria and actin filaments with time shows both the power of such a bioimage informatics tool and also how careful one must be to avoid major pitfalls. Bleaching can cause severe effects on the results of any computational analysis. Easy to identify controls should be put in place to check that this is not the case. Fig. 8 shows how the distance between a colocalising MitoTracker signal and its matched actin signal was constant through the time series but how the percentage of MitoTracker signals considered to be colocalising decreased; this shows that the code was segmenting a decreased thickness of actin cables and was not able to compensate for this.

Actin filaments have a saturation limit for fluorescent binding proteins. These probe molecules exchange with a finite cytoplasmic pool and several seconds are required for bleach recovery of filament fluorescence. Due to the significant laser power used during imaging, we feel that this decrease in segmented actin volume represents the increased proportion of bleached fluorophores within the pollen tube. We are currently comparing and developing techniques to correct for this bleaching, either at the imaging or bioimage informatics stages.

One possible compensation is to consider that all actin cables are likely to be below resolution in this particular microscopy set-up; therefore, one solution is to segment the actin cables and set the diameter of all cables to 200 nm. Another option would be to adapt the threshold for actin segmentation based on an a priori bleaching curve or, indeed, by matching a curve to some measure of bleaching (*e.g.* actin volume) for the dataset and post-processing before recomputing the actin segmentation.

A technique that has applications for many bioimage informatics problems is to consider the actin probe signal as a representation of a skeletal network (*i.e.* a network of one voxel thickness at any point). The branches of these networks would be skeletal lines one voxel thick in 3D space; this would remove issues relating to the thickness of actin cables but requires a high enough resolution to make out individual actin filaments.

Many articles on ideal bioimage informatics or image processing and analysis emphasise the issue of assumptions: without clearly stating assumptions when designing and publishing algorithms and programs for automated analysis of images, errors can begin to build in the application and evaluation of the program and its data. There is, however, a problem of accidental assumptions: algorithmic bias. We have described how intensity-based object centroid calculation has an

inherent bias, a problem we hope to address in future versions of this code. Other such biases often occur in programs designed for automated analysis. Unless they are made fully open to program users, such biases could cause misinterpretation of results (*e.g.* the dip in the second bin of figure 4.6c (*p.* 67)) without realising the cause of such results. This particular bias becomes negligible for objects much greater than one voxel in size; hence further magnification to Nyquist limits is important in the initial imaging stages.

The object-based colocalisation techniques that have been showcased in this paper show not only the power of relatively simple bioimage informatics tools but also the ability of well-designed tools to be applied to multiple imaging scenarios. As new scenarios are approached, the bioimage informaticists working with a code can improve and adapt the code, giving it more abilities or making it more accurate; these adaptations can be saved in a modular manner, allowing a single code to be adapted for multiple situations based upon the needs of the biologist, the imaging modality, issues with the images and the question being asked. With this in mind, the program used in this paper is currently being refined and adapted further; it is hoped that issues like bleaching, hot pixels and poor dynamic range can be fixed using computational and statistical methods that can be turned on and off by the user. The final program will hopefully be released as part of a larger project that is currently ongoing.

4.4.4 Future Work

This article has showcased object-based colocalisation in plant sciences. We have highlighted the power of such techniques and also the pitfalls that researchers can fall into. We also intend to continue development by including the tracking of trafficked particles (*e.g.* mitochondria) along networks such as actin filaments within time series. Expansion of the program in this way would allow biologists to study if, how and when mitochondria or any other labelled particles are trafficked to particular places, such as the growing pollen tube tip. With continuing rapid improvements in microscopes and microscopy techniques, it may be possible to view this trafficking with fantastic resolution in not just the x-y plane but, in the orthogonal and temporal planes, creating a mass of available ‘phenotypic’ data that can be used to probe existing models or to identify the functions of novel genes.

One area of bioimage informatics research that we are currently pursuing is the improvement of techniques for segmentation of networks, such as actin, from standard resolution imaging techniques. With optimally segmented actin networks and strong compensation for bleaching effects,

it will be possible to use dynamic object-based colocalisation techniques to study much more than just how many colocalise and also to study questions such as ‘Are trafficked mitochondria colocalised with actin for their entire journey?’ or ‘Is trafficking different at newly formed actin networks (*e.g.* the pollen tube tip) compared with established cables (*e.g.* far from the tip)?’ This aim is part of a larger project that is currently ongoing and we hope to release in the form of a graphical user interface-fronted program or plug-in for use by other researchers.

Epilogue

In this chapter, I have provided an extended 3D particle detection algorithm with a novel post-detection pruning step and demonstrated the performance in the context of a quantitative colocalisation assay analysis pipeline for plant cell biology experiments.

I have explained the key uncertainties with this approach and other, similar quantitative methods for a non-computing readership whilst also showing the extra information that can be gained by such an object-based, quantitative approach.

The work in this chapter is particle detection at a small scale: all the objects of interest are either sub-pixel or up to a few pixels in size. Whilst this has its challenges, the challenges of larger objects are many. For example, at this scale, object rotation can be neglected but, at a larger scale, object orientation will have an effect on the results of a particle detector. In chapter 5 I describe a new approach for detecting and counting larger scale ellipses in images.

Chapter 5

Ellipse Detection by Hilbert-Edge Detection And Ranging (HEDAR)

Prologue

In chapter 4 (*p.* 48) I looked at counting and detecting small (pixel-scale) objects in fluorescence microscopy images. In this, and the next, chapter I consider the problem of supra-pixel elliptical objects, that is particles or ‘blobs’ that are on the scale of 5–20 pixels and may or may not be rotationally symmetric. This size of object is of particular interest in bioimaging as this is often the size of nuclei signals in high-throughput and developmental biology fluorescence micrographs.

Approaches to ellipse detection in images already exist, however, many are affected by noise, clustering or overlapping of objects, and, most importantly, those methods dependent on edges often consider ellipses of major axis length around 5–20 as on the scale of noise and these objects are thus filtered out. In many biological images, nuclei are on this size scale and thus some general ellipse detection approaches are not appropriate for the analysis of microscopy images.

In this chapter I introduce HEDAR a novel contribution to the ellipse detection problem that a) does not require an edge map, b) does not require *a priori* knowledge of the number of ellipses present and c) can successfully detect elliptical objects that other approaches may consider on the scale of noise.

I go on to compare HEDAR with several other ellipse detection approaches and validate the algorithms on synthetic and real data, showing HEDAR to be as good as these comparator methods in most cases and besting other approaches in certain cases, noticeably images with large numbers of ellipses.

In this chapter, I also present a case study of the application of this new method to nuclei detection in fluorescence microscopy, a challenging task due to the high levels of heterogeneity and clustering in many microscopy images. I describe HEDAR in the context of nuclei detection before benchmarking the method on datasets from the Broad Bioimage Benchmark Collection and comparing against published results and several comparator approaches. In this case study I have chosen to focus on nuclei counting and detection, as opposed to segmentation or approximation of ellipse parameters.

Declaration: This chapter is a draft article being preparing for submission to an appropriate journal. This chapter is presented as it will be submitted, although referencing and notation has been altered and cross-referencing added for consistency throughout this thesis. Some stylistic changes have been made for consistency. The majority of the text is verbatim; however, additions to the body of text are included along with footnotes and, where appropriate, added appendices.

5.1 Introduction

Ellipse detection is a key problem in image processing and will remain an open and important research area for the far future. Due to its importance, new methods of detecting ellipses are in constant development and undergo regular improvement, *e.g.* the many variations of the Hough transform (see [118, 119] for recent reviews). Each of these methods, however, may not be suitable for all scenarios and current limitations include a reliance on pre- or post-processing steps, such as edge detection; the difficulty of detecting multiple ellipses; and distinguishing between overlapping ellipses.

Ellipse detection algorithms tend to fall into three broad groups (as discussed in [118]): least-square focused ellipse fitting methods, Hough transform-like, *i.e.* voting-based methods, and

edge following approaches. All three approaches require some amount of pre- or post-processing, *e.g.* to enhance ellipse selectivity of least-squares approaches. Further, calculating an edge image, such as by Canny edge detection, is essential to many methods. A fourth category can also be considered with methods that do not quite fall into any one of these categories, such methods tend to be heuristic or hybrid techniques.

5.1.1 Ellipse Fitting Methods

Those methods that fall under the least-squares category turn the problem of ellipse detection into an optimisation problem and many methods are extensions of the early mathematical attempts, *i.e.* [120–122]. These extensions were developed to reduce bias and noise sensitivity [123], whilst increasing the ability to deal with occlusion [124] and the stability of the system [125]. Although solutions have been proposed, least-square methods are still very sensitive to noise and thus are difficult to use in image processing. Further, these least-square methods generally work for a single ellipse, or one ellipse at a time, an obvious challenge for application to large-scale imaging scenarios where many ellipses must be detected at once [121]. These methods also require extra levels of pre- and post-processing to enhance the ellipse selectivity of the approach, increasing computational complexity and running time [126]. Finally, the iterative nature of many least-square approaches, combined with sensitivity to initialisation and noise can lead to uncertain convergence [126].

5.1.2 Voting-based Methods

Hough transform-based and voting methods search for local extrema in an N -dimensional (ND) parameter or accumulator space and are derived from Hough’s original patent [127] and early generalisations of the method [128, 129]. The early elliptical Hough transform-based methods, *e.g.* [130, 131], require an $N = n + 5$ dimension accumulator matrix, where n is the dimension of the image, from which a user-defined number of local maxima are extracted. This has huge computational and memory requirements and various adaptations have been developed to reduce the dimension of the accumulator array or decrease the number of computations required. A key improvement in voting approaches was introduced in [19]: the representation of ellipses by a $n + 1$ dimensional accumulator array. Further improvements include the randomised Hough transform [132] applied to the $n + 1$ dimensional algorithm [133].

Hough transform-based and voting methods are still popular for image processing: they are accurate, robust against occlusion and can be implemented very efficiently. However, these

algorithms require an edge image, the production of which may be a very delicate process for noisy images. Also, most Hough transform-based and voting algorithms require a known number of ellipses (and thus accumulator array maxima) to extract, a number that it is often difficult or impossible to determine before segmenting the images, especially in large-scale or high-throughput imaging scenarios. Further, voting-based methods are very sensitive to the intrinsic quantisation of the accumulator array, small changes in quantisation can lead to both false positives and false negatives [119].

5.1.3 Edge-following Methods

Many recent methods focus on the use of elliptical arcs or ellipse segments to detect and refine earlier detection of ellipses, *e.g.* [134–137]. These methods exploit edge connectivity to group edge pixels into arcs, which are then themselves grouped into ellipses. Introduction of self-correcting arc grouping [138] have helped reduce errors in the arc grouping by using a feedback loop to replace low confidence ellipses with high confidence ellipses. These approaches are the most successful for the detection of many ellipses within an image but require strong edge connectivity for the ellipses. Further, these methods often struggle to detect smaller ellipses as the edge segments tend towards the size of noise-derived edge segments.

Finally, there are several notable methods that fall outside of these three categories (see [118] for a recent survey). These methods are often statistical, *e.g.* random sample consensus (RANSAC)-based [139, 140], or heuristic in nature and many of these methods are a hybridisation of the above categories, such as [141, 142].

5.1.4 Contributions and Overview

The method presented in this paper uses aspects of both least-squares methods and Hough-like methods and is a form of object detection that relies on the symmetric nature of ellipses and a foreground-background distinction in the image. Using the framework proposed by [118] the method presented in this paper fills the “Algorithm Selection” layer. As such, we do not focus on image acquisition or preparation but rather on the algorithms strengths and weaknesses when compared to other algorithm choices. It should be noted that our method requires no pre-calculated edge maps (*c.f.* Hough transform-based methods) and is thus more robust to noise and able to detect ellipses on the scale of noise.

Key challenges in ellipse detection are: 1. the sensitivity of many approaches to noise, particularly those that require edge images and may require fine-tuning of such pre-processing stages; 2. the difficulty of detecting many ellipses within a single image; and 3. the requirement of many methods to have large number of pre- and post-processing stages, which may have multiple parameters for each stage.

In this paper, we derive our new technique for the detection and implicit measurement of ellipses. The method takes advantage of the fact that ellipses have an intrinsic level of symmetry allowing our technique to look along a single axis of the shape.

The key contributions of this technique are: 1. the ability to detect ellipses from an image without *a priori* knowledge of the number of ellipses present; 2. robustness against noise; and 3. minimal pre- and post-processing stages and minimal parameter tuning.

Through this paper we will demonstrate the development of this new method through analogous extension of granulometry (see [Granulometry of Circular Objects in 2D Images](#)), a technique that uses only morphological operations and simple mathematical operations. We then provide an algorithmic description of our method and introduce necessary changes for dealing with real and noisy images (see [Methods](#)). We compare our method to a popular elliptical Hough transform, two recent non-Hough methods and a state-of-the-art ellipse fitting method by validating the methods quantitatively on synthetic images, including noisy scenarios (see [Experiments and Discussion](#)). Finally, we demonstrate qualitatively the ability of our method to detect ellipses in a variety of real images (see [Experiments on Real Data](#)).

5.2 Granulometry of Circular Objects in 2D Images

Granulometry is a mathematical method that computes the distribution of ‘granule’ sizes in images through simple morphological operations [35]. This technique has mostly been used in the analysis of images of geological materials, although its power with regard to circle detection and extraction became clear from an early stage (see [143, 144]). Mathematically it can be defined through the following steps:

1. Given an initial circular SE of unit radius, b_0 , we then create a family of SEs, $\{b_k\}$, $k \in \mathbb{N}$, where b_k is defined as

$$b_k = b_0 \oplus_k b_0, \quad (5.1)$$

where \oplus_k denotes k sequential morphological dilations, *e.g.* for $k = 3$, $b_3 = b_0 \oplus b_0 \oplus b_0$.

2. Then, given a binary image, I , we define the set of morphological openings (\circ) of I ,

$$\{O_k\} = \{I \circ b_k\}. \quad (5.2)$$

3. The granulometry function for image I is defined as

$$G_I(k) = |O_k|, \quad (5.3)$$

i.e. the sum of pixel intensities in O_k .

4. Finally, we define the pattern distribution or granulometry curve of the image as the function $PD_I(k)$ given by,

$$PD_I(k) = G_I(k + 1) - G_I(k). \quad (5.4)$$

In order to illustrate the development of our novel method, we will work through how granulometry can be used to detect circles of different sizes in 2D images. There is a simple extension to ellipses in 2D (described below) and a straightforward extension to 3D.

We begin with an image, as shown in figure 5.1a; this image has a variety of circles of different size and position. By opening the image with circular SEs of different radius we produce images such as figure 5.1b. As can be clearly seen, objects smaller than the SE size vanish from the image through opening; classically, the sum of the image intensity is then used to gauge the distribution of objects (figure 5.1d) for all SE sizes, *i.e.* radii, used. We can look more closely by subtracting one opening from the previous opening of our image; we are now left with only ‘granules’ between these two SEs in size (*e.g.* figure 5.1c); classically, each of these subtraction images corresponds to an element of $PD_I(k)$ which, as the first derivative of the granulometric curve, forms a signal where minima denote the existence of objects of that size (*c.f.* equation (5.4) and figure 5.1e).

5.2.1 From Circles to Ellipses

By using different shape SEs this technique can be extended to other shape objects; it is, however, necessary to extend the technique if the objects are not rotationally symmetric. Thus, for a square SE, one would have a granulometric curve for each rotation of the element, producing a granulometric surface.

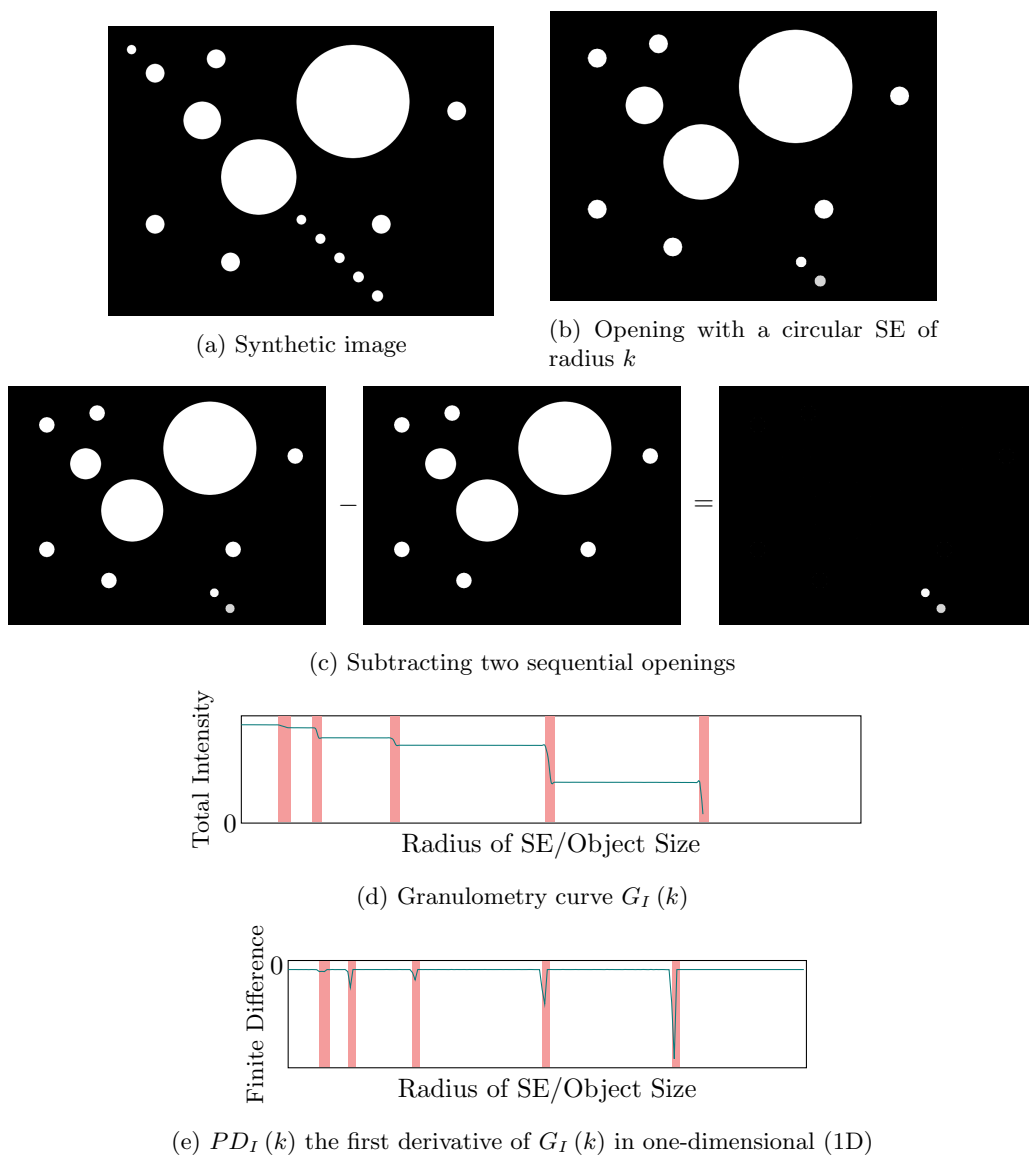


Figure 5.1: Classical granulometry for circles in 2D binary images. (a) Image of fifteen circles of known position and radius, $r \in [3, 30]$ px. (b) Opening with a circular SE of radius k causes circles with radius $r < k$ to vanish. (c) Difference of two sequential openings leaves all circles within that small band. (d) Granulometry curve $G_I(k)$ produced by the sum of intensities over the resultant image for each size SE. Sharp drops indicate the point where lots of objects of that size disappear. (e) $PD_I(k)$, finite difference of the granulometry curve showing local minima indicative of image objects.

