Image processing and analysis methods in quantitative endothelial cell biology

William S. N. Grimes

MRC Laboratory for Molecular Cell Biology
University College London

This thesis is submitted for the degree of
Doctor of Philosophy

September 2016

Declaration

I *William Grimes*, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

William S. N. Grimes
September 2016

Acknowledgements

I would like to acknowledge the support of my supervisors Prof. Dan Cutler, Dr Janos Kriston-Vizi, and Dr Hwee-Kuan Lee. Their guidance and critique through the duration of this research work has been invaluable. Grateful thanks to members of the Cutler Lab at the Laboratory for Molecular Cell Biology in London: Jess McCormack, Kim Harrison-Lavoie, David Westmoreland, Francesco Ferraro, Mafalda Silva, and Francesca Patella. Without their consistent provision of high quality experimental data and microscopy images this work would not have been possible. Thank you also to members of the Computer Vision and Pattern Discovery group at the Bioinformatics Institute in Singapore: Patrick Koh, Chao Hui Huang, Emarene Kalaw, Yuan Ing Chow, Michael Loh, Malay Singh, and Konnie Koh. Their help with computer programming, image processing, and statistical data analysis has been invaluable. The following people are mentioned for their collaborative input and helpful discussions: Karen Page, Mike Shaw, Alex Knight, Thomas Boudier, Daniel Holdbrook, Lin Gu, Vassilis Kargas, and Michael Criddle. Finally, my thanks are extended to the Medical Research Council (UK) and A*STAR (Singapore) for providing funding and resourses for this reserach.

Abstract

This thesis details the development of computerised image processing and analysis pipelines for quantitative evaluation of microscope image data acquired in endothelial vascular biology experimentation. The overarching objective of this work was to advance our understanding of the cell biology of cardiovascular processes; principally involving haemostasis, thrombosis, and inflammation.

Bioinformatics techniques are increasingly necessary to extract and evaluate information from biological experimentation. In cell biology advances in microscopy and the increased acquisition of large scale digital image data sets have created a need for automated image processing and data analysis. The development, testing, and evaluation of three computerised workflows for analysis of microscopy images investigating cardiovascular cell biology are described here.

The first image analysis pipeline extracts morphometric features from high-throughput experiments imaging endothelial cells and organelles. Segmentation of endothelial cells and their organelles followed by extraction of morphometric features provides a rich quantitative data set to investigate haemostatic mechanisms.

A second image processing workflow was applied to platelet images obtained from super-resolution microscopy, and used in a proof-of-principle study of a new platelet dense-granule deficiency diagnostic method. The method was able to efficiently differentiate between healthy volunteers and three patients with Hermansky-Pudlak syndrome. This was achieved by segmenting and counting the number of CD63-positive structures per platelet, allowing for the differentiation of patients from control volunteers with 99% confidence.

The final workflow described is a video analysis method that quantifies interactions of leukocytes with an endothelial monolayer. Phase contrast microscopy videos were analysed with a Haar-like features object detection and custom tracking method to quantify the dynamic interaction of rolling leukocytes. This technique provides much more information than a manual evaluation and was found to give a tracking accuracy of 92%.

These three methodologies provide a toolkit to further biological understanding of multiple facets of cardiovascular behaviour.

Table of contents

List of figures

List of tables

List of algorithms

List of abbreviations

ADP	adenosine diphosphate
AI	artificial intelligence
ATP	adenosine-5'-triphosphate
CARS	coherent anti-Stokes Raman spectroscopy
CCD	charged-coupled device
CHS	Chediak-Higashi syndrome
CLAHE	contrast limited adaptive histogram equalisation
CSR	complete spatial randomness
DIC	differential interference contrast
ELISA	enzyme-linked immunosorbent assay
GDP	guanosine diphosphate
GTP	guanosine-5'-triphosphate
HPS	Hermansky-Pudlak syndrome
HUVEC	human umbilical vein endothelial cell
MOTA	multiple object tracking accuracy
MOTP	multiple object tracking precision
OpenCV	open source computer vision library
PCA	principal component analysis

PBS	phosphate-buffered saline
PMA	phorbol 12-myristate 13-acetate
PSF	point spread function
SHG	second-harmonic generation microscopy
SIM	structured illumination microscopy
$siRNA$	small interfering ribonucleic acid
SNR	signal-to-noise ratio
SVM	support vector machine
TGN	trans-Golgi network
$THP-1$	A human monocytic cell line derived from an acute monocytic leukemia patient
vWD	von Willebrand disease
vWF	von Willebrand factor
WPB	Weibel Palade body

Chapter 1

Introduction

Modern biology is an increasingly quantitative science in which higher level insights are being obtained by delving deeper into the intricate mechanisms of biological processes. An abstract approach using statistics and modelling is required to infer patterns and trends that cannot be recognised intuitively. Increasingly sophisticated and diverse microscopy techniques are yielding larger and more complex data sets. Optical technologies and instrumentation are providing previously unimagined capabilities in image acquisition. For instance robotic microscopy and modern immunostaining methods have increased the speed and capacity by which microscopes gather image data from experiments. The consequences of these trends has led to the necessity of automated computerised image analysis approaches to harvest data from the image based data sets. These analyses have been facilitated by the rapid advances in computer hardware and improved analytical techniques.

This thesis describes computerised methods developed for extracting and interpreting quantitative output from images obtained from experiments investigating endothelial, haemostatic, and inflammatory processes. The overarching objective of this work was to advance our understanding of the cell biology of cardiovascular processes, principally involving haemostasis, thrombosis, and inflammation. In this chapter a general introduction to the various facets of this interdisciplinary research is provided. Starting with an introduction to endothelial cellular biology in section 1.1, followed by a description of microscopy techniques in section 1.2, and finally an introduction to digital imaging, and techniques in image processing and analysis is described in section 1.3.

1.1 Endothelial cellular biology

In complex organisms the circulatory system facilitates the transport of lymph and blood; the system supporting the circulation of lymph is referred to as the lymphatic system, and the system responsible for the circulation of blood is referred to as the cardiovascular system. The cardiovascular system is a closed network comprised of: blood, the heart, and blood vessels. It is extensive, such that nearly all cells in the majority of tissues of a vertebrate are located within 50 μm to 100 μm of a capillary [1]. The blood circulating in the cardiovascular system is comprised of: plasma, red blood cells, white blood cells, and platelets. In addition to circulating blood the cardiovascular system circulates and transports nutrients, oxygen, carbon dioxide, and hormones. This extensive transport network is crucial to allow blood circulation for the movement of oxygen and nutrients to tissue, and the removal of waste material from tissues.

Failure to maintain the physical integrity of the closed, distributed, and pressurised network forming the cardiovascular system leads to bleeding and a loss of pressure. These have serious impact on the viability of an organism, and thus there are mechanisms in place for maintaining the physical integrity of the system. Two principle mechanisms that maintain the physical integrity of the cardiovascular system and underpin the work for this thesis are:

- the system's ability to repair holes in vessel walls, and
- the ability to allow for controlled exit through the vessel wall where repair of underlying tissue is required.

Controlled exit from the vasculature may be for example at sites of tissue damage, where leukocytes are transported from the circulature into the tissue. Repairing holes in the vessel walls and controlled exit from the vasculature are both actively supported by those cells lining the vasculature, called endothelial cells.

1.1.1 Endothelial cells

Endothelial cells line the entire cardiovascular system from the heart to the smallest capillaries. They are fundamental to the formation of the vascular network, including the formation of capillaries, arteries, and veins. Development of connective tissue surrounding a blood vessel wall is initiated by signals from endothelial cells [1]. In this way arteries and veins can develop from small vessels composed only of endothelial cells and a basal lamina.