This is further extended if the SE shape is eccentric, *i.e.* the width and height are not equal. For example, granulometry for ellipses would require the user to consider rotation of the objects and the differences in the major and in the minor axes. This now becomes a heavy computational task with many morphological operations and a collection of granulometric curves in a large parameter space describing the SE shape.

It is possible to consider the extension of this technique to any arbitrary object that can be used to create a set of SEs dependent on the shape's parameters. As the shape becomes more complex the computations become greater and the data produced more elephantine.

5.2.2 Extension to Greyscale and Higher Dimensions

Not only can granulometry be extended to any arbitrary, parametrisable shapes in 2D binary images but it can be extended to such shapes in greyscale images; it is, in fact, possible to use a greyscale structuring function, *i.e.* a weighted SE, for morphological opening. Granulometry in greyscale images works under the same principle as for binary images, with binary erosions/dilations generalising to local intensity min/max functions, respectively. However, it is recommended that the signal is normalised by the intensity values of each pixel to prevent bright objects appearing more plentiful than dull objects (since the removal of a bright object will produce a larger intensity drop).

The extension to higher dimensions, *i.e.* 3D volumes or 4D+ hyper-volumes causes further memory and computation issues, namely the number of possible parameters for an arbitrary shape increases dramatically. Whilst a sphere, for example, would need only a radius, an ellipsoid would need three radii and two angles of rotation; that's five granulometric curves to analyse concurrently.

5.2.3 Granulometry to Edge Detection and Ranging

First, we note that it is possible to produce a granulometric signal, as a function of SE scale (radius) s , not just for the image $G_I(s)$ but for every pixel $G_{\vec{p}}(s)$, describing the intensity of pixel $\vec{p} = (x, y)$ after morphological opening with a disk-shaped SE b_s of scale s .

These per-pixel signals are step-like functions whose drops occur where s is equal to the radius of the object that the pixel belongs to. Contrariwise, a pixel outside of all objects would show no

step in its signal. If we exchange the morphological opening for a morphological erosion then the signals are still step-like, but now their drops occur at values of s equal to the shortest distance to the boundary of any containing object; pixels outside of all objects still have no step in their signal. This is true in both binary and greyscale cases since erosion with a disk-shaped SE of radius s sets the intensity of each pixel \vec{p} to the minimum intensity in its circular neighbourhood of radius s .

We then take this notion a step further, by replacing the disk-shaped SE of radius s with a symmetrical line of length $2r$ orientated along angle θ . Now the erosion finds the minimum intensity, not in a circular neighbourhood of each pixel but along a line of length $2r$, centred on pixel \vec{p} along a set orientation. The granulometric signal $G_{\vec{p}}(s, \theta)$ for a pixel \vec{p} inside a bright object now drops at a value of s equal to the shortest distance to the object boundary *in the direction of the line's orientation* θ . For a pixel outside of all objects, the signal still shows no step.

By eroding with lines of *all* orientations, we produce a 2D function $G_{\vec{p}}(s, \theta)$ at every pixel; this can be thought of as multiple 1D granulometric signals per pixel. Once we learn the values of s at which each of these signals drop, we have for every pixel the shortest distance to a containing object's boundary along all directions.

One can think of this approach as analogous to radar or sonar: at every pixel an intensity measuring beam is sent in all directions and the signal returned is used to detect edges and range their distance from the pixel. It is for this analogy that we have named our method *Edge Detection and Ranging*.

As a refinement, we replace the line shaped SE with the two endpoints of said line. This provides a level of robustness against speckle noise. Consider, for example, a bright object on a dark background that contains a spurious dark pixel, \vec{p}_{dark} . If a line SE centred on pixel \vec{p} with length $2S$ (where $S \geq |\vec{p} - \vec{p}_{dark}|$) ends at \vec{p}_{dark} , then \vec{p} 's granulometric signal will drop and remain low for all $s \geq S$. This is because the erosion takes the minimum intensity along the line; effectively pixel \vec{p} cannot 'see through' pixel \vec{p}_{dark} . By using the two endpoints we see only a transient dip in signal, rather than a permanent drop, as the endpoints can move past \vec{p}_{dark} and so the signal beyond \vec{p}_{dark} is no longer influenced by \vec{p}_{dark} .

5.3 Methods

In this section we introduce the algorithmic application of these granulometric signals to ellipse detection by Hilbert transform-based edge detection and ranging.

5.3.1 The Algorithm

```

Input :  $I$ 
Output:  $C, M_C, m_C, \theta_C$ 

1 forall  $\theta$  do
2   forall  $s$  do
3      $b_{s,\theta} \leftarrow b_\theta \oplus_s b_\theta$ 
4      $b_{s,\theta} \leftarrow \text{endpoints}(b_{s,\theta})$ 
5      $G_I(s, \theta) \leftarrow I \ominus b_{s,\theta}$ 
6   end
7 end

8  $PD_I(s, \theta) \leftarrow G_I(s + 1, \theta) - G_I(s, \theta)$ 
9  $d_I(\theta) \leftarrow \underset{s}{\text{argmin}}(PD_I(s, \theta))$ 

10  $M \leftarrow \max(d_I(\theta))$ 
11  $\theta_M \leftarrow \underset{\theta}{\text{argmax}}(d_I(\theta))$ 

12  $\theta_m \leftarrow \theta_M - 90$ 
13  $m \leftarrow d_I(\theta_m)$ 

14  $\{C\} \leftarrow \text{regionalMaxima}(M \cdot m)$ 
15  $\{M_C, m_C, \theta_C\} \leftarrow \text{descriptors}(C)$ 

```

Algorithm 5.1: Edge Detection And Ranging for ellipse detection in 2D images as described in the main text.

Algorithm 5.1 shows the algorithm for detecting and quantifying ellipses within a 2D binary image. For all possible radii/axis lengths, s , and all possible rotations, θ , we create a line SE, $b_{s,\theta}$ (line 5), remove all pixels except the line endpoints (line 6) and carry out morphological erosion (denoted by \ominus) on the image, I (line 7). This forms $G_I(s, \theta)$, which may be thought of as a set of 1-D granulometric signals, one for each θ at each pixel. Calculating the first order derivative (or first order finite difference; *c.f.* [Dealing with Noise](#)) with respect to s gives $PD_I(s, \theta)$, from which we determine the distance s at which $PD_I(s, \theta)$ is most negative, along all directions θ and for every pixel. This drop-off distance we call $d_I(\theta)$ (*c.f.* [figure 5.1 \(p. 85\)](#); line 12). The maximum $d_I(\theta)$ represents the major axis length, M (again for each pixel), and the angle along which this maximum occurs is the orientation of the major axis, θ_M . Likewise,

m and θ_m are the minor axis length and orientation. The values of M , m and θ at the centroids of the regional maxima, C , give the descriptors of each detected ellipse (Lines 20–21).

For ellipse detection, the angle that gives the maximum possible length of all angles is selected as the major axis. We have chosen to select the minor axis as the length orthogonal to the major axis; however, a search for the angle that gives the minimum length could also be used. The angle between the two axes could then be used as a validation check that the object is an ellipse.

Regional maxima [144], and their centroids, are then found from an image that is the element-wise multiplication (\cdot) of the length values for the major, M , and minor, m , axis measurements. The value of M , m and θ at the detected centroids gives the approximated ellipse present at each position.

5.3.2 Finding an Ellipse

Figure 5.2 shows graphically the outcomes of the algorithm described above as applied to an example ellipse. As can be seen in figure 5.2b, the major axis peak occurs along the centre of the ellipse (cyan line) at the minor axis. Likewise, a peak occurs in the centre of the ellipse along the major axis in figure 5.2c (cyan line). By multiplying figure 5.2b and figure 5.2c element-wise, it is possible to produce an image with a strong peak at the centre of the ellipsoid. Simple regional maxima can then be calculated to find the centroid.

The ellipse shown in figure 5.2 was produced using pseudo-random values for major axis full-length, minor axis full-length and orientation: 194 px, 116 px and 143° , respectively, and was centred at (100,100) px. Using a version of this code, implemented in MATLAB 2015a [145], our method extracted an ellipse centred at $(100.0 \pm 0.5, 100.0 \pm 0.5)$ px, a major axis of (194.0 ± 0.5) px, a minor axis of (116.0 ± 0.5) px and a rotation of $(144 \pm 1)^\circ$. As this is a simple scenario we compared the output to MATLAB's built-in `regionprops` [67], which returned a centroid of (100,100) px, a major axis of 195.1401 px, a minor axis of 117.4202 px and a rotation of 145.0753° . In figure 5.2D, there is a thin border of difference due to the 1° angular resolution.

5.3.3 Dealing with Noise

The extension of any mathematical method to noisy scenarios is non-trivial. For our method, the reliance on a clean, sharp drop in signal at the appropriate length and rotation is a serious

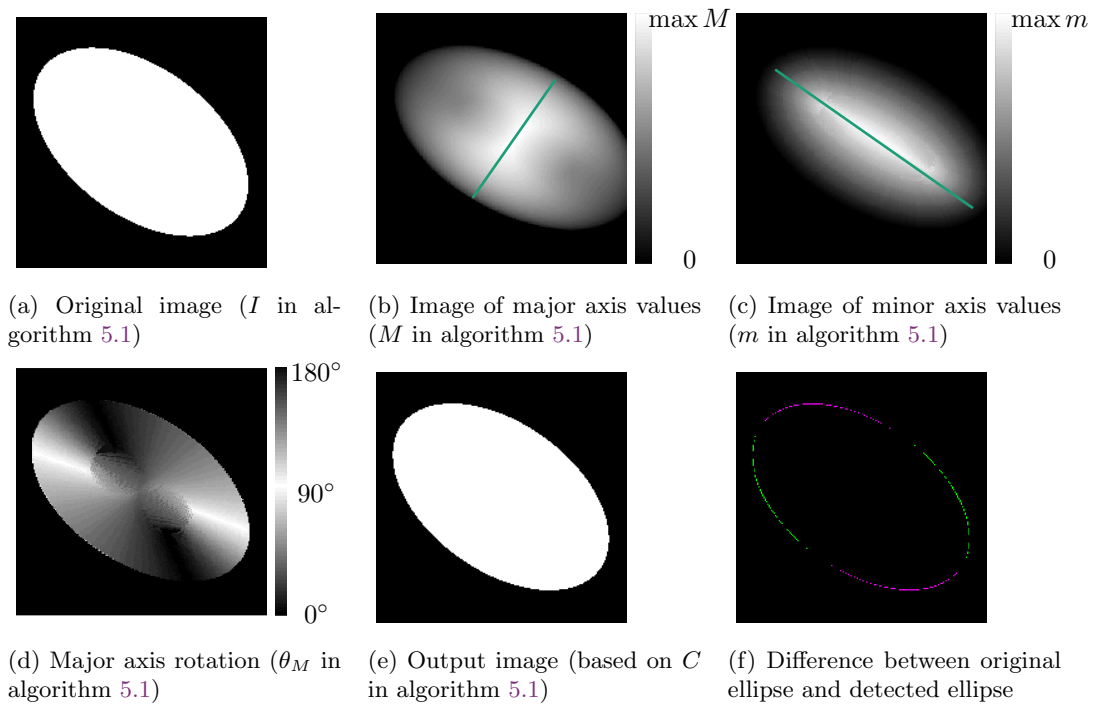


Figure 5.2: Detection of an ellipse with a randomly produced major and minor axis and a random rotation. (b) Image showing the ‘major’ axis value for every pixel of the original image, *i.e.* furthest object edge in any direction. Black indicates zero length while pure white indicates the maximum major axis length. (c) Image showing the ‘minor’ axis value for every pixel, *i.e.* closest object edge. (d) Image showing the rotation of the ‘major’ axis values (to the nearest 1°). (f) The difference between the original ellipse and the detected ellipse (green indicates undetected part of original ellipse and magenta incorrectly identified area).




Figure 5.3: Using the Hilbert transform allows for the identification of any step positions for clean (a) and both blurred and noisy (b). Regional maxima indicate steps down, whilst regional minima indicate steps up. Red curves shown the input signal, blue plots show the Hilbert transform of the input signal.

barrier when speckle noise or blur exists in the image. Speckle noise will create erroneous drops in signal that would otherwise seem clear and sharp whilst blur will weaken genuine drops between foreground and background signals. Various methods of dealing with these noisy signals were considered, including least-squares fitting of the signal to a set of Heaviside functions with drops at all s , the object size. Whilst this, and other methods, worked in various scenarios each had their limitations and all had a large impact on memory and computation time.

The problem can be considered a simple 1D edge detection issue and can be solved using the short response Hilbert transform (see appendix E), as described in 2D by [146]. Step and well signals (*c.f.* edges and ridges in 2D) have a distinct response signal in the Hilbert transform (figure 5.3a) and downward steps can be determined by finding the global maxima. Whilst blurring and speckle noise both affect the clarity of output it is still possible to determine the step position under large amounts of noise and blurring (figure 5.3b). We have limited our algorithm to downward steps, *i.e.* considering objects to be bright on a dark background.

Application of the Hilbert transform for step detection in 1D signals is not only robust but provides extra information: namely, the size of the gap between the upper and lower segments of the step (figure 5.3). This gap information is then used as a way to threshold the length information, removing background speckle and decreasing issues between clustered objects, darker objects or brighter backgrounds, making the method more resilient to image intensities.

This adaptation also allows our method to find objects of varying size, rotation and intensity within noisy and complicated greyscale images. This step gap-based thresholding is the core user parameter of our algorithm.

With the replacement of the finite difference with this Hilbert step detection, we can begin to think of the algorithm in a different way. Consider that for all pixels we consider the intensity signals along a given direction and use the Hilbert transform to range the nearest edge. We do this in all directions and for all pixels, similar to an array of satellites all being used for local radar. Once we have our range information we use further processing to determine which pixels are within ellipses and further to determine the ellipse features. It is for this feature that we have called our method *Hilbert-Edge Detection And Ranging (HEDAR)*.

Algorithm 5.2 shows the algorithm for detecting and extracting the quantified information of ellipses within a noisy, 2D greyscale image (*c.f.* algorithm 5.1 (*p.* 88)). Initial steps are identical to algorithm 5.1 (*p.* 88); however, instead of the finite difference of the granulometric signals, the Hilbert transform is calculated (line 3). The index of the maximum imaginary value of the Hilbert transform is taken as the length of any possible step in that direction (line 3). This forms $L_I(\theta)$, which may be thought of as a 1D function at every pixel giving one length for each θ at each pixel. $L_I(\theta)$ is then used to calculate the difference between the upper and lower portions of the step, $\Delta_I(\theta)$ (line 4). All signals with a step difference greater than the user parameter th are removed (set to zero; line 6). The features of possible ellipses and the centroids of ellipses are calculated as in algorithm 5.1 (*p.* 88) (lines 7–12). Finally we choose to run a simple pruning stage (not shown), that removes ellipses based on their percentage overlap. Algorithm 5.2 is representative of the implemented code used in [Experiments and Discussion](#).

5.4 Experiments and Discussion

It is important for feature detection methods to have quantified validation of annotated results, especially if the method may then be used on large-scale unknown data. In this section we present results from a large number of experiments on synthetic and real data, analysing the running time, accuracy of results and robustness against noise.

5.4.1 Comparator Methods

We have compared our algorithms performance against four alternative ellipse fitting or detection methods. First, we have selected an efficient and robust randomised Hough-based method based

Input : I, th
Output: C, M_C, m_C, θ_C

- 1 Algorithm 5.1: lines 1–7
- 2 **for** $\theta_i \in \theta$ **do**
- 3 $L_I(\theta_i) \leftarrow \underset{s}{\operatorname{argmax}}(\operatorname{imag}(\operatorname{hilbert}(G_I(s, \theta = \theta_i))))$
- 4 $\Delta_I(\theta_i) \leftarrow \operatorname{mean}(G_I(s \geq L_I(\theta_i), \theta = \theta_i)) - \operatorname{mean}(G_I(s < L_I(\theta_i), \theta = \theta_i))$
- 4 **end**
- 5 $L_I(\Delta_I < th) \leftarrow 0$
- 6 $M \leftarrow \max(L_I)$
- 7 $\theta_M \leftarrow \underset{\theta}{\operatorname{argmax}}(L_I(\Delta_I \geq th))$
- 8 $\theta_m \leftarrow \theta_M - 90$
- 9 $m \leftarrow L_I(\theta_m)$
- 10 $\{C\} \leftarrow \operatorname{regionalMaxima}(M \cdot m)$
- 11 $\{M_C, m_C, \theta_C\} \leftarrow \operatorname{descriptors}(C)$

Algorithm 5.2: HEDAR for ellipse detection in 2D images as described in the main text.

on [19, 133] (henceforth referred to as elliptical hough transform (Hough)); this method is well cited since publication and provides a strong example of Hough- and voting-based techniques. Hough- and voting-based techniques have a key weakness in that *a priori* knowledge of the number of ellipses is required otherwise some thresholding of which objects to consider and which to ignore must occur. As such, we have, where appropriate, tested both scenarios. It should be noted that our method can be run without such *a priori* knowledge. For all experiments using Hough we have included the running of Canny edge detection with automatically calculated thresholds; the algorithm is then run on the output edge image.

Second, we have chosen to compare our ellipse detection method to the results of a least-squares-like ellipse fitting method [147] (hyper renormalisation (HyperRN)); fitting methods are generally designed to fit a single ellipse, therefore we have not applied this method to all of our experiments. For all experiments using HyperRN we have included the running of gradient-based edge detection; the algorithm is then run on the output edge points in space.

Finally, we compare our method to two recent arc-fitting, geometric approaches as presented in [126] and in [137]. Like the Hough, these methods, ElliFit (ElliFit) and fast and effective ellipse detector (YAED) respectively, rely on an edge image.

To detect multiple ellipses, ElliFit iterates through each connected component of the edge image to determine whether or not an ellipse is present in that area [126]. For all experiments using

ElliFit we have included the running of boundary detection, the algorithm is then run on the connected edges. Unlike the previously described methods, ElliFit does not require any parameters however we include a maximum length parameter, which filters all detected ellipses greater than it, to make the method more comparable to the other methods.

Similarly, YAED iterates through detected arc segments in the edge image and by pairing these segments determines whether or not an ellipse is present in that area [137]. Like ElliFit, YAED does not require any parameters however we include a maximum length parameter, which filters all detected ellipses greater than it, to make the method more comparable to the other methods.

5.4.2 Running Time and Memory Use

The efficient, vectorised MATLAB implementation of our algorithm uses morphological operations and the Fourier transform that introduce a computational and memory bottleneck. This bottleneck makes our method more than an order of magnitude slower than our comparator methods. With alternative language implementations the relative gap could be reduced but the HEDAR technique will always be more computationally complex as it acts on all pixels in the image, not a subset, such as only the edge pixels.

Further, the use of morphological operations requires the signals for all pixels to be held in memory, this limits the size of the image and maximum ellipse size that can be detected on a standard desktop PC with this implementation. As such, our synthetic experiments have been run for ellipses with major axis lengths less than 50 px and images less than 128 px \times 128 px (unless otherwise stated).

Figure 5.4 (*p. 96*) shows the performance of our full Hilbert-based method (HEDAR; as described in algorithm 5.2 (*p. 93*)), the elliptical hough transform (Hough) (as described in [19, 133]), both with and without a given number of objects (Hough (w/ N) and Hough, respectively), the ellipse-fitting hyper renormalisation (HyperRN; as described in [147]), ElliFit (as described in [126]) and fast and effective ellipse detector (as described in [137]). All methods were implemented in MATLAB [67] except YAED, which is provided by the authours as a C++ project.

As expected the computational running for HEDAR is larger than the running time for the other methods, especially at small image sizes (figure 5.4a (*p. 96*)). HEDAR does, however, run in constant time for any number of objects in the image (figure 5.4b (*p. 96*)), unlike edge-based

methods where, with more objects, there are more edge pixels to process. This lends the use of HEDAR methods towards busy images with many objects, *e.g.* biological images, as opposed to smaller images with few objects, *e.g.* computer vision tasks.

5.4.3 Detection Accuracy


To investigate the detection accuracy of the proposed method we have emulated the synthetic experiments first shown in [137]. We ran our method and the comparators on automatically generated synthetic images of size 32×32 each containing a single ellipse without noise.

Figures 5.5 and 5.6 show how robust these methods are to size and eccentricity, respectively. In these plots each ‘pixel’ is given an intensity value that corresponds to the Jaccard similarity index [148]* between the input image and the result of that algorithm. Hence, a bright pixel indicates that the given method was able to correctly identify the ellipse. In figure 5.5 (p. 97) each pixel represents the minimum Jaccard similarity for ellipses of a given major axis length and axis ratio (minor axis over major axis) over all rotations. A deep blue pixel may indicate that no ellipse was detected or that more than one ellipse was detected, under which situations a Jaccard similarity of zero is assigned. In figure 5.6 (p. 98) each pixel represents the minimum Jaccard similarity for ellipses of a given major axis length and orientation over all axis ratios, regardless of the number of ellipses detected. Figure 5.7 (p. 99) shows four random examples of ellipses as generated in the above experiments.

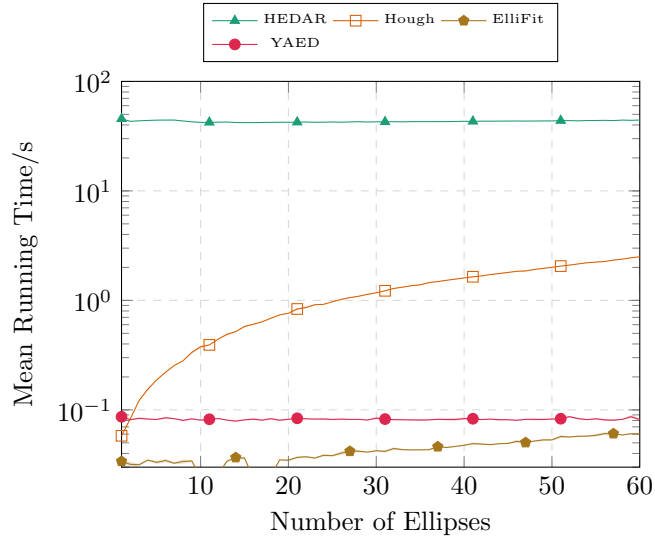
As expected, all algorithms, where ellipses are identified, show a level of robustness to orientation, regardless of eccentricity. Further, as the axis ratio tends towards zero and the ellipses tend towards line-elements, all methods begin to experience difficulties; however, figure 5.5 (p. 97) shows a better ability to detect such cases for HEDAR and Hough, as opposed to ElliFit.

Although ElliFit shows better quality results, the large space of zero-valued pixels for small size ellipses indicates an inability to deal with ellipses that might be considered on the level of noise. Both HEDAR and Hough (with a given number), are, however, able to detect and measure such small-scale ellipses. It should be noted that Hough without a given number of ellipses is unable to correctly identify only a single ellipse when run with default parameters over this range of images.

* For completeness the mathematical definition of the Jaccard index has been provided in appendix D (p. 150) along with justification for the choice of this metric over others.



(a) Performance time against image size



(b) Running time against number of ellipses within a fixed-size image

Figure 5.4: Comparison of computational running time for methods as described in the text. (a) Mean running time against image size in pixels for ten trials. Each trial image contained a single ellipse of the same properties. (b) Mean running time against the number of fixed-size, non-overlapping ellipses in a fixed-size image. Each trial contained a specific number of ellipses of the same size with varied position and orientation. Standard error upon the mean was calculated for all data points on all plots and are shown as thin bands.




Figure 5.5: Algorithm response with respect to major axis length (vertical axis; from 3 to 39 pixels) and axes ratio (minor axis length/major axis length; horizontal axis). The proposed method, HEDAR, (a), Hough with no N (b) and ElliFit (c). (d) YAED failed to reliably identify ellipses of this size scale due to the small number of edge pixels.

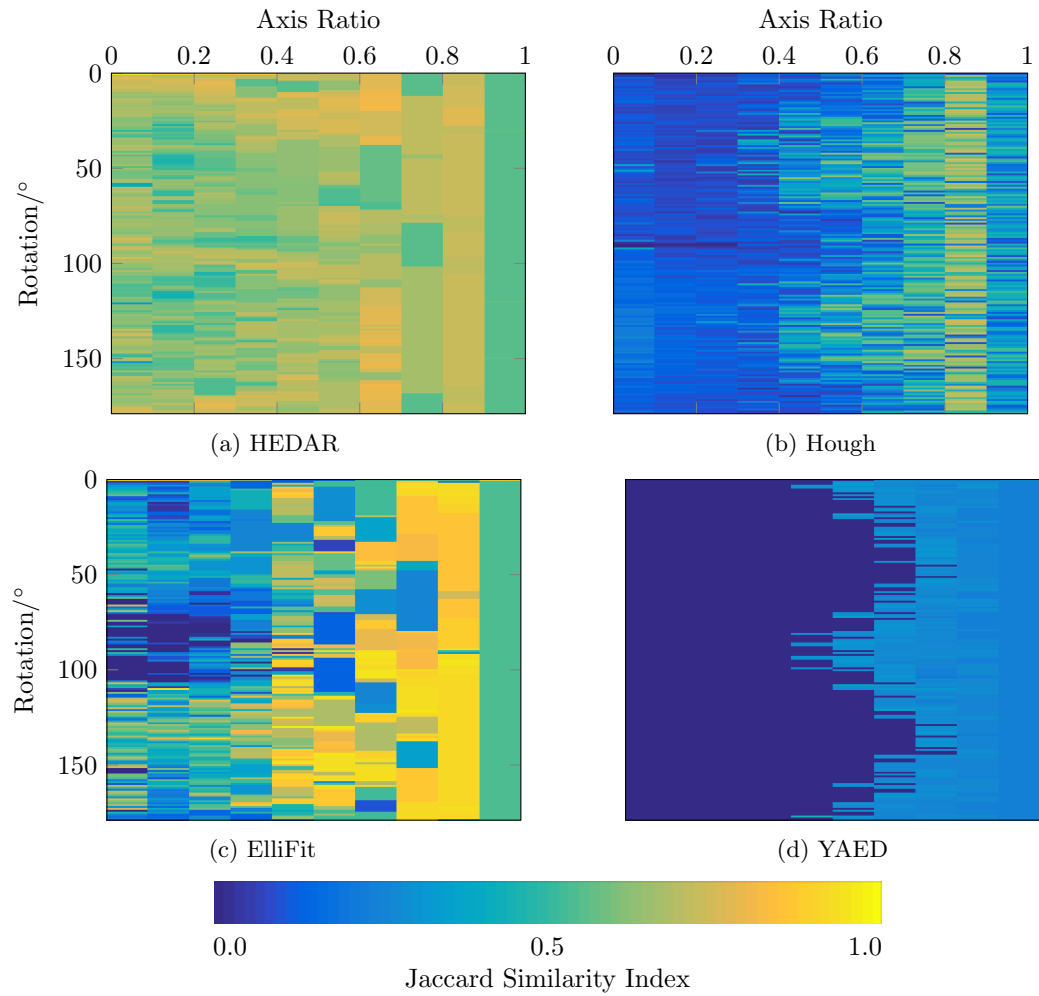


Figure 5.6: Algorithm response with respect to rotation (vertical axis; from 0° to 179°) and axes ratio (minor axis length/major axis length; horizontal axis). The proposed method, HEDAR, (a), Hough with no N (b), ElliFit (c) and YAED (d).

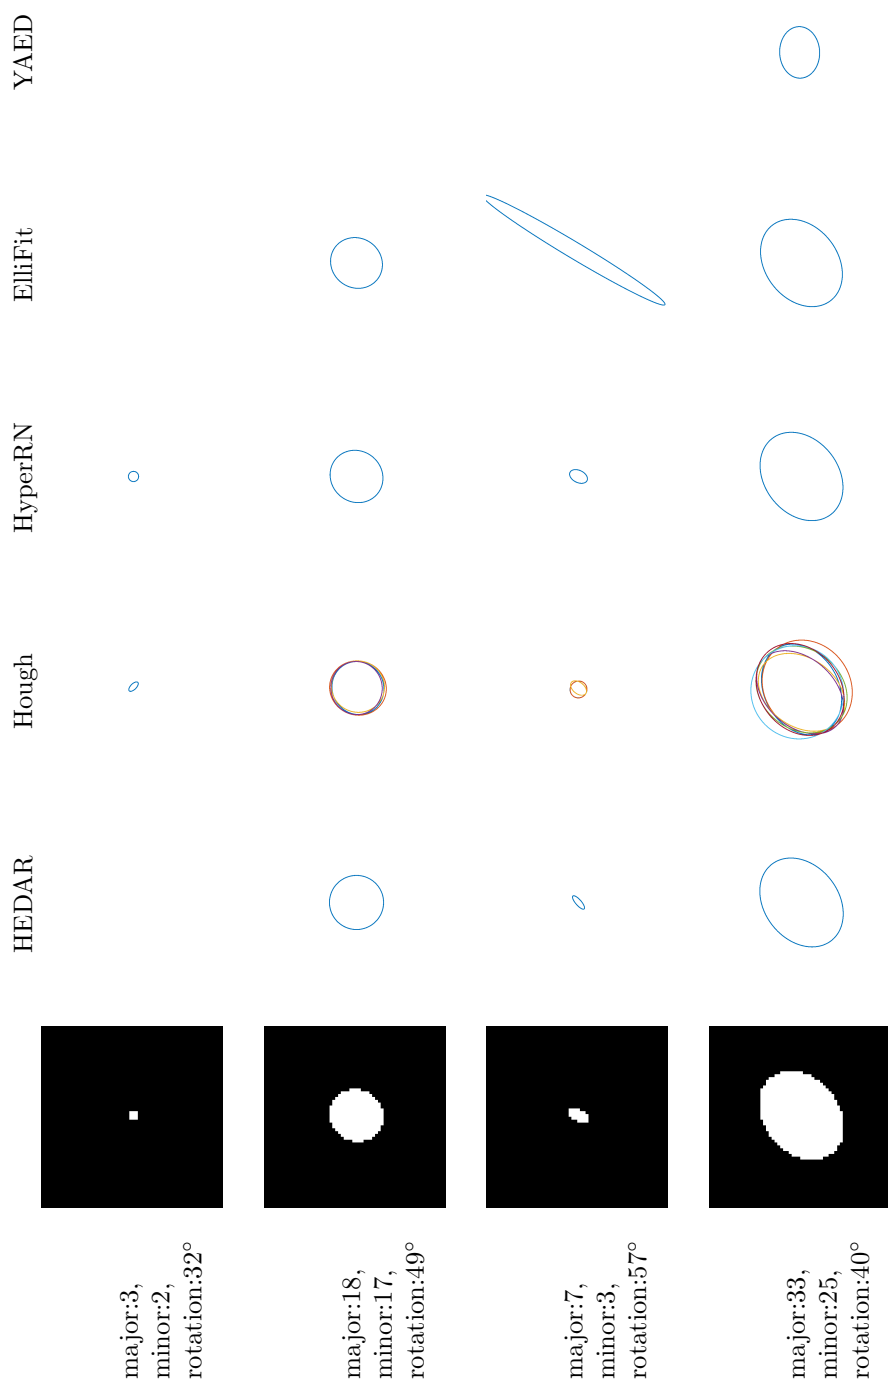


Figure 5.7: Examples of the results for each algorithm for four randomly chosen ellipses, as used in figures 5.5 and 5.6 (p. 97) and (p. 98).




Figure 5.8: HEDAR can robustly detect the number of non-overlapping, fixed-size ellipses in an image. Detected number of ellipses against actual number of ellipses for HEDAR and comparator methods. Standard error upon the mean is presented for all plots as a band.

Not only is the overall accuracy of any method important for its application but so too is the method's ability to deal with many objects. Figure 5.8 shows how HEDAR is able to correctly detect the number of ellipses in an image. We ran HEDAR on a set of synthetic images with a known number of small ellipses. Each trial image was of size 256×256 and contained n small ellipses of varied size, position and orientation. The plot shows how, as the number of ellipses increases, both Hough and ElliFit struggle to correctly count the number of objects, whilst HEDAR is able to maintain a correct count for most n .

5.4.4 Robustness Against Noise

In order to determine the usefulness of new methods, they must be tested against issues such as noise and image complexity. Figure 5.9 shows the mean Jaccard similarity index of the results for synthetic images of known peak signal-to-noise ration (PSNR) for three different types of noise: additive Guassian, multiplicative Gaussian (speckle) and salt & pepper noise. As can be




Figure 5.9: Comparing the robustness of the proposed method over a range of noisy images for different noise types. Jaccard similarity index against peak signal-to-noise ratio in decibels; standard error upon the mean is presented for all plots as a band.

seen, the proposed method is robust to large amounts of noise. Figure 5.10 shows examples of low (20 dB) and high levels of noise (0 dB).

5.4.5 Experiments on Real Data

This section presents the results of our method, HEDAR, on real data and qualitatively compares the outputs to the Hough, ElliFit and YAED methods.

All method parameters have been included in the paper repository. HEDAR requires a thresholding value related to local object brightness and, optionally, takes a maximum ellipse size in pixels; the Hough method requires a maximum ellipse size and an accumulator thresholding parameter; whereas, the ElliFit and YAED methods require only an optional maximum ellipse size for post-processing.

All images were converted to double format, with values between 0 and 1. Images with bright background were also inverted. Before Hough, ElliFit or YAED analysis, images underwent canny edge detection using the automatic MATLAB lower and upper thresholds.




Figure 5.10: Examples of HEDAR run on synthetic noisy images.

Figure 5.11 shows the results as applied to biological images (figure 5.11a), posed images (figure 5.11b) and real-world images (figure 5.11c). As can be seen, HEDAR outperforms alternative methods in accurately capturing the number of elliptical objects in the image. Clearly, HEDAR does not return perfect results and issues of connected objects, such as the cherry stems, and uneven object brightness, such as the egg shells, distort the resultant ellipse. However, we feel the results shown indicate the strengths of the HEDAR approach over classical and state-of-the-art approaches. Other examples can be seen in figure 5.12 (*p.* 105).

Figure 5.13 (*p.* 105) demonstrates a key limitation of the HEDAR method: that objects of interest are considered bright on a dark background. Obviously many images, such as that in figure 5.13 (*p.* 105), have both bright and dark objects on a medium background. HEDAR is able to capture the bright object when run on the original image and the dark object when run on the inverted image. It is the authors' opinions that further development of the Hilbert-edge detection step (*Dealing with Noise*) will allow the HEDAR approach to capture both bright objects on darker backgrounds and dark objects on brighter background within a single image.

5.5 Automatic Nuclei Counting in Fluorescent Microscopy

5.5.1 Background

Quantitative biology relies heavily on the use of fluorescence microscopy and associated tools. Key to a large number of assays is the ability to detect and count the number of cells or nuclei in an image or region of interest. Frequently, such as for high-throughput and large-scale imaging experiments, the first step for automated image analysis is nuclei or cell detection and segmentation. As such, and given the broad range of fluorescence microscopy technologies, methods and assays, the scientific community requires new automatic approaches to this essential step.