As a lining to the cardiovascular system, endothelial cells form the interface between plasma circulating in the blood and the surrounding tissues as illustrated in Figure 1.1a. At this interface position the thin endothelial cell lining assumes the role of an active selective permeability barrier, controlling the movement of materials and white blood cells to and from the bloodstream [1]. The endothelial function as a permeable barrier is aided by its organisation into a thin layer only a single cell thick. Cells in the endothelial monolayer also tend to be flat with a thin profile as can be seen in Figure 1.1b.

Vascular endothelial cells must respond rapidly to the changing vascular environment. They respond to fluid flow by aligning and elongating in the direction of flow [2]. Endothelial cells also respond to physical and chemical stimuli within the circulation to regulate: haemostasis, the vasomotor tone, and the immune and inflammatory responses. Endothelial cells play a crucial role in angiogenesis, the formation of new blood vessels [3].

The endothelial contribution to processes in haemostasis and the inflammatory response are largely supported by the presence of a unique storage organelle. These organelles, called Weibel-Palade bodies (WPBs) provide a rapid-response initiation of haemostasis and inflammatory responses within the dynamic and rapidly changing vascular environment.

1.1.2 Weibel-Palade bodies

The WPB was first discovered in 1964 by Ewald Weibel and George Palade, it was described as a 'rod-shaped cytoplasmic component, which consists of a bundle of fine tubules, enveloped by a tightly fitted membrane' [4]. They were later shown to be secretory vesicles and to contain a set of molecular cargoes that can be released in response to endothelial activation from an external stimulus. Activation initiates an exocytic response to mediate endothelial processes via cellular signalling or mechanical stress. WPB organelles measure between 0.1 μm to 0.3 μm in diameter and 0.5 μm to 5.0 μm in length [5]. They carry molecular cargoes that can be secreted to act as mediators for inflammation and haemostasis.

Two key molecules released in WPB exocytosis are von Willebrand factor (vWF) and P-selectin, which are critical to haemostasis and inflammation, respectively. Tubules of vWF are stored within WPBs, giving WPBs a distinctive 'cigar shape', as illustrated in the diagram in Figure 1.1c. The central role of WPBs in haemostasis is provided by its content of the platelet-recruiting cargo protein vWF [6]. Strands of vWF unfurl to capture circulating platelets and form a platelet plug, these strands of vWF can be seen in the WPB undergoing exocytosis in Figure 1.1c. This mechanism is discussed further in subsection 1.1.3. As well

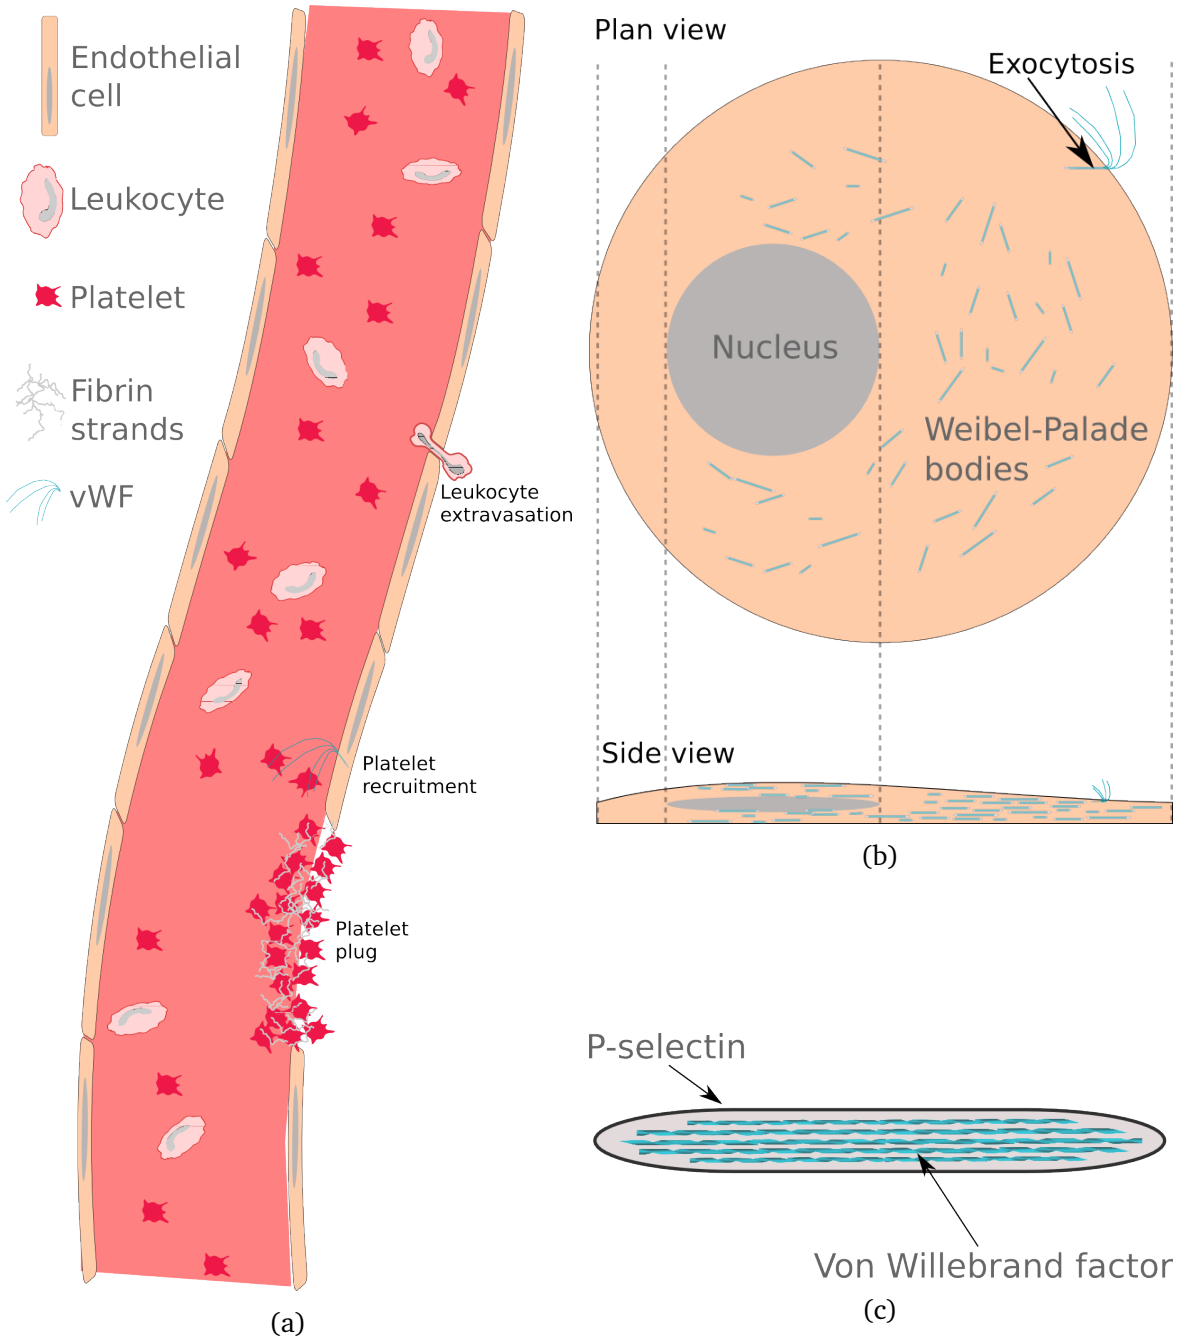

Fig. 1.1 Diagrams to elucidate the machinery for cardiovascular repair at a vascular, cellular and organelle scale. The diagram in (a) shows the endothelial lining of a blood vessel with circulating platelets and leukocytes, responsible for providing *first aid* in the vascular environment. In this diagram is an example of von Willebrand Factor (vWF) recruitment of platelets at the site of vascular injury and leukocyte extravasation. Diagram (b) depicts an endothelial cell and its content of the storage organelles Weibel-Palade bodies, where one Weibel Palade body is undergoing exocytosis. Finally, diagram (c) shows the *cigar-shape* WPB and two proteins: von Willebrand factor essential for the haemostatic response and P-selectin important in the mechanisms of inflammation and the immune response.

as underpinning primary haemostasis the multimeric vWF protein acts to chaperone the blood clotting protein factor VIII.

It was discovered in 1989 that WPBs also contain the leukocyte adhesion receptor P-selectin in their limiting membrane [7, 8]. This discovery revealed a second role of WPBs in the recruitment of leukocytes to the site of injury to support an inflammatory response (see subsection 1.1.4). The combination of the haemostatic response provided by vWF and the inflammatory response provided by P-selectin makes WPBs in Denisa Wagner's words 'the perfect first aid kit after an insult to the vasculature' [9].

1.1.3 The role of Von Willebrand Factor in haemostasis

The physiological response at the site of blood vessel injury that reduces and stops bleeding by creating a barrier to blood loss is called haemostasis. A haemostatic equilibrium is integral to the circulatory system, where an imbalance can lead to morbidity and mortality [10]. Secretion of vWF from endothelial cells initiates haemostasis after injury and provides several essential functions:

- to bind and transport blood clotting protein factor VIII,
- to mediate platelet adhesion to reactive surfaces, and
- to mediate platelet aggregation and thrombus growth [11].

Haemostasis

The process of haemostasis involves multiple contributions from the vascular system, the coagulation system, the fibrinolytic system, from platelets, the kinin system, from serine protease inhibitors, and from the complement system [12]. This complex interconnected process is orchestrated via a sequence of three events:

1. vascular constriction,
2. the formation of a platelet plug, and
3. the formation of a blood clot as a result of blood coagulation [13].

The first event triggered following vessel injury is vascular constriction or vasoconstriction. This is a localised response by the smooth muscle surrounding the endothelium that causes the vessel to narrow, reducing the diameter of the lumen and the flow of

blood to the wounded area, thereby reducing the loss of blood via the wound. To prevent further extensive blood loss additional haemostatic mechanisms are initiated leading to the formation of a blood clot.

Within seconds after vascular injury platelets interact with each other and with the injured vessel. They start to adhere to and accumulate at the site of the damage to form a primary haemostatic plug. Endothelial cells surrounding the vessel injury initiate the capture of platelets by releasing their vWF molecular cargo. In order to recruit circulating platelets strands of vWF are anchored to the plasma membrane of the endothelial cell, and to exposed subendothelial matrix. The accumulated platelet plug acts to temporarily arrest bleeding, it can however easily be broken or dislodged from the vessel wall. Therefore a secondary haemostatic response is required [10].

Secondary haemostasis is the third part of the haemostatic mechanism, involving the strengthening and stabilising of the primary haemostatic plug. This is achieved by the addition of insoluble strands of fibrin, occurring at approximately 15 s to 20 s after severe injury or 1 min to 2 min after minor injuries [13]. The fibrin is produced through a series of biochemical reactions involving coagulation factors associating with the platelet plug and injured vessel. The addition of fibrin to the primary haemostatic plug creates the secondary haemostatic clot [14]. The secondary haemostatic platelet clot stabilised by strands of fibrin can be seen in Figure 1.1a. This secure secondary haemostatic clot remains until the vessel injury has healed, upon which it is degraded and removed by further components of the haemostatic system, in a process known as fibrinolysis.

The cardiovascular system inhibits clot formation in the absence of injury by maintaining a nonreactive environment for the components of the haemostatic system. This is achieved by modulating functions that form and prevent blood clots. In the absence of vessel injury, electrostatic repulsion exists between the surface of the endothelial cell and proteins or platelets. Following an endothelial injury however, high shear flow conditions occur, due to perturbed flow, endothelial cells are activated, and platelets come in contact with subendothelial collagen and vWF. This causes platelet adhesion and initiates the process of primary haemostasis [3].

The production and processing of vWF is therefore one of the primary thrombogenic functions of endothelial cells. During vessel injury, vWF is secreted into the plasma from the luminal side of the endothelium. VWF plays an important role in the initial stage of clot formation by binding to collagen fibers in the extracellular matrix and supporting the binding of platelets, as well as recruiting platelets through long-range contacts extending from the endothelial surface into the plasma [15].

Von Willebrand Factor

The large, multimeric, multi-domain glycoprotein vWF is expressed by endothelial cells and platelets. In endothelial cells the vWF is produced as ultra-large vWF in WPBs, in platelets vWF is present in α-granules, and it is also present in subendothelial connective tissue following basolateral constitutive secretion from the adjacent endotheium [16]. It has been established that synthesis of vWF is a prerequisite for WPB formation, this is demonstrated by the absence of WPBs in endothelial cells from vWF deficient animals [17]. In the bloodstream vWF is a cofactor to factor VIII, preventing its proteolysis when vWF is depleted factor VIII is rapidly removed from the blood stream [16].

In endothelial cells vWF is synthesised at the endoplasmic reticulum as a monomer known as *pre-pro-VWF*. Further processing along the secretory pathway occurs following its arrangement into dimers in the endoplasmic reticulum. At the trans-Golgi network (TGN) the vWF is folded into helical tubules [18], that then co-associate to drive the formation of organelles recognisable as WPBs. At this time, cleavage of the pro-peptide and the formation of vWF concatamers of up to 50 dimers long begins to occur, and continues from early formation to full maturity.

1.1.4 The role of P-selectin in the inflammatory response

A number of physiological changes occur following tissue injury, where tissue injury may occur from traumatic injury, infection or cellular necrosis for example. These physiological changes can include local vasodilation, increased permeability of the capillaries to plasma proteins, and infiltration of the damaged tissues by leukocytes [13]. This set of changes constitutes the inflammatory response, in which the primary objective is to localise and eradicate the irritant and repair the surrounding tissue. The infiltration of leukocytes to damaged tissue is facilitated largely by the endothelial cells in a process known as extravasation, which is begun via the leukocyte adhesion cascade. Specifically the leukocyte receptor P-selectin present on WPBs initiates the primary step of leukocyte capture in the inflammatory response.

Inflammatory response

Leukocytes migrate from circulating in the blood to the surrounding tissue at the site of tissue damage or inflammation. This transfer of leukocytes out of the vasculature is known as extravasation and forms part of the innate immune response. The arrest of circulating leukocytes and their exit from the vasculature relies on appropriate expression

and coordination of cell surface adhesion receptors on the leukocytes and endothelium [19]. The series of adhesion and activation events from the arrest of leukocytes leading to their extravasation is known as the leukocyte adhesion cascade and is shown in Figure 1.2. A sequence of five steps have been identified in this process. From the capture of the leukocyte, to rolling, slow rolling, adhesion strengthening and paracellular or transcellular migration. The leukocytes are then able to migrate through the basement membrane to the site of damage [20].