Manual cell counting and annotation, by visual inspection, is difficult, labour intensive, time-consuming and subjective to the annotator involved. As such, a significant number of groups have proposed a variety of approaches for automatic nuclei detection in fluorescent microscopy images. In recent years proposed methods have included:

- variations upon thresholding-based segmentation, which are limited in performance due to intensity inhomogeneity and nuclei/cell clustering [150];

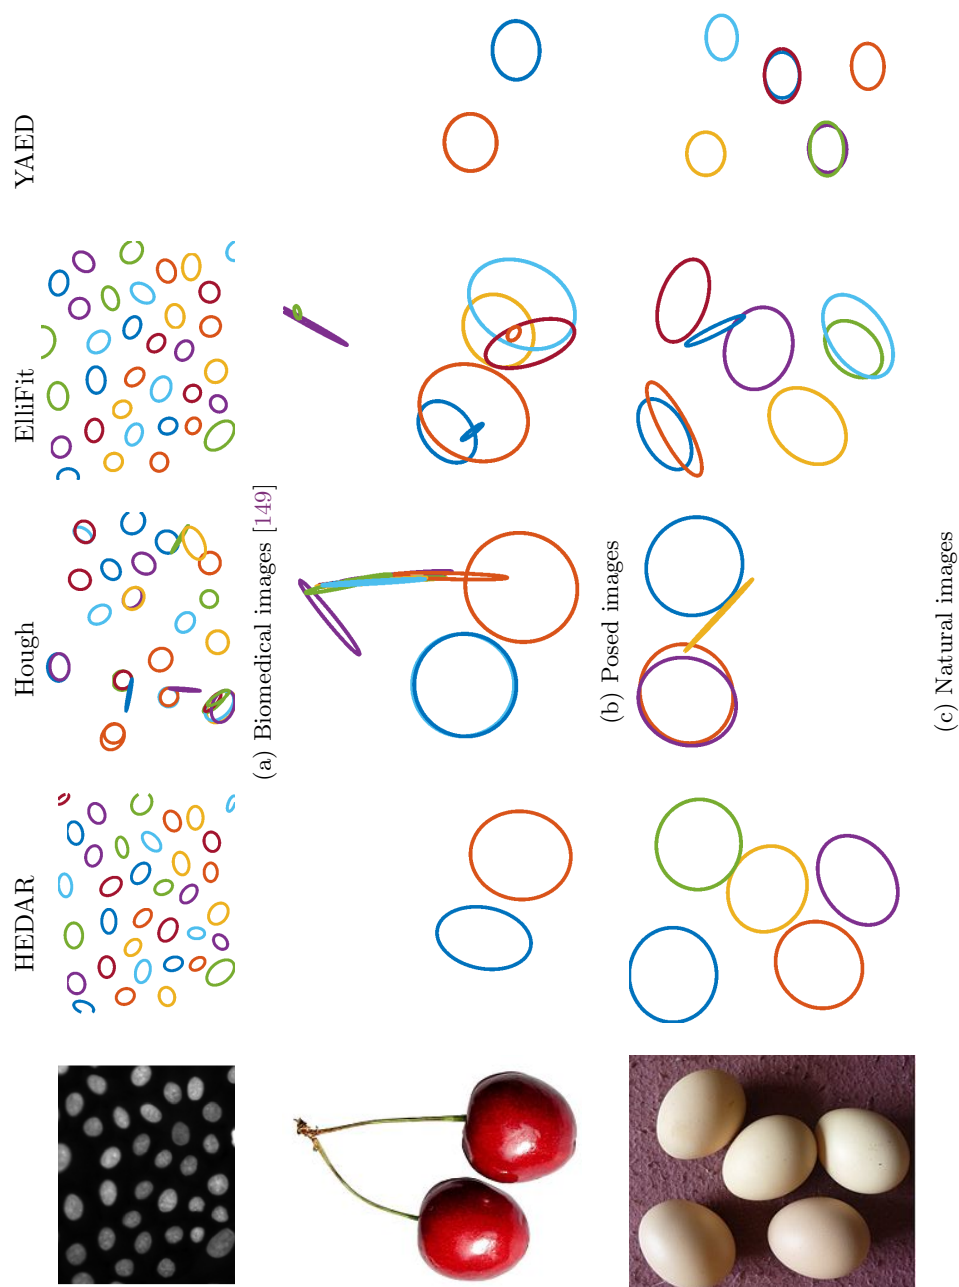


Figure 5.11: Examples of HEDAR and comparator methods run on real images. (a) Cell nuclei captured by fluorescent microscopy [149]. (b) A 'posed' photograph of cherries. (c) A more 'natural' photograph of eggs. All images were converted to greyscale before analysis.




Figure 5.12: Examples of HEDAR on real greyscale images with bright and dark objects of interest. (a) A coin. (b) A Frisbee. (c) A comet. (d) Mars.




Figure 5.13: Example of HEDAR on real greyscale images with bright and dark objects of interest. (a) Using the original image. (b) Using the inverted image (results shown on the original image).

- H-minima transforms [151];
- voting-based techniques [152], which both show good results but are sensitive to parameters;
- gradient vector flow tracking and thresholding [153, 154], such PDE-based methods require strong stopping and reinitialisation criteria, set in advance, to achieve smooth curves for tracking;
- using LoG filters [20], which is low computational complexity but struggles with variation in size, shape and rotation of objects within an image; and
- graph-cut optimisation approaches [155], which requires the finding of initial seed points for each nucleus.

In this case study we apply a new ellipse detection algorithm, Hilbert-Edge Detection And Ranging (HEDAR) [156] to fluorescence microscopy images from the Broad Bioimage Benchmark Collection [21]. HEDAR requires no preprocessing on these data sets and consists of two stages: a) automatic detection of ellipse-like objects and simultaneous annotation of position, orientation, major and minor axes; b) simple post-detection pruning.


We show that HEDAR is a competitive algorithm for counting nuclei in fluorescent microscopy images of cells. We also report those scenarios under which HEDAR fails to count nuclei so that the reader can easily determine the possible limitations when using this method.

5.5.2 Experiments & Results

In this section we use one additional step in the processing of images: in post-processing, we remove all unduly small objects based on minimum nuclei size and then cycle through all detected ellipses, removing the more overlapped ellipse of any pair overlapping by greater than 50%.

We have tested HEDAR on three data sets, all freely available from the Broad Bioimage Benchmark Collection [21]. The first set, BBBC001v1 [157], consists of six Hoechst 33342-labelling channel images of human HT28 colon cancer cells (figure 5.14b). All images are 512×512 pixels and stored as 8-bit TIFF files.

The second data set, BBBC002v1 [157], consists of 50 images each of five distinct samples of *Drosophila melanogaster* Kc167 cells, again with Hoechst 33342 labelling (figure 5.14b). All images are 512×512 pixels and stored as 8-bit TIFF files.



(a) The Broad Bioimage Benchmark Collection dataset BBBC001 contains images of the HT29 human colon cancer cell lines labelled by Hoechst 33342, a DNA labeller used to highlight the nucleus.

(b) The Broad Bioimage Benchmark Collection dataset BBBC002 contains samples of *Drosophila melanogaster* Kc167 cells also labelled by Hoechst 33342.

(c) The Broad Bioimage Benchmark Collection dataset BBBC004 contains synthetic images generated by the SIMCEP programme. This image is from the dataset generated with a 60% chance of generated nuclei overlap.

Figure 5.14: Examples of images from the Broad Bioimaging Benchmark Collection, as used in this thesis.

The final data set, BBBC004v1 [158], consists of 20 images each of five distinct subsets of synthetic fluorescent cell populations with different clustering/overlapping probabilities generated with SIMCEP [159, 160] (figure 5.14b). All images are 950×950 pixels and stored as 8-bit TIFF files.

Optimal parameters for HEDAR (maximum object size and step thresholding level) were manually determined and we used a single set of parameters for each distinct set of images.

Figure 5.15 shows examples of HEDAR results on BBBC001, Hoechst 33342-labelling of human HT28 colon cancer cells. Most nuclei are accurately counted (figure 5.15a) giving a low mean percentage error (table 5.1). HEDAR fails to detect a small number of cells, specifically irregular cells (*e.g.* the nucleus in figure 5.15b) and clustered or overlapping cells (*e.g.* figure 5.15c), including cells that have recently undergone mitosis.

Figure 5.16 (*p.* 109) shows examples of HEDAR results on BBBC002, *Drosophila melanogaster* Kc167 cells. Whilst HEDAR detects most nuclei (figure 5.16a (*p.* 109)) there are more false negatives than in the other datasets. Here HEDAR fails due to the large and complex nature of these nuclei, with cells appearing multi-nucleated (figure 5.16d (*p.* 109)) and with large,

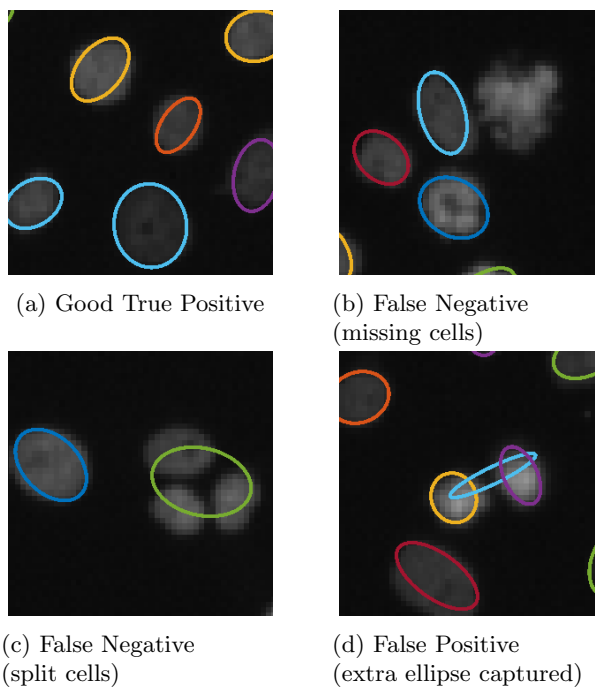


Figure 5.15: Examples of results for benchmark dataset BBBC001.

Data Set	Mean % Error (StDev)		Human % Variation
	HEDAR	[157]	
BBBC001v1	1.93 (1.25)	6.2	11
BBBC002v1	12.37 (14.16)	17	16
BBBC004v1	3.59 (0.03)	-	-

Table 5.1: Mean percentage counting error for all datasets.

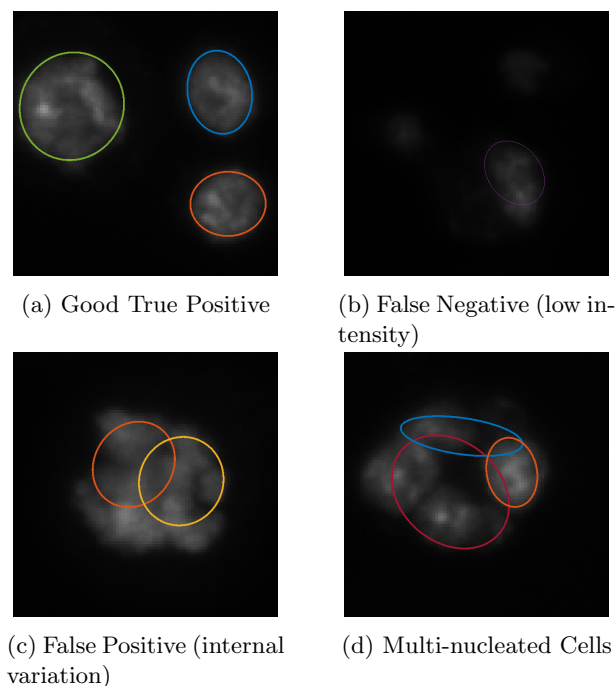


Figure 5.16: Examples of results for benchmark dataset BBBC002. Note how the large, high variation cells cause issues for this method of ellipse detection.

inhomogeneous nuclei (figure 5.16c). Despite this, the mean percentage error is in line with other published results and variation between the human ground truth counts (table 5.1 (p. 108)).

Using the BBBC004 dataset, we were also able to characterise how well the proposed method can cope with increasing degrees of clustering. Figure 5.17 shows the results of HEDAR on each subset of synthetic images from BBBC004; for comparison, the reported results of two other methods are also plotted [161, 162]. As can be seen, HEDAR-based counting is much less affected by clustering and overlap than the comparison methods.

These results show that HEDAR is a competitive tool for nuclei counting and detection. HEDAR performs at the same or better level of error than commonly used solutions and recent state-of-the-art solutions and gives results that agree with ground truth data with a high level of accuracy.




Figure 5.17: Mean percentage error in nuclei counting results for dataset BBBC004, a synthetic dataset with images of increasing overlap probability.

5.6 Conclusion

We have detailed a new method of ellipse detection in images, giving the mathematical form of this approach and describing the extension needed to cope with noisy image data. We have compared this algorithm to an efficient and accurate Hough transform-based ellipse detector, two state-of-the-art geometric ellipse detection methods and a state-of-the-art ellipse fitting technique and shown that, whilst being computationally heavier, our novel method is invariant in running time with the number of ellipses or the levels of noise and is as or more robust to noise in images and clustering of objects than the alternative approaches.

This new method can work with no *a priori* knowledge; this is unlike other methods, such as Hough-based approaches or convolution approaches where some knowledge of the elliptical objects or the image properties is required to get useful information from the image.

In this paper we have highlighted both the benefits of this method and those scenarios where this approach may falter. We believe this information to be essential before users should use this method on their own data. We expect that further development of this approach, such as optimisation for larger nuclei and adaptations to the Hilbert-edge ranging stages, will further improve the performance and the applicability of this technique.

In future work we hope to minimise the computational complexity by 1) rotating the image rather than each SE, 2) using a sliding window approach to replace the multiple SEs needing for

the opening approach, 3) incorporating multiple signals into a single Hilbert transform to reduce the number of Fourier transforms required.

Epilogue

In this chapter, I have introduced a novel method for the detection, counting and extraction of ellipses in greyscale images. I have highlighted the benefits and limitations of this new method and validated the approach on synthetic and real data, comparing against several other methods. This new approach, HEDAR, performs satisfactorily in all these experiments and is shown to be a more powerful ellipse counter than these comparator approaches.

I have described the application of this method to nuclei counting and detection in fluorescence microscopy. This task has the challenges of low signal-to-noise ratio, high heterogeneity within and between the objects of interest and often high clustering rates. Although many methods for this problem already exist, their success is often heavily linked to the dataset used for development and new, robust approaches to this issue are still needed. A key advantage of HEDAR is the lack of a reliance on edge detection, a step which is often very sensitive to contrast artefacts and image noise. I have shown HEDAR to be a satisfactory ellipse detector with a better ability to count nuclei in clustered or noisy scenarios compared to comparator approaches.

In the last two chapters, I have described two methods to object detection and counting in fluorescence microscopy covered in the (sub-)pixel and supra-pixel range. I have shown that both methods are satisfactory object detectors and, further, that HEDAR is powerful when it comes to counting elliptical objects in noisy and complex images.

In the next chapter, I move beyond object detection to object measurement, again utilising mathematical morphology, to quantify medical images. This approach again builds upon granulometry but focusses on measuring the thickness of, in this case, vascular systems.

Chapter 6

Granulometry for Thickness Measurement of Cerebral Vasculature

Prologue

In the previous chapters, I have introduced an image enhancement method (chapter 3 (p. 20)) and two object detection methods (chapters 4–5) built upon the ideas of mathematical morphology. In this chapter, I illustrate the possibility of using this approach for the quantitative measurement of objects of interest, here, in the context of measuring the thickness of vasculature for the identification of abnormalities, such as cerebral aneurysms.

First, I introduce the challenges for image-based quantification of vascular abnormalities, before describing my approach to segmenting and quantifying the vascular network from a series of magnetic resonance angiography images. I validate these measurements by comparing them to manual measurements across the network before proposing future extensions to this pipeline.




Figure 6.1: Maximum intensity projection of TOF MRA data showing the brain vasculature and other vessels and features. Key vessels have been highlighted.

6.1 Introduction

As one of the world's leading causes of death, understanding, diagnosing and treating the causes of stroke is of vital importance [163]. Strokes may be caused by clotting or haemorrhaging within the cerebral vasculature, both of which have an effect on the shape and thickness of the blood vessel network. Digital image analysis of MRA data, such as that in figure 6.1, allows us to measure and compare morphology and dynamics within brain vasculature to identify changes within a patient's vasculature and identify abnormalities between patients, *e.g.* [164, 165].

Improved image analysis tools will allow for enhanced research into the development and formation of intracranial vascular abnormalities and characteristics predisposing aneurysm rupture or clotting. Such tools can be used to support the extensive literature surrounding modelling of such vascular networks. Further, these methods can be utilised directly in diagnostic screening and,

as the volume of medical imaging data increases, pre-screening tasks that may help neurologists and surgeons focus on areas of the MRA images [166].

The current most popular existing tool for quantification and analysis of cerebral vasculature images is the open source Vascular Modelling Toolkit (VMTK) [29, 167], which is capable of segmentation of regions of vasculature using level-set approaches, geometric analysis of extracted surfaces and modelling with computational fluid dynamics.

Unfortunately, VMTK, like many other existing tools, focusses on the user identifying and defining volumes of interest [168–170]. This decreases its usefulness for automated analysis, as global approaches for the segmentation of all vasculature, calculation of metrics across the brain and highlighting of areas with potential abnormalities are required.

We propose that, unlike existing softwares, a global approach may be used to identify volumes as the candidate regions for other, local approaches [171]. Such global metrics might include image vesselness or blobness measures, such as [41], which can then be used to identify regions for local image and morphology features to be calculated, as in [172].

This paper introduces a new global approach for the interactive, semi-automated segmentation of complete brain vasculature, measurement of vessel thickness and the visual highlighting of potential abnormalities.

Through this paper, we will describe this global image analysis approach (section 6.2), introducing the techniques applied at each stage, before demonstrating the results on a dataset of real MRA images with and without aneurysms (section 6.3 (*p.* 118)). We show how this approach can highlight aneurysms and aneurysm-like regions of interest and discuss its further development (section 6.4 (*p.* 120)).

6.2 Methods

Our interactive vasculature segmentation and automated measurement and aneurysm highlighting framework follows a series of steps, as illustrated in figure 6.2. The following sections detail the method for semi-automated vascular segmentation, thickness measurement, putative aneurysm highlighting and further analyses.

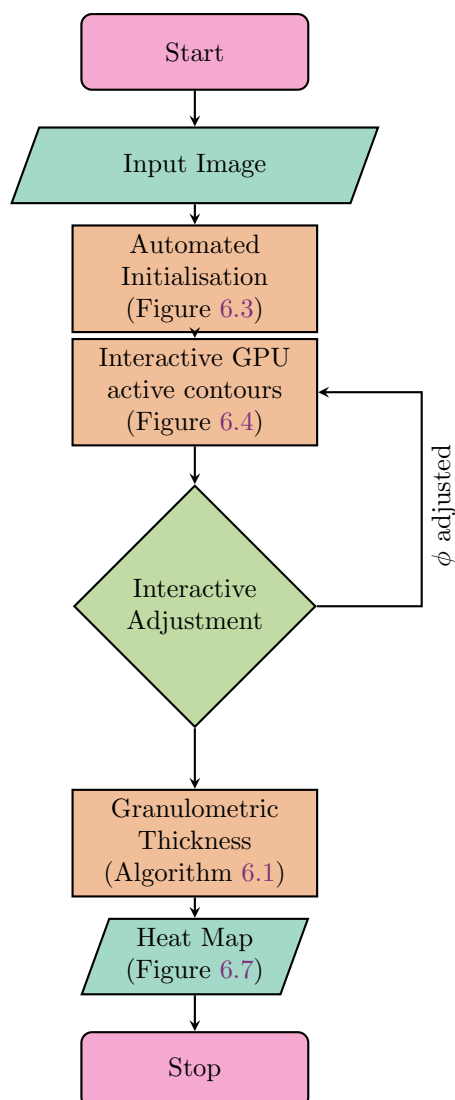


Figure 6.2: Flow chart of the proposed global approach. The raw MRA image is processed to automatically detect seed points (figure 6.3) that are used for automated GPU active contour segmentation (figure 6.4). The user can then choose to keep the segmentation or manually edit the level set contour. The final segmentation then undergoes granulometric thickness measurement (algorithm 6.1) before being visualised as a heat map across the surface of the vasculature (figure 6.7).

6.2.1 Automated and Interactive Vasculature Level-Set Segmentation

To simplify the process for the user, we automatically identify seeding shapes for the initialisation of the level-set segmentation. These seeds are calculated by detecting the blood vessels that enter the base of the brain, *e.g.* the internal carotid arteries (see figure 6.1). Ideally these should be four bright spots on the first slice of the image; however, in practice, we find that one or more spots may be dull or appear missing from the first slices; we also often capture additional vessels that are bright and low down the stack. For our segmentation to work, we need to capture at least two arteries from the internal carotid and vertebral arteries. To automatically detect these arteries, we identify all areas with intensity in the top 5% of a maximum projection of the base 5% slices, *e.g.* the bottom five slices of a 100 slice stack. We then prune those areas based on their size (area ≥ 30 pixels), eccentricity (≤ 0.9) and solidity (defined as the ratio of the actual area to the area of the convex hull; ≥ 0.8) and whether or not they are in the central region of the MRA. Figure 6.3 shows some representative results from this stage. The thresholds for size, eccentricity and solidity were empirically chosen.

These seeding shapes are used as the initial ϕ image for our interactive GPU active contours (IGAC) segmentation method [6] based on [173]*. The IGAC programme can act in an interactive manner, where the user can fine-tune three intuitive parameters and utilise the novel brush functions presented in [6], or, as used initially here, in an automated way with set parameters that work well for a broad range of MRA images. For the first, automated run of the segmentation programme, we automatically remove all connected components not connected with one or more of our initial seeds.

The user is then presented with the segmentation results for confirmation of correct segmentation; should the user declare the segmentation to be poor, they are directed to the interactive variation of the segmentation tool with the current results as the initialising ϕ . At this stage, no automated pruning is carried out and the user is given full control of the final, segmented vascular network.

6.2.2 Measuring Vessel Thickness with Mathematical Morphology

Once segmented the resultant binary image undergoes analysis by mathematical morphological processing, a form of granulometric analysis used to determine the vessel thickness at any point, as shown in algorithm 6.1 (*p.* 118). Iterative morphological opening with spherical kernels allows the accurate measuring of thickness at all positions in the vasculature network; for MRA data, which is usually anisotropic in z , elliptical kernels are used.

* See appendix F (*p.* 152) for a brief summary of this level-set approach.

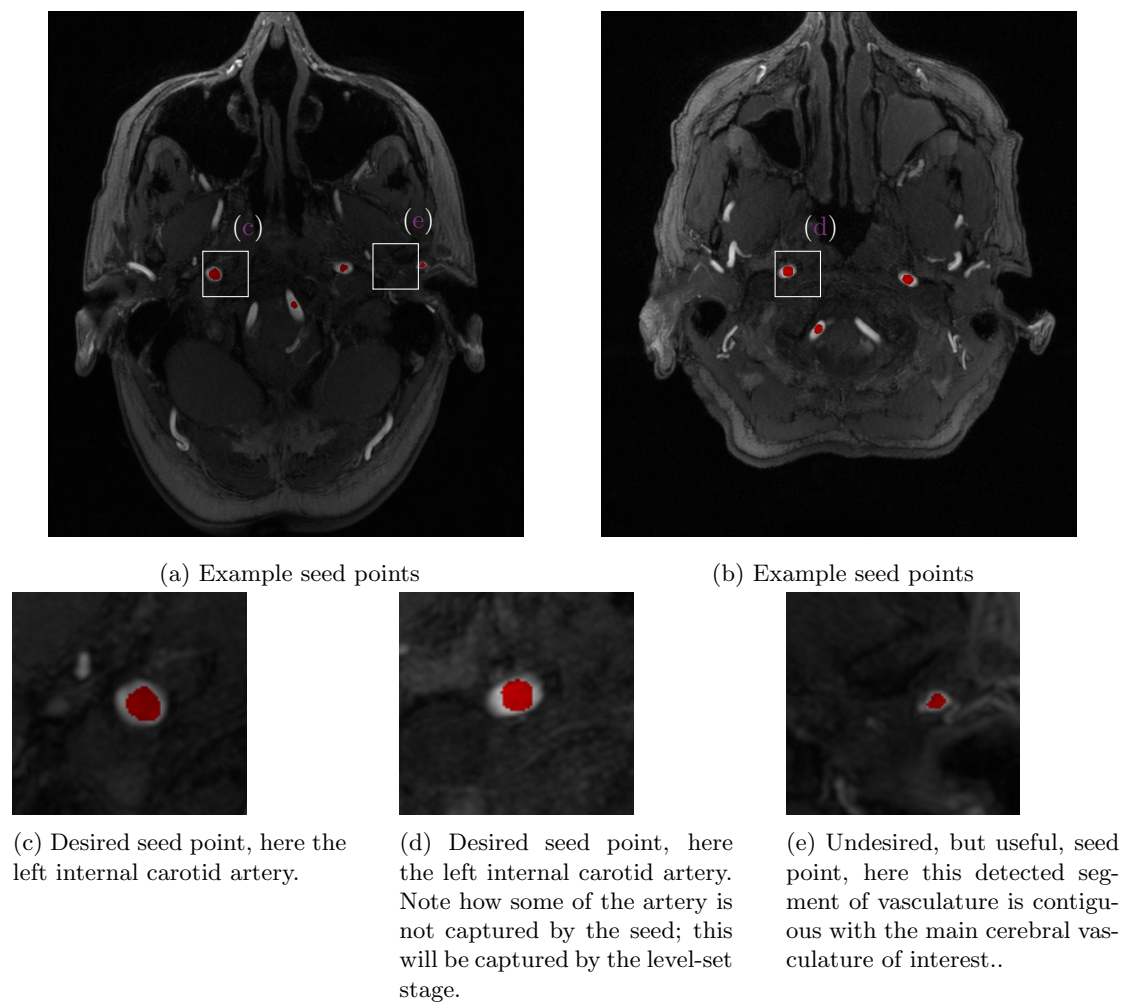


Figure 6.3: Automatically identified vascular material (red pixels in (a) and (b)) for ϕ initialisation. Note how, in some images, non-central vessels may also be captured, *e.g.* (e); such artefacts generally have little effect on the output.

```
Input :  $I, s_{\max}$   
Output:  $I_{thickness}$   
1  $r \leftarrow \text{getResolution}(I)$   
2 for  $s \in [1, s_{\max}]$  do /* for sequential sphere sizes */  
3    $b \leftarrow \text{anisotropicSphere}(s, r)$   
4    $I_s \leftarrow I \circ b$   
5 end  
6  $I_{thickness} \leftarrow \max_s(I_s)$ 
```

Algorithm 6.1: Measuring segmented vessel thickness with sequential openings allows us to get a robust measure of the thickness at all positions. Those areas with larger thicknesses than their surrounding vasculature may indicate the presence of aneurysms. In our pipeline, I is the output of the level-set segmentation process, ϕ .

6.2.3 Visualisation

This thickness information is then mapped onto the surface of the segmented vasculature. We find that mapping the thickness of the vessels to a sequential colormap, such as the commonly used jet colormap, highlights abnormalities which generally appear thicker than their adjacent vessel segments.


6.2.4 Implementation

The proposed approach is implemented using a combination of C++ and MATLAB 2016a [67]. Segmentation codes are implemented in C++ and are built on the GLFW OpenGL library for GPU parallelisation; segmentation codes can be found at [174]. Wrapper codes and all further analyses are implemented in MATLAB.


6.3 Results

We report results on ten 3D TOF 5T MRA data sets. All image slices are 580×640 pixels in size with an xy-resolution of 0.31 mm and a z-resolution of 0.57 mm. The number of slices per image is 136. The data contains four aneurysms at differing locations in the vasculature.

Of the ten sample images, seven gave visually high-quality segmentations purely through the automated initialisation figure 6.4. One segmentation had small but noticeable errors, *i.e.* missing the posterior communicating arteries (see figure 6.1) in the Circle of Willis figure 6.4a



(a) Poor automatic segmentation where the user can visually see parts of the posterior communicating arteries of the Circle of Willis are missing (white arrows).



(b) Good automatic segmentation where the user cannot identify any missing vasculature, any loss of connectivity or any extraneous segmentation.

Figure 6.4: Semi-automated segmentation of brain vasculature using our IGAC. Automatic segmentation is started from automatically detected seeds (figure 6.3). The results may be accurate representations of the vasculature (*c.f.* the maximum Z-projection (left) with the resultant segmentation (right) in (b)) or the segmentation may be missing components (a) at which point the user may take control and interactively adjust the segmentation to get a better results. The choice of 'good' and 'poor' is a visual assessment that may vary between users and between datasets.

(white arrows); however, these structures were difficult to perceive in the raw image and thus difficult to fix in the interactive stage.


Two segmentations failed to capture large sections of the network automatically, *i.e.* missing the majority of a hemisphere's vasculature. In both cases this was due to significant obstructions (hole/block) in the image figure 6.5; due to these obstructions, it was impossible to correct the segmentation during the interactive adjustment phase fully. For one image, this prevented the vasculature containing the aneurysm from being segmented; this is a clear limitation of this approach that must be overcome in moving forward.

As can be seen in figure 6.6 (*p.* 122), the calculated thickness is as expected with those vessels entering the brain being measured as thicker than the smaller vessels that branch throughout the brain. To validate this thickness we selected five random regions of each MRA image and measured the thickness of the vasculature in the raw image, comparing them to the automatically calculated thickness. We manually measured the thickness orthogonal to the vessel and calculated the mean thickness after accounting for resolution anisotropy. We found the mean error in thickness to be approximately a 0.66 mm under estimate of the thickness. This is likely due to the segmented vasculature being thinner than the human-measured vessels in the real image due to blurred and unsharp edges in the raw image.


6.4 Discussion & Conclusions

In this paper, we have introduced a new set of tools for the semi-automated segmentation of brain, and other, vascular systems and the automatic measurement of thickness across the network. This allows users to easily visualise possible sites of abnormalities, such as aneurysms, for further probing.

This approach should be easily extendable to other vascular systems such as the abdominal aorta; however, refinements are required to overcome the issue of obstructions in the segmentation stage (figure 6.5) and to better segment smaller vessels at the tips of the network. Some of this may be accomplished by the application of a 3D vessel enhancement method, such as an extension of chapter 3 (*p.* 20).



(a) Poor automatic segmentation due to a 'hole'



(b) Poor automatic segmentation due to an obstruction

Figure 6.5: Severe artefacts in the image (white arrows; right) can prevent accurate segmentation (left) leading to missing or incorrect analysis later in the pipeline. Some of these issues can be accounted for in the interactive adjustment stages but severe obstructions such as those in (a) and (b) may prevent accurate segmentation.







Figure 6.6: Results of thickness calculation on a full segmentation of brain vasculature shown in 3D (a) and overlaid on the 2D maximum Z-projections (b).



(a) Thickness overlaid on the 3D segmented surface highlighting an aneurysm



(b) Thickness projected onto a maximum Z-projection of the original MRA

Figure 6.7: Results of thickness calculation on a vasculature network with an aneurysm clearly highlighted as a thicker area relative to the neighbourhood. Here the thickness is shown in 3D (a) and overlaid on the 2D maximum Z-projections (b) and the white arrow indicates the aneurysm.

Epilogue

In this chapter I have applied mathematical morphology to the challenging task of measuring the thickness of vascular networks in MRA images. I have shown this approach to give accurate values for the thickness of vessels and allows the visualisation of possible vascular abnormalities such as aneurysms.

Over the preceding chapters, I have described several approaches that use mathematical morphology and related techniques to enhance images, detect objects and quantitatively measure objects of interest. Across this thesis, I have demonstrated the power of these approaches and the importance of remembering more classical, mathematical approaches in a time when ‘deep-learning’ is being pushed forward as *the* solution to all problems, which it may not be.

Chapter 7

Concluding Remarks

In this thesis, I have introduced new techniques for the enhancement of images and the detection and measurement of objects in images. These methods have been built upon the ideas and concepts of mathematical morphology.

Relying on the nature of the shape of vessel-like structures in images, I have introduced a new vessel enhancement method, the bowler-hat transform (chapter 3 (p. 20)), validating the results on synthetic data, comparing to existing methods and demonstrating the enhancement ability of this approach on challenging real data. I have shown that the bowler-hat transform is a successful enhancement technique that equals or betters comparator methods in synthetic and real scenarios.

Building, similarly, upon the nature of symmetry in elliptical objects in images, I have introduced a new ellipse detection algorithm, HEDAR (chapter 5 (p. 79)). Again, this approach has been validated on synthetic images, compared against state-of-the-art approaches and run on challenging real-world data. Further, I have shown the success of HEDAR for the challenging task of detecting and counting nuclei in fluorescence microscopy images (section 5.5 (p. 103)).

Finally, I have developed an early pipeline for segmenting and measuring the thickness of vasculature and other vessel-like networks based. Here, I have applied this to brain blood vessels and the visual highlighting of abnormalities such as aneurysms (chapter 6 (p. 112)).

7.1 Contributions to the Field

This thesis constitutes a series of contributions to the research area of image processing and analysis, especially within the field of mathematical morphology.

First, in chapter 3 (*p. 20*), I have introduced a new 2D technique for vessel enhancement in biomedical images, the bowler-hat transform, this technique is simple and effective and we have already begun extending the concept to 3D. The bowler-hat transform has been validated against comparator methods and is shown to give better results under both synthetic and real scenarios. Further, the bowler-hat transform has been applied to the challenging medical problem of retinal vasculature segmentation.

In chapter 4 (*p. 48*), I have extended a sub-pixel particle detection method, using the local image intensity information for post-identification pruning and selection. This approach has been explained and published in such a way as to widen the understanding of this approach and its limitations with regard to colocalisation assays in biology, making this method easily usable by the biological community.

I have then proposed, in chapter 5 (*p. 79*), a new method for detection of ‘large’ particles but ‘small’ ellipses based on directional granulometry and mathematical morphology, HEDAR. These size objects are often on the scale of noise and, as such, this is a challenging situation that HEDAR moves a step towards solving. I have moved towards resolving the problem of noise on these directional signals based on a novel use of the Hilbert Transform to detect true steps in these 1D signals. I have validated this approach, comparing to state-of-the-art methods, on both synthetic and real data. Further, HEDAR has been applied to the challenging biological problem of nuclei counting and detection.

Finally, in chapter 6 (*p. 112*), I have introduced a new tool for the segmentation and thickness measurement of brain vasculature. This method can be used as an early, global image measurement step, allowing easy visual and quantitative identification of possible aneurysms or abnormalities. I have shown the results on a small dataset of images, highlighting the failures and successes of the approach.

7.2 Conclusions

Although now a long established field, mathematical morphology still provides a wide scope for research, adaptation and utilisation for image processing, analysis and annotation. The ability

to build complex filters from sequential, simpler operations also allows such methods to be easily implemented, optimised and parallelised.

In this thesis, I have demonstrated a selection of ways that mathematical morphology can be used to enhance and quantify images. For example, the bowler-hat transform (chapter 3 (p. 20)), is a strong vessel enhancer for images, it is fast and straightforward and provides results equal to or better than other vessel enhancement techniques.

Whilst, convolutional approaches can provide excellent results, in the area of elliptical object detection, these convolutions become complex and the parameter space vast. With HEDAR I have demonstrated a reduced-parameter approach to detecting elliptical objects in images by building on signal processing of a collection of 1D signals, extracted through an approach analogous to granulometry. Whilst computationally complex, HEDAR is a powerful and successful algorithm for the detection and counting of ellipse-like objects in images (chapter 5 (p. 79)). It has even been demonstrated that HEDAR can deal with high levels of noise and clustered or overlapping objects (section 5.5 (p. 103)).

These new techniques have been shown to work well on noisy and complex data, such as medical and biological images and, despite the current trend towards deep learning, mathematical approaches like these will likely have a major role to play in biomedical image analysis for years to come.

7.3 Future Work

The majority of work in this thesis has been developed for 2D images, a key development for the uptake of these methods in the field is the extension to 3D. This extension is already ongoing for the bowler-hat transform and HEDAR. Early pilot experiments show both approaches to be as successful in 3D as in 2D.

Further, due to the computational complexity of HEDAR, some attempts to parallelise the algorithm may be required. Porting this approach to GPU architecture should be achievable with only a little rethinking of the algorithm. Further optimisation may also be possible by combining the 1D signals before analysis with the Hilbert transform, as the required Fourier transform step and its inverse are a computational bottleneck in the algorithm.

Bibliography

1. Sazak, Ç., Nelson, C. J. & Obara, B. The Multiscale Bowler-Hat Transform for Vessel Enhancement in Biomedical Images. *IEEE Transactions in Image Processing*. Submitted (2017) (pp. iii, 21).
2. Nelson, C. J., Duckney, P., Hawkins, T. J., Deeks, M. J., Laissue, P. P., Hussey, P. J. & Obara, B. Blobs and curves: object-based colocalisation for plant cells. *Functional Plant Biology* **42**, 471–485 (2015) (pp. iii, 48).
3. Nelson, C. J., Dixon, M., Laissue, P. P. & Obara, B. *Speeding up active mesh segmentation by local termination of nodes* in *Medical Image Understanding and Analysis* (July 2014) (p. iv).
4. Jackson, P. T. G., Nelson, C. J., Schiefele, J. & Obara, B. *Runway detection in High Resolution remote sensing data* in *Image and Signal Processing and Analysis (ISPA), 2015 9th International Symposium on* (Sept. 2015), 170–175 (p. iv).
5. Willcocks, C., Jackson, P., Nelson, C. & Obara, B. Extracting 3D Parametric Curves from 2D Images of Helical Objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PP**, 1–1 (2016) (p. iv).
6. Willcocks, C. G., Jackson, P. T. G., Nelson, C. J., Nasrulloh, A. V. & Obara, B. Interactive GPU Active Contours for Segmenting Inhomogenous Objects. *Journal of Real-Time Image Processing*. Submitted (2017) (pp. iv, 116, 152–153).
7. Griffin, G., Holub, A. & Perona, P. *Caltech-256 object category dataset* <http://www.vision.caltech.edu/Image%5C_Datasets/Caltech256/> (p. v).
8. Hooke, R. *Micrographia: or some physiological descriptions of minute bodies made by magnifying glasses, with observations and inquiries thereupon* (Royal Society, 1665) (p. 2).