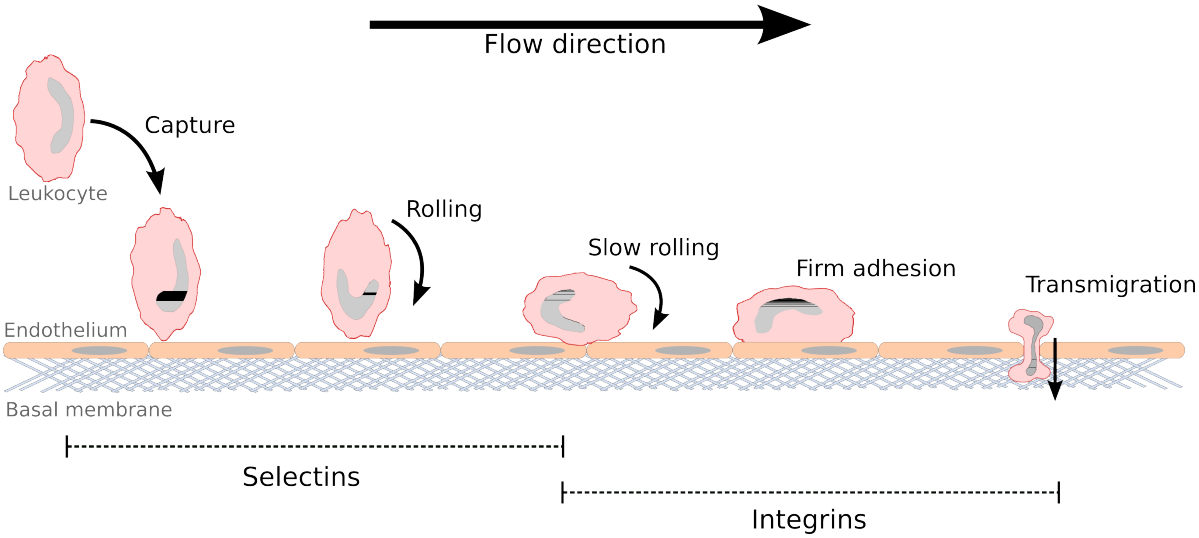

Fig. 1.2 An illustration of the key stages in the leukocyte adhesion cascade, along with an approximate timeline for the involvement of selectins and integrins in leukocyte adhesion. Transmigration maybe paracellular or transcellular, here the paracellular case is shown. Adapted from Ley et al., 2007 [20]

P-selectin

A variety of adhesion receptors have been implicated in leukocyte endothelial interactions and are thought to play roles in recruitment, arrest, and extravasation of leukocytes. These can broadly be split into selectins and integrins, where selectins tend to be involved in the earlier stages of the leukocyte adhesion cascade and integrins in later stages (see Figure 1.2).

The selectins are a family of adhesion receptors that have been shown to mediate rolling of leukocytes. There are three known selectins; L-Selectin, E-Selectin and P-Selectin, which are encoded by three closely linked genes [21]. The integral membrane protein leukocyte receptor P-selectin is stored within the membrane of WPBs of endothelial cells [7, 8], from where it is delivered to the cell surface within minutes after secretagogue-triggered

exocytosis [22]. P-selectin thus plays a key early role in the inflammatory trafficking of leukocytes, being the first receptor involved in recruiting leukocytes from flowing plasma to the endothelial surface [23]. In chapter 4 the role of P-selectin in leukocyte extravasation and the inflammatory response is investigated via a novel approach to quantitative analysis in fluid flow assays.

At later stages in leukocyte extravasation important roles are played by the family of $\beta 2$ or leukocyte integrins [24, 19]. The genetic disease leukocyte adhesion deficiency I (LADI) is caused by a mutation of the $\beta 2$ integrin gene. The increased concentration of circulating leukocytes in patients with LADI implies that $\beta 2$ or leukocyte integrins play an important role in leukocyte extravasation. Patients with LADI are susceptible to frequent bacterial and fungal infections [25]. Effective arrest of leukocytes at an inflammation site requires local activation of circulating leukocytes. The circulating leukocytes transiently interact with the activated endothelium, reducing their velocity in the direction of flow, allowing activation of their $\beta 2$ integrins, leading to arrest and extravasation. This transient interaction is manifested as rolling of the leukocytes along the vessel wall in areas of inflammation [19, 26].

1.2 Microscopy

There are broadly three categories of microscopy: light microscopy, electron microscopy, and scanning probe microscopy. The oldest being light microscopy, which uses a series of lenses to create a magnified image of an object via the diffraction, reflection, or refraction of electromagnetic waves in the visible light spectrum. Developments in light microscopy since the 17th century have lead to the plethora of diverse and sophisticated instruments available today [27]. More recent alternatives to light microscopy are electron microscopy, and scanning probe microscopy. Electron microscopy images a sample by the interaction of an accelerated beam of electrons with the sample. An accelerated electron beam gives a higher resolving power than that available with visible light microscopy. This higher resolution can be achieved because excited electrons can have a shorter wavelength than visible light (see subsection 1.2.1). The third microscopy technique is scanning probe microscopy, this diverse family of methods scans the specimen to study the surface properties via a physical probe. Scanning tunnelling microscopes (STMs) and atomic force microscopes (AFMs) are two prominent types of scanning probe microscopes.

Light microscopy is a key tool in cell biology because: the resolution is well suited to imaging cellular structures, light is relatively nonperturbing meaning that living cells can be imaged over time, and a variety of fluorescent probes can be used to mark proteins and organelles [28]. There are many types of light microscopy. A useful distinction is between microscopy techniques which require the use of labels such as fluorescent dyes and proteins and label-free methods. The latter label free methods include bright field, dark field, phase contrast, differential interference contrast as well as non linear methods such as coherent anti-Stokes Raman spectroscopy (CARS) and second-harmonic generation microscopy (SHG) [29].

Bright-field is the most elementary microscopy technique where a sample is illuminated and imaged by its effect on the light incident on it as it absorbs, scatters, or diffracts the light. The extent of absorbency of transmitted light by the sample gives rise to varying contrast in the image. Since cells tend to be thin and transparent they often transmit light well, resulting in subtle phase differences and contrast differences. Additional optics may be added to amplify the phase shift of light induced by the cells to see extra detail in the sample. Two common techniques used to amplify this phase shift are phase contrast microscopy or differential interference contrast (DIC) microscopy.

The second class of light microscopy called fluorescence microscopy uses a different method to illuminate the sample. Fluorescence microscopy makes use of the property of

chemical compounds called fluorophores. These fluorophores absorb light at a certain wavelength and emit light at another longer wavelength. Fluorophores can be introduced by the experimenter and used to fluorescently label certain proteins, cellular structures or organelles. The fluorophores may be introduced either by genetically encoding a fluorescent protein or by binding of a fluorescently labeled antibody [28]. Fluorescence techniques are especially powerful because multiple fluorescent labels can be applied to simultaneously label specific structures. Different labelled structures can be imaged by changing the wavelength of incident light or by separating the emission spectra. Fluorophores are also detectable with a very low abundance [30].

Images acquired by light microscopy methods form the basis for image processing and analysis in this thesis. Although electron microscopy was also used to image specimens at higher resolutions the resulting images were used in a more qualitative sense. In chapter 2 endothelial cells were imaged by a confocal high-throughput microscopy approach. A structured illumination microscopy technique to obtain high resolution images of platelets was used in chapter 3. Videos of leukocytes interacting with the endothelium in chapter 4 were acquired via a phase contrast microscopy method. A large determinant in choosing the appropriate microscopy technique was to obtain images with sufficient contrast and a resolution appropriate to the smallest structures to be imaged.