-
9. Tomer, R., Khairy, K., Amat, F. & Keller, P. J. Quantitative high-speed imaging of entire developing embryos with simultaneous multiview light-sheet microscopy. *Nature Methods* **9**, 755–763 (2012) (p. 2).
 10. Myers, G. Why bioimage informatics matters. *Nature Methods* **9**, 659–660 (2012) (p. 2).
 11. Duncan, J. S. & Ayache, N. Medical image analysis: progress over two decades and the challenges ahead. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**, 85–106 (2000) (p. 2).
 12. Setiawan, A. W., Mengko, T. R., Santoso, O. S. & Suksmono, A. B. *Color retinal image enhancement using CLAHE* in *International Conference on ICT for Smart Society* (2013), 1–3. doi:10.1109/ICTSS.2013.6588092 (p. 2).
 13. Jerman, T., Pernuš, F., Likar, B. & Špiclin, Ž. Enhancement of Vascular Structures in 3D and 2D Angiographic Images. *IEEE Transactions on Medical Imaging* **35**, 2107–2118 (2016) (pp. 3, 24, 26, 34).
 14. Manders, E., Stap, J., Brakenhoff, G., Van Driel, R. & Aten, J. Dynamics of three-dimensional replication patterns during the S-phase, analysed by double labelling of DNA and confocal microscopy. *Journal of Cell Science* **103**, 857–862 (1992) (pp. 3, 49).
 15. Lachmanovich, E., Shvartsman, D., Malka, Y., Botvin, C., Henis, Y. & Weiss, A. Co-localization analysis of complex formation among membrane proteins by computerized fluorescence microscopy: application to immunofluorescence co-patching studies. *Journal of Microscopy* **212**, 122–131 (2003) (pp. 3, 52).
 16. Kuhn, H. W. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* **2**, 83–97 (1955) (pp. 3, 60–61, 64, 149).
 17. Irshad, H., Veillard, A., Roux, L. & Racoceanu, D. Methods for Nuclei Detection, Segmentation, and Classification in Digital Histopathology: A Review — Current Status and Future Potential. *IEEE Reviews in Biomedical Engineering* **7**, 97–114 (2014) (p. 6).
 18. Wienert, S., Heim, D., Saeger, K., Stenzinger, A., Beil, M., Hufnagl, P., Dietel, M., Denkert, C. & Klauschen, F. Detection and Segmentation of Cell Nuclei in Virtual Microscopy Images: A Minimum-Model Approach. *Scientific Reports* **2**, 503 (2012) (p. 6).
 19. Xie, Y. & Ji, Q. *A new efficient ellipse detection method* in *Proceedings of the 16th International Conference on Pattern Recognition* **2** (2002), 957–960 (pp. 6, 81, 93–94).
 20. Xu, H., Lu, C., Berendt, R., Jha, N. & Mandal, M. Automatic Nuclei Detection based on Generalized Laplacian of Gaussian Filters. *IEEE Journal of Biomedical and Health Informatics* **PP**, 1–12 (2016) (pp. 6, 106).

21. Ljosa, V., Sokolnicki, K. L. & Carpenter, A. E. Annotated high-throughput microscopy image sets for validation. *Nature Methods* **9**, 637 (2012) (pp. 6, 106).
22. Jones, T. R., Carpenter, A. & Golland, P. Voronoi-based segmentation of cells on image manifolds in *International Workshop on Computer Vision for Biomedical Image Applications* (2005), 535–543 (p. 6).
23. Meijering, E., Dzyubachyk, O., Smal, I. & van Cappellen, W. A. Tracking in cell and developmental biology. *Seminars in Cell & Developmental Biology* **20**, 894–902 (2009) (p. 6).
24. Lo, R. C., Lu, B., Fokkema, M. T., Conrad, M., Patel, V. I., Fillinger, M., Matyal, R. & Schermerhorn, M. L. Relative importance of aneurysm diameter and body size for predicting abdominal aortic aneurysm rupture in men and women. *Journal of Vascular Surgery* **59**, 1209–1216. ISSN: 0741-5214 (2014) (p. 7).
25. Bae, J.-H., Kim, W.-S., Rihal, C. S. & Lerman, A. Individual Measurement and Significance of Carotid Intima, Media, and Intima–Media Thickness by B-Mode Ultrasonographic Image Processing. *Arteriosclerosis, Thrombosis, and Vascular Biology* **26**, 2380–2385 (2006) (p. 7).
26. Yamakado, K., Tanaka, N., Nakagawa, T., Kobayashi, S., Yanagawa, M. & Takeda, K. Renal Angiomyolipoma: Relationships between Tumor Size, Aneurysm Formation, and Rupture. *Radiology* **225**, 78–82 (2002) (p. 7).
27. Reyes-Aldasoro, C. C., Björndahl, M. A., Akerman, S., Ibrahim, J., Griffiths, M. K. & Tozer, G. M. Online chromatic and scale-space microvessel-tracing analysis for transmitted light optical images. *Microvascular Research* **84**, 330–339 (2012) (p. 7).
28. Barber, P., Vojnovic, B., Ameer-Beg, S., Hodgkiss, R., Tozer, G. & Wilson, J. Semi-automated software for the three-dimensional delineation of complex vascular networks. *Journal of microscopy* **211**, 54–62 (2003) (p. 7).
29. Antiga, L. & Steinman, D. *The Vascular Modeling Toolkit (VMTK)* <<http://www.vmtk.org>> (pp. 7, 114).
30. Tyrrell, J. A., di Tomaso, E., Fuja, D., Tong, R., Kozak, K., Jain, R. K. & Roysam, B. Robust 3-D Modeling of Vasculature Imagery Using Superellipsoids. *IEEE Transactions on Medical Imaging* **26**, 223–237 (Feb. 2007) (p. 7).
31. Serra, J. *Image Analysis and Mathematical Morphology* (Academic Press, 1982) (pp. 10, 16).

-
32. Sidiropoulos, N. D., Baras, J. S. & Berenstein, C. A. Optimal filtering of digital binary images corrupted by union/intersection noise. *IEEE Transactions on Image Processing* **3**, 382–403 (1994) (p. 10).
 33. Lee, J., Haralick, R. & Shapiro, L. Morphologic edge detection. *IEEE Journal on Robotics and Automation* **3**, 142–156 (1987) (p. 10).
 34. Lantuéjoul, C. *Sur le modèle de Johnson-Mehl généralisé* Internal Report 514 (Centre de Morphologie Mathématique, France, 1977) (p. 10).
 35. Matheron, G. *Random Sets and Integral Geometry* (John Wiley & Sons, 1975) (pp. 10, 19, 83).
 36. Beucher, S. & Lantuéjoul, C. *Use of watersheds in contour detection* in *International Workshop on Image Processing: Real-time Edge and Motion Detection/Estimation* (1979) (p. 10).
 37. Maragos, P. in *Handbook of Image and Video Processing* (ed Bovik, A.) 2nd. Chapter 3.3 (Elsevier Academic, 2005). ISBN: 9780080533612 (p. 16).
 38. Zana, F. & Klein, J. C. Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation. *IEEE Transactions on Image Processing* **10**, 1010–1019 (2001) (pp. 18, 29, 34).
 39. Gonzalez, R. C. & Woods, R. E. *Digital Image Processing* 3rd edition. ISBN: 013168728X (Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2006) (p. 18).
 40. Fraz, M., Remagnino, P., Hoppe, A., Uyyanonvara, B., Rudnicka, A., Owen, C. & Barman, S. Blood vessel segmentation methodologies in retinal images – A survey. *Computer Methods and Programs in Biomedicine* **108**, 407–433 (2012) (p. 21).
 41. Frangi, A. F., Niessen, W. J., Vincken, K. L. & Viergever, M. A. *Multiscale vessel enhancement filtering* in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Springer, 1998), 130–137 (pp. 21, 24–25, 34, 114, 143).
 42. Krissian, K., Malandain, G., Ayache, N., Vaillant, R. & Troussset, Y. Model-based detection of tubular structures in 3D images. *Computer Vision and Image Understanding* **80**, 130–171 (2000) (pp. 21, 24).
 43. Obara, B., Fricker, M., Gavaghan, D. & Grau, V. Contrast-independent curvilinear structure detection in biomedical images. *IEEE Transactions on Image Processing* **21**, 2572–2581 (2012) (pp. 21, 26–27, 34).

44. Pisano, E. D., Zong, S., Hemminger, B. M., DeLuca, M., Johnston, R. E., Muller, K., Braeuning, M. P. & Pizer, S. M. Contrast limited adaptive histogram equalization image processing to improve the detection of simulated spiculations in dense mammograms. *Journal of Digital Imaging* **11**, 193–200 (1998) (pp. 21, 27, 34).
45. Chen, B., Chen, Y., Shao, Z. & Luo, L. *Retinal vessel enhancement using multi-dictionary and sparse coding in IEEE International Conference on Acoustics, Speech and Signal Processing* (2016), 893–897 (pp. 21, 29).
46. Meijering, E., Jacob, M., Sarria, J.-C., Steiner, P., Hirling, H. & Unser, M. Design and validation of a tool for neurite tracing and analysis in fluorescence microscopy images. *Cytometry Part A* **58A**, 167–176 (2004) (pp. 21, 24–25).
47. Hoover, A. D., Kouznetsova, V. & Goldbaum, M. Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response. *IEEE Transactions on Medical Imaging* **19**, 203–210 (2000) (pp. 21, 38).
48. Mendonca, A. M. & Campilho, A. Segmentation of retinal blood vessels by combining the detection of centerlines and morphological reconstruction. *Transactions on Medical Imaging* **25**, 1200–1213 (2006) (pp. 23, 29).
49. Krissian, K., Ellsmere, J., Vosburgh, K., Kikinis, R. & Westin, C.-F. *Multiscale segmentation of the aorta in 3D ultrasound images in IEEE Engineering in Medicine and Biology Society* **1** (2003), 638–641 (p. 24).
50. Levitan, I. B., Kaczmarek, L. K. *et al. The neuron: cell and molecular biology* (Oxford University Press, 2002) (p. 25).
51. Capowski, J. J. *Computer techniques in neuroanatomy* (Springer US, 1989) (p. 25).
52. Morrone, M. C., Ross, J., Burr, D. C. & Owens, R. Mach bands are phase dependent. *Nature* **324**, 250–253 (1986) (p. 26).
53. Kovese, P. *Phase congruency detects corners and edges in The Australian Pattern Recognition Society Conference* (2003), 309–318 (p. 27).
54. Kovese, P. Phase congruency: A low-level image invariant. *Psychological Research* **64**, 136–148 (2000) (p. 27).
55. Kovese, P. Image features from phase congruency. *Videre: Journal of Computer Vision Research* **1**, 1–26 (1999) (p. 27).
56. Kovese, P. *Invariant measures of image features from phase information* PhD thesis (University of Western Australia, 1996) (p. 27).

-
57. Freeman, W. T. *Steerable filters and local analysis of image structure* technical report (DTIC Document, 1992) (p. 28).
 58. Freeman, W. T. & Adelson, E. H. *Steerable filters for early vision, image analysis, and wavelet decomposition* in *International Conference on Computer Vision* (1990), 406–415 (p. 28).
 59. Freeman, W. T. & Adelson, E. H. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis & Machine Intelligence* **13**, 891–906 (1991) (p. 28).
 60. Perona, P. Steerable-scalable kernels for edge detection and junction analysis. *Image and Vision Computing* **10**, 663–672 (1992) (p. 28).
 61. Jalba, A. C., Wilkinson, M. H. & Roerdink, J. B. Morphological hat-transform scale spaces and their use in pattern classification. *Pattern Recognition* **37**, 901–915 (2004) (p. 29).
 62. De, I., Chanda, B. & Chattopadhyay, B. Enhancing effective depth-of-field by image fusion using mathematical morphology. *Image and Vision Computing* **24**, 1278–1287 (2006) (p. 29).
 63. Chen, T., Wu, Q., Rahmani-Torkaman, R. & Hughes, J. A pseudo top-hat mathematical morphological approach to edge detection in dark regions. *Pattern Recognition* **35**, 199–210 (2002) (p. 29).
 64. Miri, M. S. & Mahloojifar, A. Retinal Image Analysis Using Curvelet Transform and Multistructure Elements Morphology by Reconstruction. *IEEE Transactions on Biomedical Engineering* **58**, 1183–1192 (2011) (p. 29).
 65. Fraz, M., Barman, S., Remagnino, P., Hoppe, A., Basit, A., Uyyanonvara, B., Rudnicka, A. & Owen, C. An approach to localize the retinal blood vessels using bit planes and centerline detection. *Computer Methods and Programs in Biomedicine* **108**, 600–616 (2012) (p. 29).
 66. Lu, C.-Y., Jing, B.-Z., Chan, P. P., Xiang, D., Xie, W., Wang, J. & Yeung, D. S. *Vessel enhancement of low quality fundus image using mathematical morphology and combination of Gabor and matched filter* in *International Conference on Wavelet Analysis and Pattern Recognition* (2016), 168–173 (p. 29).
 67. The Mathworks, Inc. *MATLAB* <https://www.mathworks.com>. Massachusetts, USA, 2016 (pp. 34, 89, 94, 118).
 68. Sazak, C. *The Bowler Hat Transform* GitHub. <To%20Be%20Released> (p. 34).
 69. Lindeberg, T. *Scale-space theory in computer vision* (Springer Science & Business Media, 2013) (p. 34).

70. Staal, J., Abramoff, M., Niemeijer, M., Viergever, M. & van Ginneken, B. Ridge based vessel segmentation in color images of the retina. *IEEE Transactions on Medical Imaging* **23**, 501–509 (2004) (p. 38).
71. Yang, Y., Li, Y. & Wu, C. Genomic resources for functional analyses of the rice genome. *Current Opinion in Plant Biology* **16**, 157–163 (2013) (p. 49).
72. Cobb, J. N., DeClerck, G., Greenberg, A., Clark, R. & McCouch, S. Next-generation phenotyping: requirements and strategies for enhancing our understanding of genotype–phenotype relationships and its relevance to crop improvement. *Theoretical and Applied Genetics* **126**, 867–887 (2013) (p. 49).
73. Thomasson, M. S. & Macnaughtan, M. A. Microscopy basics and the study of actin–actin-binding protein interactions. *Analytical Biochemistry* **443**, 156–165 (2013) (p. 49).
74. Sappl, P. G. & Heisler, M. G. Live-imaging of plant development: latest approaches. *Current Opinion in Plant Biology* **16**, 33–40 (2013) (p. 49).
75. Swedlow, J. R., Goldberg, I., Brauner, E. & Sorger, P. K. Informatics and quantitative analysis in biological imaging. *Science* **300**, 100–102 (2003) (p. 49).
76. Cardona, A. & Tomancak, P. Current challenges in open-source bioimage informatics. *Nature Methods* **9**, 661–665 (2012) (p. 49).
77. Obara, B., Jabeen, A., Fernandez, N. & Laissue, P. P. A novel method for quantified, super-resolved, three-dimensional colocalisation of isotropic, fluorescent particles. *Histochemistry and Cell Biology* **139**, 391–402 (2013) (pp. 50, 59).
78. Eliceiri, K. W., Berthold, M. R., Goldberg, I. G., Ibáñez, L., Manjunath, B. S., Martone, M. E., Murphy, R. F., Peng, H., Plant, A. L., Roysam, B. *et al.* Biological imaging software tools. *Nature Methods* **9**, 697–710 (2012) (p. 50).
79. Allan, C., Burel, J.-M., Moore, J., Blackburn, C., Linkert, M., Loynton, S., MacDonald, D., Moore, W. J., Neves, C., Patterson, A. *et al.* OMERO: flexible, model-driven data management for experimental biology. *Nature Methods* **9**, 245–253 (2012) (p. 50).
80. Kvilekval, K., Fedorov, D., Obara, B., Singh, A. & Manjunath, B. Bisque: a platform for bioimage analysis and management. *Bioinformatics* **26**, 544–552 (2010) (p. 50).
81. Cho, B. H., Cao-Berg, I., Bakal, J. A. & Murphy, R. F. OMERO. searcher: content-based image search for microscope images. *Nature Methods* **9**, 633–634 (2012) (p. 50).

-
82. Goff, S. A., Vaughn, M., McKay, S., Lyons, E., Stapleton, A. E., Gessler, D., Matasci, N., Wang, L., Hanlon, M., Lenards, A., Muir, A., Merchant, N., Lowry, S., Mock, S., Helmke, M., Kubach, A., Narro, M., Hopkins, N., Micklos, D., Hilgert, U., Gonzales, M., Jordan, C., Skidmore, E., Dooley, R., Cazes, J., McLay, R., Lu, Z., Pasternak, S., Koesterke, L., Piel, W. H., Grene, R., Noutsos, C., Gendler, K., Feng, X., Tang, C., Lent, M., Kim, S.-j., Kvilekval, K., Manjunath, B. S., Tannen, V., Stamatakis, A., Sanderson, M., Welch, S. M., Cranston, K., Soltis, P., Soltis, D., O'Meara, B., Ane, C., Brutnell, T., Kleibenstein, D. J., White, J. W., Leebens-Mack, J., Donoghue, M. J., Spalding, E. P., Vision, T. J., Myers, C. R., Lowenthal, D., Enquist, B. J., Boyle, B., Akoglu, A., Andrews, G., Ram, S., Ware, D., Stein, L. & Stanzone, D. The iPlant collaborative: cyberinfrastructure for plant biology. *Frontiers in Plant Science* **2**, 00034 (2011) (p. 50).
83. Lobet, G., Draye, X. & Périlleux, C. An online database for plant image analysis software tools. *Plant Methods* **9**, 1 (2013) (pp. 50–51).
84. Roeder, A. H., Chickarmane, V., Cunha, A., Obara, B., Manjunath, B. & Meyerowitz, E. M. Variability in the control of cell division underlies sepal epidermal patterning in *Arabidopsis thaliana*. *PLoS Biology* **8**, e1000367 (2010) (p. 50).
85. Lichius, A., Goryachev, A. B., Fricker, M. D., Obara, B., Castro-Longoria, E. & Read, N. D. CDC-42 and RAC-1 regulate opposite chemotropisms in *Neurospora crassa*. *Journal of Cell Science* **127**, 1953–1965 (2014) (p. 50).
86. Obara, B., Grau, V. & Fricker, M. D. A bioimage informatics approach to automatically extract complex fungal networks. *Bioinformatics* **28**, 2374–2381 (2012) (p. 50).
87. Pound, M. P., French, A. P., Wells, D. M., Bennett, M. J. & Pridmore, T. P. CellSeT: novel software to extract and analyze structured networks of plant cells from confocal images. *The Plant Cell* **24**, 1353–1361 (2012) (p. 50).
88. Rolland-Lagan, A.-G., Amin, M. & Pakulska, M. Quantifying leaf venation patterns: two-dimensional maps. *The Plant Journal* **57**, 195–205 (2009) (p. 50).
89. Dhondt, S., Van Haerenborgh, D., Van Cauwenbergh, C., Merks, R. M., Philips, W., Beemster, G. T. & Inzé, D. Quantitative analysis of venation patterns of *Arabidopsis* leaves by supervised image analysis. *The Plant Journal* **69**, 553–563 (2012) (p. 50).
90. Pridmore, T. P., French, A. P. & Pound, M. P. What lies beneath: underlying assumptions in bioimage analysis. *Trends in Plant Science* **17**, 688–692 (2012) (p. 51).
91. Carpenter, A. E., Kametsky, L. & Eliceiri, K. W. A call for bioimaging software usability. *Nature Methods* **9**, 666 (2012) (p. 51).