1.2.1 Resolution

Optical resolution describes the ability of a system to resolve detail in the object that is being imaged. The optical resolution of an imaging system depends on its components such as: the lens, digital sensors, and display components. The minimum resolving power or resolution limit is defined as the minimum distance apart of two points in order for them to be resolved as separate objects [31].

A fundamental feature of image formation is the point spread function (PSF). The ideal PSF is the three-dimensional diffraction pattern of light from an infinitely small point source in the sample, which is transmitted onto the image plane through an infinitely high numerical aperture objective. In the image plane light waves converge and interfere at the focal point to produce a diffraction pattern. An example diffraction pattern and a three-dimensional representation of the intensity distribution are shown in Figure 1.3. The series of concentric rings in the diffraction pattern is known as the Airy pattern, and the central bright disk is termed the Airy disk. The resolving power of a lens can be calculated

by measuring the diameter of the Airy disk, where the radius of the disk is determined by the numerical aperture.

For a microscope objective the limit of resolution refers to its ability to distinguish between two closely spaced Airy disks in the diffraction pattern. All optical systems have a resolution limit from diffraction occurring at the objective lens aperture. This diffraction for a circular aperture causes infinitely small points within a sample to be seen as Airy disks in the image. These Airy disks are observed as bright disks surrounded by concentric rings [27].

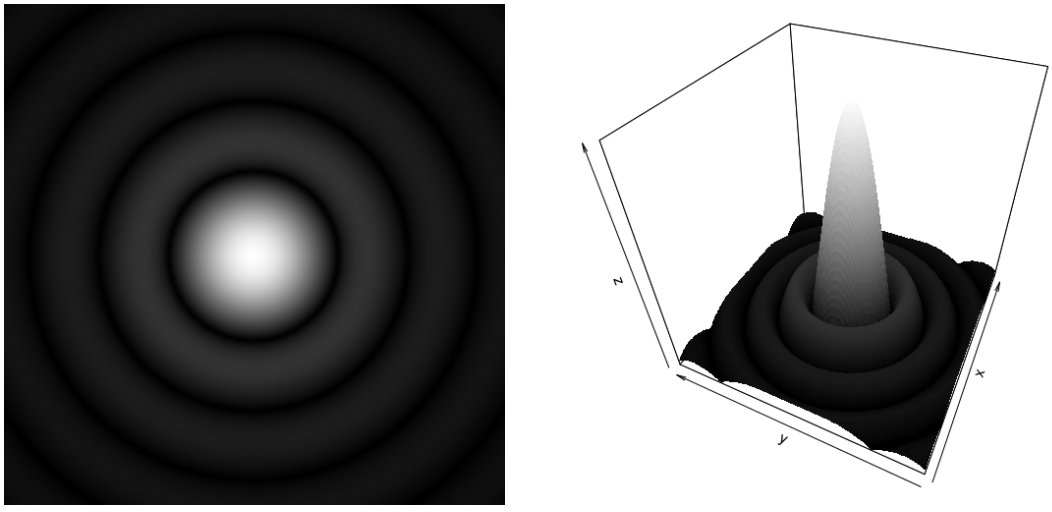

(a) (b)

Fig. 1.3 The diffraction pattern from a point spread function is plotted in image (a), and plotted in three dimensions in image (b), where the relative intensity is plotted in the z-axis as a function of spatial position for the point spread function

Principal determinants of the resolution are the wavelength of incident particles and numerical aperture of the microscope. The numerical aperture of a lens is a measure of its ability to gather light and to resolve fine specimen detail while working at a fixed object distance. There are several equations commonly used to describe the relationship between the resolution (r), the wavelength (λ), the objective lens numerical aperture (NA_{obj}), and the condensor lens numerical aperture (NA_{cond}):

$$r = 0.61\frac{\lambda}{NA_{obj}}, \tag{1.1}$$

$$r = 0.5\frac{\lambda}{NA_{obj}}, \tag{1.2}$$

and,

$$r = 1.22\frac{\lambda}{NA_{obj} + NA_{cond}}. \tag{1.3}$$

These equations should not be considered as absolute but rather guidelines that can vary according to theoretical predictions of lens behaviour, the difference in multiplication factor between Equation 1.1 and Equation 1.2 demonstrates this.

Imaging at an appropriate magnification and resolution are large determining factors when choosing the optical setup for an experiment. Spatial resolution requirements in an imaging system are defined by the separation of objects in the sample.

1.2.2 Confocal microscopy

In 1957 Marvin Minsky patented the concept for a stage-scanning confocal optical system [32]. A confocal microscope offers several key advantages over conventional wide-field microscopy, including: the ability to control the depth of field, eliminating out of focus image-degrading light, and being able to collect sequential optical sections enabling the reconstruction of three-dimensional structures.

Viewing a sample with a conventional wide-field fluorescence microscope floods the sample evenly in light. All fluorophores in the specimen are excited and the resulting fluorescence is detected at the microscope camera. This gives a plane of in-focus light, as well as degrading contributions from out of focus planes over the depth of field. A confocal microscope on the other hand only images light in focus, this is achieved using point illumination of the sample and a pinhole aperture in a conjugate plane in front of the detector. The pinhole is positioned in a plane conjugate to the focal point of the lens thus it is a confocal pinhole.

As can be seen in Figure 1.4 the grey ray paths that are not from the in-focus focal plane do not pass through the detector pinhole aperture so are not imaged. Only light from fluorescence axially close to the focal plane is detected. Confocal microscopes have an image resolution in the axial depth plane much finer than a wide-field microscope. Since a pinhole aperture blocks much of the light from reaching the detector, longer exposures are often required to obtain images with a suitable signal-to-noise ratio. Longer exposures are also required for suitable signal-to-noise because a confocal laser scanning microscope is point scanning. The pixel dwell time is much less than equivalent in a (parallel) widefield system in which all camera pixels receive photons for the duration of the total image exposure. Additionally, to increase the signal intensity the confocal optical setup often

includes a photomultiplier tube before the light detector to amplify the light signal (see Figure 1.4).

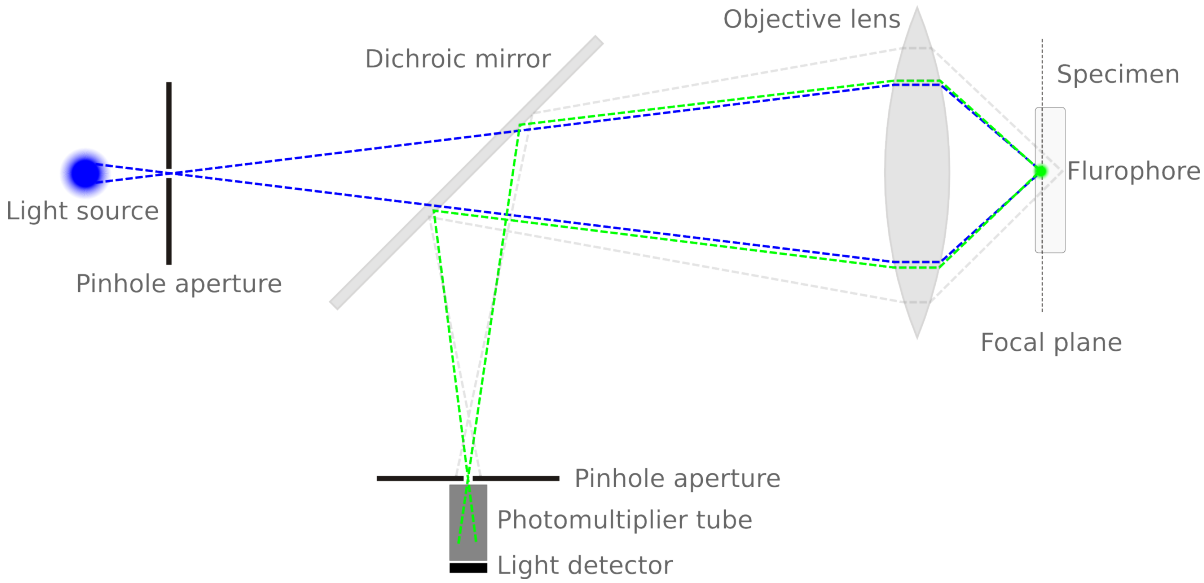

Fig. 1.4 Schematic to illustrate the optics of an epi-fluorescence confocal microscope, with a fluorophore that absorbs blue wavelength light and emits in the green wavelength. Pinhole apertures at the light source and detector allow the focus of light rays only at the point in focus on the specimen to pass to the detector. The dashed grey paths show light that is not in the focal plane of the image so does not not travel through the pinhole aperture to the detector and is thereby eliminated in the final image.

An important feature of confocal microscopes is the dichroic or dichromatic mirror. In order to excite the fluorophores in the sample in Figure 1.4 the light incident on the sample needs to be blue, a longer wavelength green light is then emitted by the fluorophores. A dichroic mirror reflects light shorter than a certain wavelength, and only transmits light longer than that wavelength. In this way at the detector only the emitted green light is observed and not the scattered blue light.

The schematic in Figure 1.4 is capable of imaging a single point source, but in order to obtain a full image from point illumination the sample is scanned in a regular raster. Two methods were originally proposed for obtaining this raster image, either by moving the sample or using a system of oscillating scanning mirrors [32]. The oscillating mirrors are the more common method since they allow for more precise control in the horizontal plane and do not disturb the sample. A series of two-dimensional rasters can also be acquired at sequential focal depths and reconstructed into a three-dimensional *z-stack*, or displayed as a maximum intensity projection. The obtainable resolution in the vertical direction will

usually be less than that in the horizontal plane and will depend on he wavelength of the light and the numerical aperture of the lens.

The development of lasers and the first confocal laser scanning microscope, along with steady improvements in raster scanning has made confocal microscopy a hugely practical and widespread technique in the life sciences [33]. In this body of work confocal microscopy has been used extensively with an automated high-content screening approach as described in chapter 2.

1.2.3 Structured illumination microscopy

The family of technically innovative methods of increasing the limits of optical resolution by at least a factor of two are collectively known as super-resolution microscopy techniques. Super-resolution microscopy is a form of light microscopy that allows samples to be imaged at resolutions higher than the resolution limit as stated by Ernst Abbe in 1873 (see Equation 1.2) [31]. These techniques have increased the maximum attainable resolution from light microscopy from $\sim$250 nm to $\sim$10 nm [34]. Advances in super-resolution technologies are either based on tailored illumination, nonlinear fluorophore responses, or the precise localisation of single molecules [35]. A commonly used method in super-resolution microscopy is structured illumination microscopy (SIM).

The super-resolution technique SIM images a sample using a series of sinusoidal or striped illumination patterns [36]. Typically the striped illumination is produced by a laser passing through a grating and projected by the objective lens onto the sample. When illuminated with this patterned light an interference pattern called a Moiré fringe is produced, which shifts high spatial frequency information to lower frequencies that are detectable by the microscope. Figure 1.5a shows a Moiré fringe arising from two superimposed patterns, in the area of overlap between the two patterns a lower frequency coarse pattern can be observed. By illuminating a sample with patterned light in this way high frequency information beyond the resolution limit of the objective lens is shifted to lower frequencies.

To construct an image using the high frequency information contained in Moiré fringes a series of images with different Moiré fringes are acquired. This is achieved by capturing images of the sample under illumination with different known patterns, the high frequency information encoded in the Moiré patterns can be seperated and shifted back to its true origin in Fourier space. Since the illumination pattern is known and can be mathematically described, it can be deconvolved from the Moiré pattern to observe higher frequency

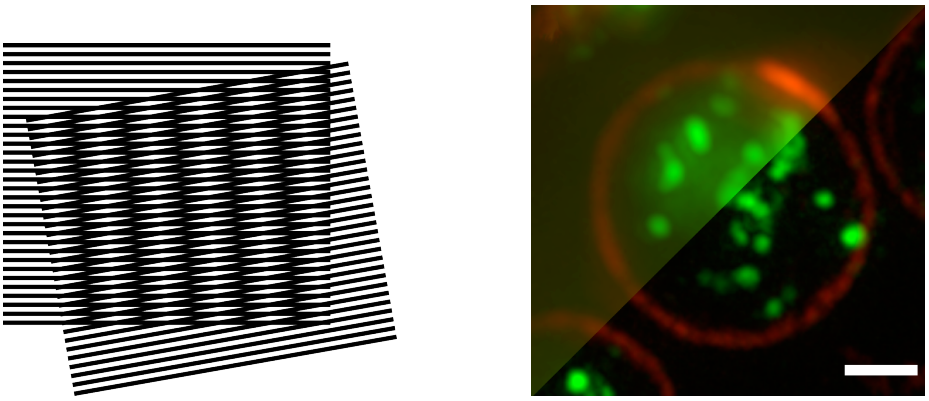

(a) (b)

Fig. 1.5 In image (a) the Moiré effect can be seen from two sets of striped lines rotated and offset, this is a key concept behind structured illumination microscopy. The image in (b) shows a platelet imaged both under conventional wide-field light microscopy method in the upper left region and a structured illumination microscopy technique in the lower right region, adapted from Westmoreland et al., 2016 [37]. Scale bar: 1.0 µm.

information that is otherwise irresolvable. A high-resolution of the underlying structure can be acquired by using algorithms to deconvolve and combine the acquired images.

In this work SIM has been used to image platelets and their dense-granules to form the foundation of a new tool for diagnosis of patients with dense-granule disorders. The maximum obtainable resolution by linear SIM is limited to a two-fold increase on that available by conventional microscopy since the periodicity of the illumination pattern is created by diffraction [36]. SIM increases the image resolution to ∼100 nm [38]. This increased resolution was necessary to image platelet dense-granules for that cannot easily be resolved through conventional light microscopy (see Figure 1.5b).

1.2.4 Phase contrast microscopy

In cell biology many specimens are highly transparent with little natural pigmentation. The differences in light absorption between different cellular structures are minimal, and much of the contrast and detail between structures is difficult to detect by conventional microscopy. Although the various cellular organelles tend not to be good light absorbers they can have considerable variations in refractive index [29]. Light travelling through these structures with different refractive indices is slowed by differing amounts, this can be measured as a phase shift in the light waveform. Phase contrast microscopes allow us

to visualise differences in optical path length, which is the product of the geometric path length and the refractive index of the medium.

The interaction of light travelling through a medium causes changes in the amplitude and phase of the wave dependent on properties of the medium. The amplitude changes occur from scattering and absorption of light, which is visible as the brightness of the light. Phase changes in the light are not visible without specialised equipment. Phase contrast microscopes feature an optical system that transforms differences in the phase of non-diffracted light, and light diffracted by a specimen into contrast differences in the image (see Figure 1.6). In this way transparent structures that may be otherwise unobservable are visible as if they had been optically stained. A phase contrast microscope measures the phase shift or retardation caused by the specimen in the optical path relative to background light. The velocity (v) of light in a medium is given by $v = c/n$, where (c) is the speed of light in a vacuum, and (n) the refractive index of the medium. Light waves traveling through an object travel slower and emerge phase shifted relative to the background rays.