92. Rajaram, S., Pavie, B., Hac, N. E., Altschuler, S. J. & Wu, L. F. SimuCell: a flexible framework for creating synthetic microscopy images. *Nature Methods* **9**, 634–635 (2012) (p. 51).
93. Gelasca, E. D., Obara, B., Fedorov, D., Kvilekval, K. & Manjunath, B. A biosegmentation benchmark for evaluation of bioimage analysis methods. *BMC Bioinformatics* **10**, 368 (2009) (p. 51).
94. Brunkard, J. O., Runkel, A. M. & Zambryski, P. C. Plasmodesmata dynamics are coordinated by intracellular signaling pathways. *Current Opinion in Plant Biology* **16**, 614–620 (2013) (p. 51).
95. Deeks, M. J., Calcutt, J. R., Ingle, E. K., Hawkins, T. J., Chapman, S., Richardson, A. C., Mentlak, D. A., Dixon, M. R., Cartwright, F., Smertenko, A. P., Oparaka, K. & Hussey, P. J. A superfamily of actin-binding proteins at the actin-membrane nexus of higher plants. *Current Biology* **22**, 1595–1600 (2012) (pp. 51–52, 62–64, 68, 72, 74).
96. Simpson, C., Thomas, C., Findlay, K., Bayer, E. & Maule, A. J. An Arabidopsis GPI-anchor plasmodesmal neck protein with callose binding activity and potential to regulate cell-to-cell trafficking. *The Plant Cell* **21**, 581–594 (2009) (p. 52).
97. Kleine-Vehn, J., Wabnik, K., Martiniere, A., Langowski, L., Willig, K., Naramoto, S., Leitner, J., Tanaka, H., Jakobs, S., Robert, S. *et al.* Recycling, clustering, and endocytosis jointly maintain PIN auxin carrier polarity at the plasma membrane. *Molecular Systems Biology* **7**, 540 (2011) (p. 52).
98. Hepler, P. K., Vidali, L. & Cheung, A. Y. Polarized cell growth in higher plants. *Annual Review of Cell and Developmental Biology* **17**, 159–187 (2001) (p. 52).
99. Cheung, A. Y. & Wu, H.-M. Structural and signaling networks for the polar cell growth machinery in pollen tubes. *Annual Review of Plant Biology* **59**, 547–572 (2008) (p. 52).
100. Cai, G., Faleri, C., Del Casino, C., Emons, A. M. C. & Cresti, M. Distribution of callose synthase, cellulose synthase, and sucrose synthase in tobacco pollen tube is controlled in dissimilar ways by actin filaments and microtubules. *Plant Physiology* **155**, 1169–1190 (2011) (p. 52).
101. Moscatelli, A., Idilli, A., Rodighiero, S. & Caccianiga, M. Inhibition of actin polymerisation by low concentration Latrunculin B affects endocytosis and alters exocytosis in shank and tip of tobacco pollen tubes. *Plant Biology* **14**, 770–782 (2012) (p. 52).

-
102. Zheng, M., Wang, Q., Teng, Y., Wang, X., Wang, F., Chen, T., Šamaj, J., Lin, J. & Logan, D. C. The speed of mitochondrial movement is regulated by the cytoskeleton and myosin in *Picea wilsonii* pollen tubes. *Planta* **231**, 779–791 (2010) (p. 54).
 103. Van Gestel, K., Köhler, R. & Verbelen, J.-P. Plant mitochondria move on F-actin, but their positioning in the cortical cytoplasm depends on both F-actin and microtubules. *Journal of Experimental Botany* **53**, 659–667 (2002) (p. 54).
 104. Zheng, M., Beck, M., Müller, J., Chen, T., Wang, X., Wang, F., Wang, Q., Wang, Y., Baluška, F., Logan, D. C. *et al.* Actin turnover is required for myosin-dependent mitochondrial movements in *Arabidopsis* root hairs. *PLoS One* **4**, e5961 (2009) (pp. 54, 75).
 105. Sparkes, I. A., Teanby, N. A. & Hawes, C. Truncated myosin XI tail fusions inhibit peroxisome, Golgi, and mitochondrial movement in tobacco leaf epidermal cells: a genetic tool for the next generation. *Journal of Experimental Botany* **59**, 2499–2512 (2008) (p. 54).
 106. Avisar, D., Abu-Abied, M., Belausov, E., Sadot, E., Hawes, C. & Sparkes, I. A. A comparative study of the involvement of 17 *Arabidopsis* myosin family members on the motility of Golgi and other organelles. *Plant Physiology* **150**, 700–709 (2009) (p. 54).
 107. Romagnoli, S., Cai, G., Faleri, C., Yokota, E., Shimmen, T. & Cresti, M. Microtubule- and actin filament-dependent motors are distributed on pollen tube mitochondria and contribute differently to their movement. *Plant and Cell Physiology* **48**, 345–361 (2007) (pp. 54, 75).
 108. Murashige, T. & Skoog, F. A revised medium for rapid growth and bio assays with tobacco tissue cultures. *Physiologia Plantarum* **15**, 473–497 (1962) (p. 54).
 109. Boavida, L. C. & McCormick, S. TECHNICAL ADVANCE: Temperature as a determinant factor for increased and reproducible in vitro pollen germination in *Arabidopsis thaliana*. *The Plant Journal* **52**, 570–582 (2007) (p. 54).
 110. Nyquist, H. Certain topics in telegraph transmission theory. *Proceedings of the IEEE* **47**, 617–644 (1928) (p. 55).
 111. Shannon, C. E. Communication in the presence of noise. *Proceedings of the IRE* **37**, 10–21 (1949) (p. 55).
 112. Babaloukas, G., Tentolouris, N., Liatis, S., Sklavounou, A. & Perrea, D. Evaluation of three methods for retrospective correction of vignetting on medical microscopy images utilizing two open source software tools. *Journal of Microscopy* **244**, 320–324 (2011) (p. 56).
 113. Likar, B., Maintz, J. A., Viergever, M. A. & Pernus, F. Retrospective shading correction based on entropy minimization. *Journal of Microscopy* **197**, 285–295 (2000) (p. 56).

114. Stockman, G. & Shapiro, L. G. *Computer Vision* 1st (Prentice Hall PTR, Upper Saddle River, NJ, USA, 2001) (p. 56).
115. Sezgin, M. & Sankur, B. Survey over image thresholding techniques and quantitative performance evaluation. *Journal of Electronic Imaging* **13**, 146–168 (2004) (p. 57).
116. Obara, B., Byun, J., Fedorov, D. & Manjunath, B. *Automatic nuclei detection and data flow in Bisquik system* in *Workshop on Bio-Image Informatics: Biological Imaging, Computer Vision and Data Mining, Santa Barbara, CA, USA* (2008), 17–18 (p. 59).
117. Schneider, C. A., Rasband, W. S. & Eliceiri, K. W. NIH Image to ImageJ: 25 years of image analysis. *Nature Methods* **9**, 671–675. ISSN: 1548-7091 (July 2012) (p. 63).
118. Wong, C. Y., Lin, S. C. F., Ren, T. R. & Kwok, N. M. *A survey on ellipse detection methods* in *IEEE International Symposium on Industrial Electronics* (2012), 1105–1110 (pp. 80, 82).
119. Mukhopadhyay, P. & Chaudhuri, B. B. A survey of Hough Transform. *Pattern Recognition* **48**, 993–1010 (2015) (pp. 80, 82).
120. Rosin, P. L. A note on the least squares fitting of ellipses. *Pattern Recognition Letters* **14**, 799–808 (1993) (p. 81).
121. Gander, W., Golub, G. H. & Strebler, R. Least-squares fitting of circles and ellipses. *BIT Numerical Mathematics* **34**, 558–578 (1994) (p. 81).
122. Fitzgibbon, A. W., Fisher, R. B., Hill, F. & Eh, E. Direct Least Squares Fitting of Ellipses. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **21**, 1–15 (1996) (p. 81).
123. Maini, E. S. Enhanced Direct Least Square Fitting of Ellipses. *International Journal of Pattern Recognition and Artificial Intelligence* **20**, 939–953 (2006) (p. 81).
124. Kanatani, K. & Rangarajan, P. Hyper least squares fitting of circles and ellipses. *Computational Statistics and Data Analysis* **55**, 2197–2208 (2011) (p. 81).
125. Hal, R. & Flusser, J. *Numerically stable direct least squares fitting of ellipses* in *International Conference in Central Europe on Computer Graphics and Visualization* **6** (1998), 125–132 (p. 81).
126. Prasad, D. K., Leung, M. K. H. & Quek, C. ElliFit: An unconstrained, non-iterative, least squares based geometric Ellipse Fitting method. *Pattern Recognition* **46**, 1449–1465 (2013) (pp. 81, 93–94).
127. Hough, P. V. C. *Method and means for recognizing complex patterns* US Patent 3069654. 1962 (p. 81).

-
128. Duda, R. O. & Hart, P. E. Use of the Hough transform to detect lines and curves in pictures. *Communications of the Association Computing Machinery* **15**, 11–15 (1972) (p. 81).
 129. Ballard, D. H. Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition* **13**, 111–122 (1981) (p. 81).
 130. Nair, P. S. & Saunders, A. T. Hough transform based ellipse detection algorithm. *Pattern Recognition Letters* **17**, 777–784 (1996) (p. 81).
 131. Guil, N. & Zapata, E. Lower order circle and ellipse Hough transform. *Pattern Recognition* **30**, 1729–1744 (1997) (p. 81).
 132. McLaughlin, R. A. Randomized Hough Transform: Improved ellipse detection with comparison. *Pattern Recognition Letters* **19**, 299–305 (1998) (p. 81).
 133. Basca, C. A., Talos, M. & Brad, R. *Randomized Hough Transform for Ellipse Detection with Result Clustering* in *EUROCON. The International Conference on “Computer as a tool”* **2** (2005), 1397–1400 (pp. 81, 93–94).
 134. Kim, E., Haseyama, M. & Kitajima, H. *Fast and Robust Ellipse Extraction from Complicated Images* in *International Conference on Information Technology and Applications* **1** (2002), 357–362 (p. 82).
 135. Qiao, Y. & Ong, S. H. Arc-based evaluation and detection of ellipses. *Pattern Recognition* **40**, 1990–2003 (2007) (p. 82).
 136. Mai, F., Hung, Y. S., Zhong, H. & Sze, W. F. A hierarchical approach for fast and robust ellipse extraction. *Pattern Recognition* **41**, 2512–2524 (2008) (p. 82).
 137. Fornaciari, M., Prati, A. & Cucchiara, R. A fast and effective ellipse detector for embedded vision applications. *Pattern Recognition* **47**, 3693–3708 (2014) (pp. 82, 93–95).
 138. Chia, A. Y., Rahardja, S., Rajan, D. & Leung, K. A Split and Merge Based Ellipse Detector With Self-Correcting Capability. *IEEE Transactions on Image Processing* **20**, 1991–2006 (2011) (p. 82).
 139. Bolles, R. C. & Fischler, M. a. *A RANSAC-based approach to model fitting and its application to finding cylinders in range data* in *International Joint Conference on Artificial Intelligence* (1981), 637–643 (p. 82).
 140. Song, G. & Wang, H. in *Computer Analysis of Images and Patterns* (eds Kropatsch, W. G., Kampel, M. & Hanbury, A.) 669–676 (2007) (p. 82).
 141. Cicconet, M., Gunsalus, K., Geiger, D. & Werman, M. *Ellipses from triangles* in *International Conference on Image Processing* (2014), 3626–3630 (p. 82).

142. Liu, H. & Wang, Z. Geometric property based ellipse detection method. *Journal of Visual Communication and Image Representation* **24**, 1075–1086 (2013) (p. 82).
143. Vanrell, M. & Vitria, J. M. *Mathematical morphology, granulometries, and texture perception* in *SPIE Proceedings* **2030** (1993), 152–161 (p. 83).
144. Vincent, L. *Fast grayscale granulometry algorithms* in *Mathematical Morphology and its Applications to Image Processing* **2** (1994), 265–272 (pp. 83, 89).
145. Nelson, C. J. *Ellipse Detection by Hilber-Edge Detection and Ranging (HEDAR)* GitHub. <<https://github.com/ChasNelson1990/Ellipse-Detection-by-Hilbert-Edge-Detection-and-Ranging>> (p. 89).
146. Pei, S. C., Ding, J. J., Huang, J. D. & Guo, C. *Short response Hilbert transform for edge detection* in *IEEE Asia Pacific Conference on Circuits and Systems* (2008), 340–343 (p. 91).
147. Kanatani, K., Al-Sharadqah, A., Chernov, N. & Sugaya, Y. Hyper-renormalization: Non-minimization approach for geometric estimation. *IPSN Transactions on Computer Vision and Applications* **6**, 143–159 (2014) (pp. 93–94).
148. Jaccard, P. The Distribution of Flora in the Alpine Zone. *New Phytologist* **11**, 37–50 (1912) (p. 95).
149. Jarrin, M., Young, L., Wu, W., Girkin, J. M. & Quinlan, R. A. in *Methods in Enzymology* (eds Omary, M. B. & Liem, R. K. H.) 581–611 (2015) (p. 104).
150. Gurcan, M. N., Pan, T., Shimada, H. & Saltz, J. *Image Analysis for Neuroblastoma Classification: Segmentation of Cell Nuclei* in *IEEE Engineering in Medicine and Biology Society* (2006), 4844–4847 (p. 103).
151. Jung, C. & Kim, C. Segmenting Clustered Nuclei Using H-minima Transform-Based Marker Extraction and Contour Parameterization. *IEEE Transactions on Biomedical Engineering* **57**, 2600–2604 (2010) (p. 106).
152. Xu, H., Lu, C. & Mandal, M. An Efficient Technique for Nuclei Segmentation Based on Ellipse Descriptor Analysis and Improved Seed Detection Algorithm. *IEEE Journal of Biomedical and Health Informatics* **18**, 1729–1741 (2014) (p. 106).
153. Kong, J., Wang, F., Teodoro, G., Liang, Y., Zhu, Y., Tucker-Burden, C. & Brat, D. J. *Automated cell segmentation with 3D fluorescence microscopy images* in *IEEE International Symposium on Biomedical Imaging* (2015), 1212–1215 (p. 106).

-
154. Li, G., Liu, T., Tarokh, A., Nie, J., Guo, L., Mara, A., Holley, S. & Wong, S. T. C. 3D cell nuclei segmentation based on gradient flow tracking. *BMC Cell Biology* **8**, 1–10 (2007) (p. 106).
 155. Nandy, K., Chellappa, R., Kumar, A. & Lockett, S. J. Segmentation of Nuclei From 3D Microscopy Images of Tissue via Graphcut Optimization. *IEEE Journal of Selected Topics in Signal Processing* **10**, 140–150 (2016) (p. 106).
 156. Nelson, C. J., Jackson, P. T. G. & Obara, B. Ellipse Detection by Hilbert-Edge Detection and Ranging (HEDAR). *PLoS one*. Submitted (2017) (p. 106).
 157. Carpenter, A. E., Jones, T. R., Lamprecht, M. R., Clarke, C., Kang, I. H., Friman, O., Guertin, D. A., Chang, J. H., Lindquist, R. A., Moffat, J., Golland, P. & Sabatini, D. M. CellProfiler: image analysis software for identifying and quantifying cell phenotypes. *Genome Biology* **7**, 1–11 (2006) (pp. 106, 108).
 158. Ruusuvuori, P., Lehmussola, A., Selinummi, J., Rajala, T., Huttunen, H. & Yli-Harja, O. Benchmark set of synthetic images for validating cell image analysis algorithms in *European Signal Processing Conference* (2008), 1–5 (p. 107).
 159. Lehmussola, A., Ruusuvuori, P., Selinummi, J., Huttunen, H. & Yli-Harja, O. Synthetic Images of High-Throughput Microscopy for Validation of Image Analysis Methods. *Proceedings of the IEEE* **96**, 1348–1360 (2008) (p. 107).
 160. Lehmussola, A., Ruusuvuori, P., Selinummi, J., Huttunen, H. & Yli-Harja, O. Computational Framework for Simulating Fluorescence Microscope Images With Cell Populations. *IEEE Transactions on Medical Imaging* **26**, 1010–1016 (2007) (p. 107).
 161. Wirjadi, O., Kim, Y.-J., Stech, F., Bonfert, L. & Wagner, M. Bayesian model for detection and classification of meningioma nuclei in microscopic images. *Journal of Microscopy*, n/a–n/a (2016) (pp. 109–110).
 162. Khan, H. A. & Maruf, G. M. Counting clustered cells using distance mapping in *International Conference on Informatics, Electronics & Vision* (2013), 1–6 (pp. 109–110).
 163. Vlak, M. H., Algra, A., Brandenburg, R. & Rinkel, G. J. Prevalence of unruptured intracranial aneurysms, with emphasis on sex, age, comorbidity, country, and time period: a systematic review and meta-analysis. *The Lancet Neurology* **10**, 626–636 (2011) (p. 113).
 164. Lauric, A., Miller, E. L., Baharoglu, M. I. & Malek, A. M. Rupture Status Discrimination in Intracranial Aneurysms Using the Centroid-Radii Model. *IEEE Transactions on Biomedical Engineering* **58**, 2895–2903 (Oct. 2011) (p. 113).