The specialised phase contrast microscope has a phase annulus that generates a hollow cone of non-diffracted illumination. This hollow cone of light is projected onto the back focal plane of the objective (see Figure 1.6). The phase plate in the objective of the phase contrast microscope performs two functions: it phase shifts the light by a quarter wavelength ($\lambda/4$), and it attenuates or absorbs the non-diffracted light by 70 % to 80 % [27]. The light is focused on the image plane and in regions where the specimen is present, constructive interference occurs between the non-diffracted and diffracted light, this results in an increase in contrast compared to regions not containing the specimen.

Phase contrast microscopes are well suited to imaging thin specimens ($<5\,\mu m$), for thicker specimens the *halo-effect* can reduce the image quality [27]. The formation of halos around objects can arise due to refractive gradients passing the non-diffracted light outside the phase contrast ring and makes it difficult to determine the edges of cell structures [39].

In this work phase contrast microscopy has been used to acquire multiple sequential frames in a video sequence of leukocytes flowing over an endothelial monolayer. Interactions of the leukocytes with the monolayer were quantified and used to infer biological function, for this live cell assay phase contrast microscopy was well suited, this is described further in chapter 4.

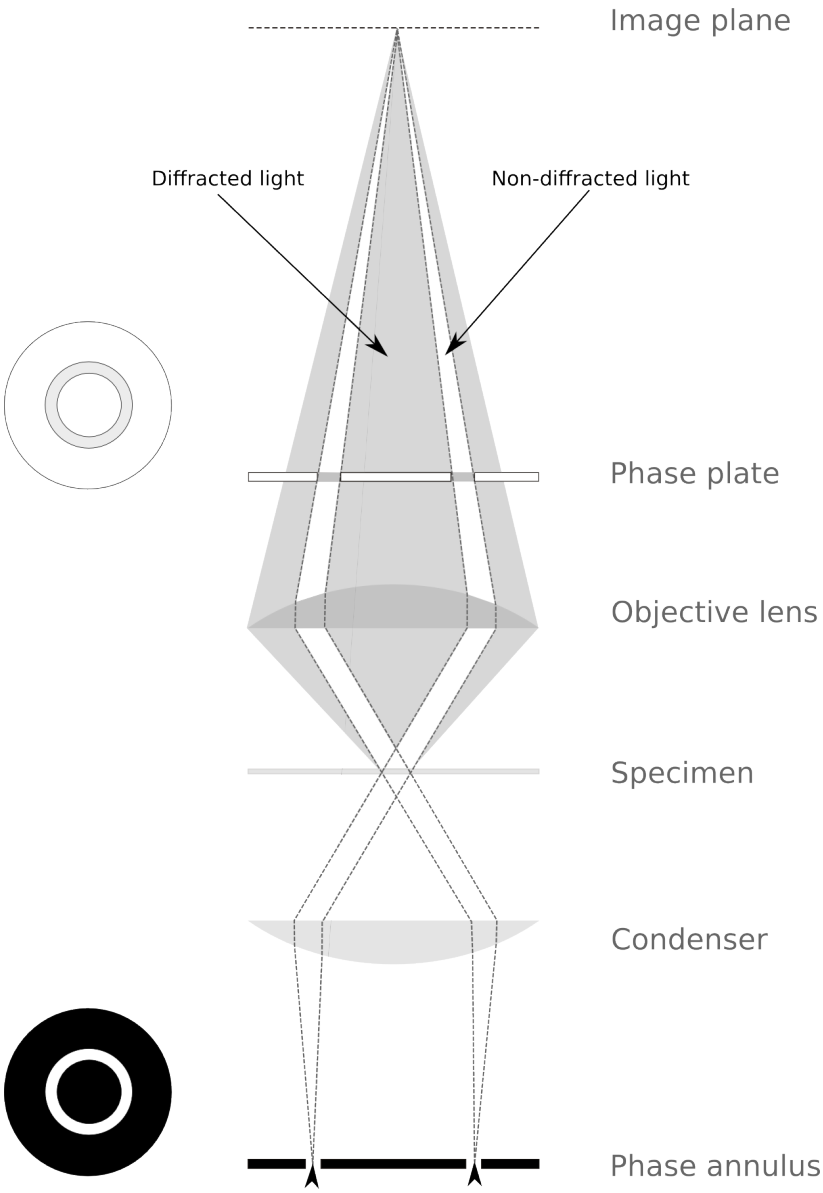

Fig. 1.6 Schematic to illustrate the optics of a phase contrast microscope. The light paths for diffracted and non-diffracted light through the condenser annulus and phase plate are shown, a plan view of the phase annulus and phase ring is also shown.

1.3 Digital image processing

The earliest discovered fossilised eyes date from the lower Cambrian period approximately 540 million years ago [40]. From a patch of photoreceptors the visual system has evolved into a complex and essential sensory system. In humans the visual system is adept at intuitively perceiving and interpreting visual information to understand a scene. It has the ability to seamlessly filter irrelevant information such as noise or illumination variations within a scene. Computer vision uses techniques from fields such as image processing, pattern recognition and machine learning to replicate functions of the human visual system. In computer vision images are stored in computer memory as digital raster images, containing a fixed number of rows and columns of pixels (picture elements).

Computer vision and image processing are becoming increasingly relevant in biological tasks such as in measurement, evaluation, and interpretation of microscopy images. These computerised approaches further biological insight by performing image processing, segmentation, recognition, and interpretation. As modern biology shifts towards experiments involving large, highly quantitative data sets manual approaches to image interpretation are becoming antiquated. With such large experimental data sets an expert visual inspection and annotation is not feasible. Furthermore, human inspection is not suited to detecting subtle population differences over tens, or hundreds of thousands of images.

Nascent computerised image processing and analysis have become a vital part of the experimental workflow. An automated or semi-automated computational approach to image processing is able to process large data sets in a fraction of the time it would take to perform the analysis manually. Further additional benefits to computerised image analysis over a manual visual inspection include:

- providing unbiased highly quantitative scoring,
- being highly reproducible,
- being able to handle large data sets,
- being able to detect subtle population differences,
- eliminating tedious manual labour, and
- being able to extract and simultaneously collect multiple features [41].

The detection and measurement of cellular structures in microscopy images is essential to cell biology. Segmentation and feature extraction are also important topics in computer

vision and image processing. Thus developing algorithms to automatically identify and measure structures of interest within biological images provides the basis for many productive collaborations between biologists and computational image processing algorithm developers. As yet general purpose image analysis methods have not been established due to the diversity of microscopy and image acquisition methods employed in cell biology. Image analysis approaches remain tailored to the specific experimental setup, however as artificial intelligence (AI) and machine learning methods continue to develop more generic methods could become viable.

A general approach to experimentation in life sciences is to investigate defined research questions or hypotheses via quantitative analyses of experimental data. A single experiment or a series of experiments may be constructed to investigate a hypothesis. For a hypothesis to be accepted results must be reproducible over multiple experiments and corroborated by alternative experimental methods.