165. Raghavan, M. L., Ma, B. & Harbaugh, R. E. Quantified aneurysm shape and rupture risk. *Journal of Neurosurgery* **102**, 355–362 (2005) (p. 113).
166. Doi, K. Current status and future potential of computer-aided diagnosis in medical imaging. *The British Journal of Radiology* (2014) (p. 114).
167. Piccinelli, M., Veneziani, A., Steinman, D. A., Remuzzi, A. & Antiga, L. A Framework for Geometric Analysis of Vascular Structures: Application to Cerebral Aneurysms. *IEEE Transactions on Medical Imaging* **28**, 1141–1155 (Aug. 2009) (p. 114).
168. Yang, X., Blezek, D. J., Cheng, L. T. E., Ryan, W. J., Kallmes, D. F. & Erickson, B. J. Computer-Aided Detection of Intracranial Aneurysms in MR Angiography. *Journal of Digital Imaging* **24**, 86–95 (2011) (p. 114).
169. Lauric, A., Miller, E., Frisken, S. & Malek, A. M. Automated detection of intracranial aneurysms based on parent vessel 3D analysis. *Medical Image Analysis* **14**, 149–159 (2010) (p. 114).
170. Kobashi, S., Kondo, K. & Hata, Y. Computer-Aided Diagnosis of Intracranial Aneurysms in MRA Images with Case-Based Reasoning. *IEICE Transactions on Information and Systems* **E89-D**, 340–350 (Jan. 2006) (p. 114).
171. Hentschke, C. M., Beuing, O., Nickl, R. & Tönnies, K. D. *Automatic cerebral aneurysm detection in multimodal angiographic images* in *IEEE Nuclear Science Symposium Conference Record* (Oct. 2011), 3116–3120 (p. 114).
172. Hentschke, C. M., Tönnies, K. D., Beuing, O. & Nickl, R. *A new feature for automatic aneurysm detection* in *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)* (June 2012), 800–803 (p. 114).
173. Wang, L., He, L., Mishra, A. & Li, C. Active Contours Driven by Local Gaussian Distribution Fitting Energy. *Signal Processing* **89**, 2435–2447 (Dec. 2009) (pp. 116, 152–153).
174. Willcocks, C. G. *Interactive GPU Active Contours* GitHub. <To%20Be%20Released> (p. 118).
175. Obara, B., Fricker, M. & Grau, V. *Coherence enhancing diffusion filtering based on the Phase Congruency Tensor* in *IEEE International Symposium on Biomedical Imaging* (2012), 202–205 (p. 147).
176. Chan, T. F. & Vese, L. A. Active contours without edges. *IEEE Transactions on image processing* **10**, 266–277 (2001) (p. 153).

Appendix A

The Hessian Matrix for Image Processing

A.1 The Scaled 2D Hessian Matrix

With respect to its use in image processing, the Hessian matrix, named after the German mathematician Ludwig Otto Hesse, is a square matrix calculated at each pixel whose elements are formed of the second-order partial derivatives in x and y for any 2D image $I(\vec{p})$, where $\vec{p} = (x, y)$,

$$H = \begin{bmatrix} H_{11} & H_{12} \\ H_{12} & H_{22} \end{bmatrix} = \begin{bmatrix} \frac{\partial^2 I}{\partial x^2} & \frac{\partial^2 I}{\partial x \partial y} \\ \frac{\partial^2 I}{\partial x \partial y} & \frac{\partial^2 I}{\partial y^2} \end{bmatrix}. \quad (\text{A.1})$$

Where all second-order partial derivatives of $I(\vec{p})$ exist and are continuous in the image domain.

In [41], Frangi and colleagues extend the Hessian to a multi-scale filter by convolving with Gaussians of a known size or scale, s ,

$$H_s = \begin{bmatrix} \frac{\partial^2 I}{\partial x^2} * G_s & \frac{\partial^2 I}{\partial x \partial y} * G_s \\ \frac{\partial^2 I}{\partial x \partial y} * G_s & \frac{\partial^2 I}{\partial y^2} * G_s \end{bmatrix}. \quad (\text{A.2})$$

This scaled Hessian describes the local intensity variations around pixel \vec{p} at scale s .

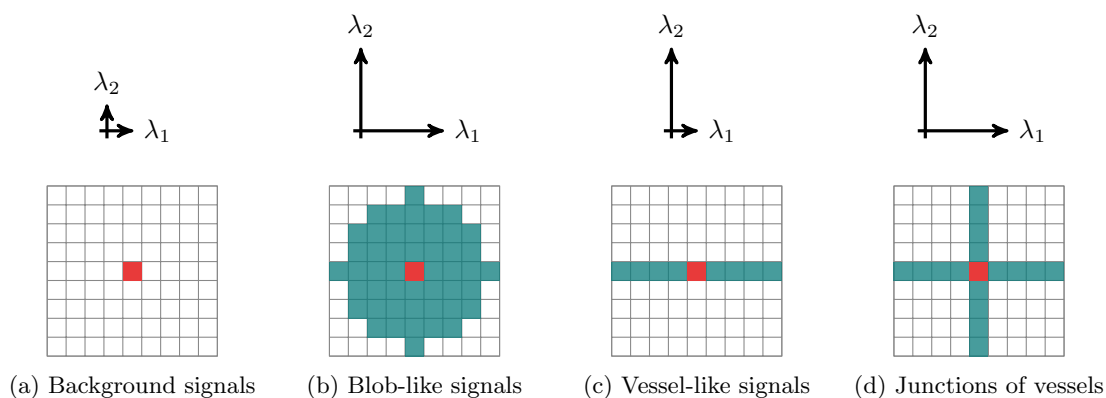


Figure A.1: Visualising the Hessian matrix eigenvalues for blobs and vessels in 2D around a specific pixel (red).

A.2 Eigenvalues of the 2D Hessian Matrix

The sorted eigenvalues of the Hessian matrix, *i.e.* $\|\lambda_1\| \leq \|\lambda_2\|$, can be used in filtering for blob-like and vessel-like features in images as shown in chapter 3 (*p.* 20). In figure A.1, these eigenvalues have been visualised for background signals, blob-like signals, vessel-like signals and vessel junctions around a known pixel (red) in 2D space.

Due to the blob suppressing nature of the first Hessian-based vessel enhancement approaches they often fail at junctions (*c.f.* the eigenvalue ratios of figure A.1b and figure A.1d). This fundamental problem is overcome by considering only the λ_2 eigenvalues; however, this leads to the ‘enhancement’ of vessels in areas of background and blob-like signals. The work in chapter 3 (*p.* 20) goes some way to combatting this issue.

Appendix B

The Phase Congruency Tensor for Image Processing

B.1 Phase Congruency

In image processing, phase congruency is an intensity- and contrast-independent feature metric that can be used for edge, corner and other feature detection. Phase congruency relies on the feature phases in the frequency domain of the image (figure B.1). At fetures such as edges and corners many frequency components are in phase regardless of the contrast between foreground and background in the intensity domain.

Mathematically, we define the phase congruency, $PC_s(\vec{p})$, for a set scale, s , over several orientations, θ and at a specific pixel, $\vec{p} = (x, y)$, as,

$$PC_s(\vec{p}) = \sum_{\theta} PC_{s,\theta}(\vec{p}), \quad (\text{B.1})$$

where $PC_{s,\theta}(\vec{p})$, the phase congruency for each orientation, is defined as,

$$PC_{s,\theta}(\vec{p}) = \frac{\sum_s W_s(\vec{p}) [A_{s,\theta}(\vec{p}) \Delta \varphi_{s,\theta}(\vec{p}) - t]}{\sum_s A_{s,\theta}(\vec{p}) + \epsilon}. \quad (\text{B.2})$$

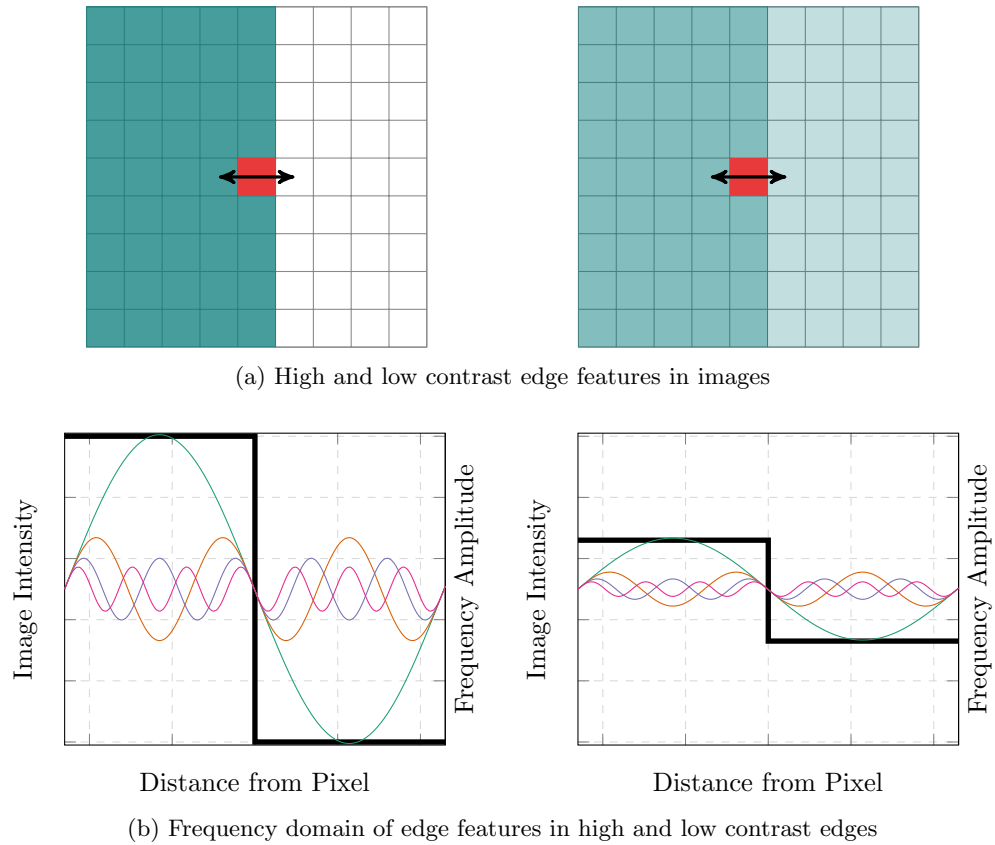


Figure B.1: The frequency domain phase features of high (left) and low (right) contrast edges. Note how the frequency component phases line up at the edge regardless of intensity

Here $A_{s,\theta}(\vec{p})$ is the amplitude of the image component at scale s , t is a noise reducing threshold and ϵ a small factor to prevent division by zero. Here $\lfloor \cdot \rfloor$ is defined as,

$$\lfloor f(x) \rfloor = \begin{cases} f(x) & f(x) \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (\text{B.3})$$

The phase deviation, $\Delta\varphi_{s,\theta}(\vec{p})$, quantifies the difference between the phase at scale s , $\varphi_{s,\theta}(\vec{p})$, and the mean phase over all scales, $\bar{\varphi}_\theta(\vec{p})$, such that,

$$\Delta\varphi_{s,\theta}(\vec{p}) = |\cos(\varphi_{s,\theta}(\vec{p}) - \bar{\varphi}_\theta(\vec{p})) - \sin(\varphi_{s,\theta}(\vec{p}) - \bar{\varphi}_\theta(\vec{p}))|. \quad (\text{B.4})$$

Finally, $W(\vec{p})$ is a sigmoid weighting term penalising narrow frequency regions based on γ and c , gain and cut-off, respectively, $A_{\max}(\vec{p})$, which is the maximum response over all scales and S , the total number of scales. Mathematically,

$$W_s(\vec{p}) = \left(1 + e^{\gamma \left(c - \frac{1}{S} \left(\frac{\sum_s A_s(\vec{p})}{A_{\max}(\vec{p}) + \epsilon} \right) \right)} \right). \quad (\text{B.5})$$

B.2 Phase Congruency Tensor

The contrast-independent nature of phase congruency makes it a powerful tool for many image processing tasks. In [175] the phase congruency metric was adapted to produce a local, tensor-based representation of the phase congruency in the image. The tensor can thus be exploited in a similar manner to the Hessian matrix. Here we assume that the phase congruency is high in the direction perpendicular to a salient feature and close to zero along the direction of that feature.

For a given set of scales, s , and orientations, θ , the phase congruency tensor at point \vec{p} is defined as,

$$T_{PC} = \sum_{\theta} PC_{\theta}(\vec{p}) \vec{n}_{\theta} \vec{n}_{\theta}^T, \quad (\text{B.6})$$

where \vec{n} is the normal vector along angle θ and $\alpha = \frac{1}{N-1}$, where N is the dimensionality of the image.

This tensor can be eigen-decomposed to produce eigenvalues and eigenvectors comparable to those used in Hessian-based image processing approaches.

Appendix C

Using the Hungarian Algorithm for Matching Particles

C.1 The Matching Problem

In the context of protein colocalisation we have a simple scenario: objects in image A need to be paired with objects from image B. We do this by posing the problem of matching pairs as an assignment problem where we use the distance between objects in space as the cost of pairing. For example, if we have objects a , b and c at $(1,1)$, $(4,3)$ and $(2,4)$, respectively, in image A and objects α , β and γ , at $(2,2)$, $(4,3)$, and $(2.5,4.5)$ respectively, in image B we might have the following assignment table:

	a	b	c
α	2	5	4
β	13	0	5
γ	14.5	4.5	0.5

Table C.1: Protein colocalisation framed as an assignment problem, values are distance squared.

	a	b	c
α	0	3	2
β	13	0	5
γ	14	4	0

(a) Subtract row and column minima

	a	b	c
α	0	3	2
β	13	0	5
γ	14	4	0

(b) Cover all zeros with matches

Table C.2: Matching pairs of proteins based on separation distance. (a) The minimum row and column values are subtracted from the costs, here the distances. (b) Matches are made such that the minimum number of matches are made to cover the maximum number of zeros.

C.2 Summary of the Hungarian Algorithm

Working with matrices the Hungarian algorithm consists of four steps [16]. First, the minimum value in each row is subtracted from all values in that row, *e.g.* our α row becomes $[0, 0, 4]$. Second, if any columns still have a non-zero minimum, the minimum value of each column is subtracted from all values in that column. Third, cover zeros, *i.e.* add matches between zero components using the minimum number of matches to cover all zeros in the matrix. Finally, if there are unmatched pairs, subtract the smallest value in the matrix from all uncovered elements and repeat the matching stage.

Appendix D

The Jaccard Similarity Index

The Jaccard index is a statistical measure of similarity of two sample sets. Given two sets, A and B , the Jaccard similarity coefficient is given by the intersection over the union of two sets,

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}. \quad (\text{D.1})$$

The resultant coefficient lies in the range $[0, 1]$ where a value of zero implies no overlap and a value of one implies 100% overlap between the two samples.

As a binary image comparison metric, *i.e.* for comparing a segmented image to the ground truth, we consider A to be the set of ‘true’ pixels in the segmented image and B the set of ‘true’ pixels in the ground truth, or:

$$J(I_{segmentation}, I_{groundtruth}) = \frac{TP}{TP + FN + FP}, \quad (\text{D.2})$$

where TP (FN) is the number of true positives (false negatives) in $I_{segmentation}$.

The Jaccard index can also be used to calculate the Jaccard dissimilarity index (or Jaccard distance),

$$d_J(A, B) = 1 - J(A, B). \quad (\text{D.3})$$

Appendix E

The Hilbert Transform

The Hilbert transform, first introduced in 1905 by David Hilbert, is an important and commonly used function in signal processing. The Hilbert transform can be considered as a multiplier operator in the Fourier domain such that,

$$\mathbb{F}(H(u))(w) = (-i\text{signum}(w)) \cdot \mathbb{F}(u)(w), \quad (\text{E.1})$$

where \mathbb{F} denotes the Fourier transform, H the Hilbert transform and signum the signum function.

The effect of applying the Hilbert transform is to shift the phase component of the function by $\pi/2$ rad.

Appendix F

Interactive GPU Level-Sets

Segmentation of objects in biological and medical images is a challenging task. Such images are often very complex and noisy and the objects of interest may be internally inhomogenous and with broken or occluded boundaries. A commonly used approach to counter these problems is the field of active contour segmentation.

Active contours, and level-set methods, allow a user to define a starting contour (or surface in 3D) and a set of energy functionals. The curve is then deformed based on these energy functionals and the image properties. Level-set methods are a particular branch of active contour implementations that deform an $(N + 1)$ -dimensional surface and take the 0^{th} level-set or isocontour as the desired N -dimensional surface.

Although active contour and level-set approaches will give accurate and precise segmentations there are several primary problems with current algorithms. Implementations are often a) too specific, relying on very restrictive assumptions or *a priori* knowledge; b) too slow, with some approaches taking several hours to complete enough iterations; and c) too difficult to use, requiring complex parameters that are not always intuitive or related to the image or object of interest.

In [6] we propose a novel GPU implementation of a specific active contour model: the local Gaussian distribution fitting (LGDF) energy model proposed by [173]. The mathematics of

this method are described fully in [173] and the implementation is introduced in [6]. The key, mathematical feature of this model is the LGDF energy term $E^{LGDF}(I, \phi)$, described in the image domain, Ω , as the sum over all pixels, \vec{p} , of a weighted sum of foreground/background ($i = 1, 2$) subregion probabilities, $p_{i, \vec{p}}(I(\vec{q}))$.

$$E^{LGDF}(I, \phi) = \int_{\vec{p} \in \Omega} E^{LGDF}(I, \phi, \vec{p}), \quad (\text{F.1})$$

$$E^{LGDF}(I, \phi, \vec{p}) = - \int_{\vec{q} \in \Omega} \omega(\vec{q} - \vec{p}) \log(p_{1, \vec{p}}(I(\vec{q}))) M_1(\phi(\vec{q})) d\vec{q} \\ - \int_{\vec{q} \in \Omega} \omega(\vec{q} - \vec{p}) \log(p_{2, \vec{p}}(I(\vec{q}))) M_2(\phi(\vec{q})) d\vec{q}, \quad \text{where} \quad (\text{F.2})$$

$$p_{i, x}(I(\vec{q})) = \frac{1}{\sqrt{2\pi}\sigma_i(\vec{p})} \exp\left(-\frac{(u_i(\vec{p}) - I(\vec{q}))^2}{2\sigma_i(\vec{p})^2}\right). \quad (\text{F.3})$$

Here $u_i(x)$ is the local foreground/background ($i = 1, 2$) intensity mean and $\sigma_i(x)$ the local foreground/background standard deviation. Where the foreground/background areas can be delimited by the C^∞ regularised Heaviside function, such that $M_1(\phi(\vec{q})) = H(\phi(\vec{q}))$ and $M_2(\phi(\vec{q})) = 1 - H(\phi(\vec{q}))$. Finally, $\omega(\vec{q} - \vec{p})$ is a normalised Gaussian weighting function centred on \vec{p} .

This level-set approach allows ϕ to change at all points each iteration, preventing the level-set from being trapped in a local energy minima [176]. Further, this method is able to segment objects with inhomogenous internal intensity distributions rather than segmenting foreground areas as those areas with uniform intensity. This makes this approach ideal for vasculature where the intensity of blood vessels varies throughout the system.