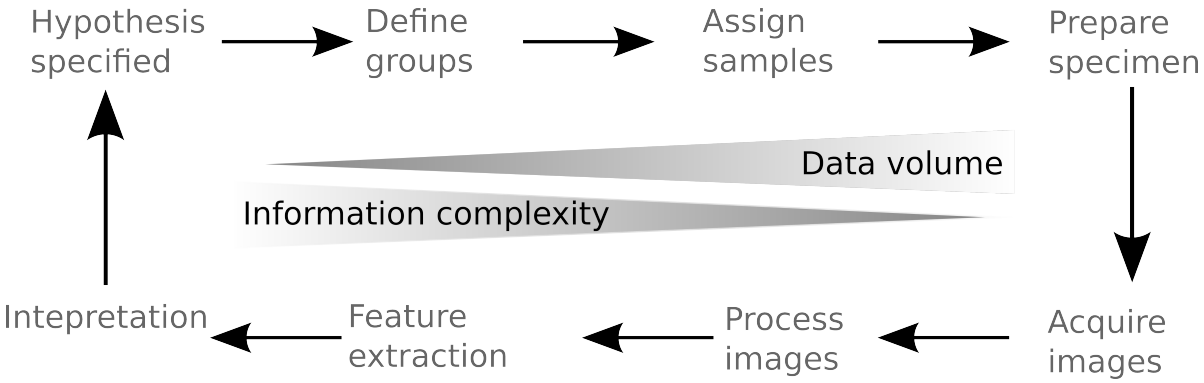

Fig. 1.7 A conceptual loop depicting an experimental workflow in life sciences involving microscopy imaging. The inverse correlation between data volume and information complexity is shown. Adapted from Prodanov and Verstreken, 2012 [42].

The experimental loop displayed in Figure 1.7 shows typical stages in experiment design and analysis. The top row of processes deals with the biological hypothesis and experimental setup, and the bottom row of processes are relating to image analysis and interpretation of the results. The bottom row of processes comprises the image analysis workflow, which converts the raw data into a form that can quantitatively address the initial hypothesis. Collection of data, aggregation of data and additional analysis allows for interpretation of the data. The loop is self-perpetuating since often the investigation of a hypothesis leads to investigation of further hypotheses.

This section gives some background to digital images, pre-processing, segmentation, morphological processing and feature extraction methods that have been developed in

this work. These methods and their applications are described in more detail in chapter 2, chapter 3, and chapter 4.

1.3.1 Image acquisition and digitisation

Digital images are acquired by the conversion of information transmitted as electromagnetic radiation into electrical signals that are measured and stored as values in a digital format. Incident electromagnetic radiation in the form of photons are converted at photosensitive detectors on a charged-coupled device (CCD) into electrons generating a voltage. The voltage can be measured and stored as a number indicating the intensity of incident light.

The intensity values measured by the CCD are stored as a matrix of values in a digital image, this can be processed by a computer and displayed. The formation of a digital image requires careful consideration of the image acquisition approach and method of digitisation. Optimisation of image acquisition and digitisation ensures the greatest efficacy of subsequent image processing and data analysis. Image analysis can rarely extract good data from poor quality images, and it is often easier to improve the data quality by optimising the image acquisition than by performing additional image processing [43].

An important factor in image digitisation is to consider the sampling. Sampling is the process of converting a continuous analogue signal into a discretised finite set of values. Spatial sampling is given by the image resolution, whilst amplitude or intensity sampling is given by the image bit-depth and is known as quantisation. The effects of spatial and intensity sampling are illustrated in Figure 1.8.

Image resolution depends on the number of pixels recorded, the optics of the system, and the electronic properties of the CCD. The image resolution may be limited by the microscope optics (see subsection 1.2.1) or by the spatial sampling rate. The higher the image resolution the truer the digital representation is to the original continuous signal. A low resolution image can distort the real observations. In the context of microscopy, resolution is the conversion of physical space into a set of equally spaced grid points on an image [44]. The finer the two dimensional grid the greater the detail that can be observed in the image.

Image quantisation determines the set of allowed numbers to represent intensity values of pixels in an image [45]. The quantisation level is expressed as bit-depth, where the greater the bit-depth the more allowed light intensity values there are. Having more allowed intensity values allows for the representation of more shades and colours in the image. Frequently used raster image bit-depths are 8-bit and 16-bit; an 8-bit image allows

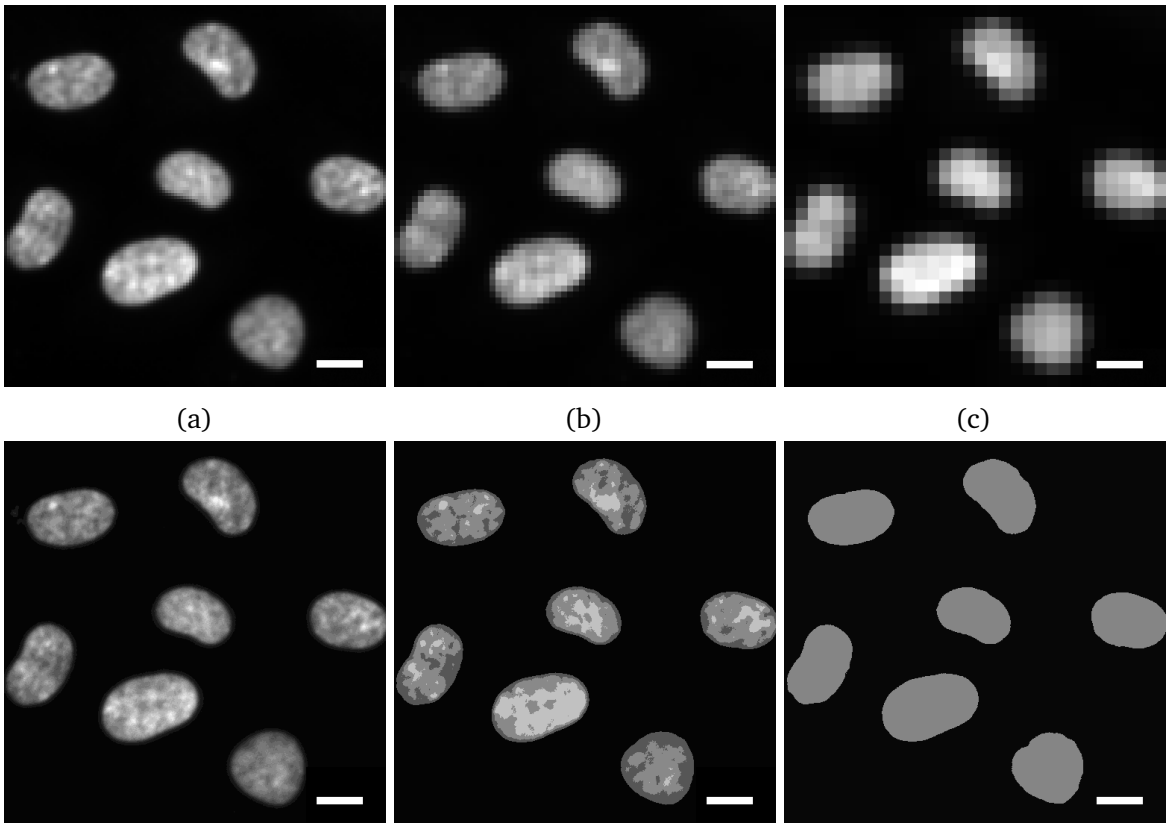

Fig. 1.8 A depiction of the effects of sampling spatially and by intensity amplitude in digitisation of images. The top row of images show the effect of down-sampling the spatial resolution, where image resolutions are: (a) 128×128, (b) 64×64, and (c) 32×32. Images have been resized to illustrate the loss of detail. The bottom row of images show the impact of quantisation where the number of intensity levels are: (a) 16, (b) 4, and (c) 2. Scale bars: $10.0\,\mu$m.

for 2^8 or 256 intensity values and a 16-bit image allows for 2^{16} or 65536 intensity values. Some microscopes also use 12-bits per channel, giving 2^{12} or 4096 possible grey values. The number of quantisation levels should be high enough for human perception of fine shading details in the image.

The process of digitisation leads to noise in images manifest as variance in the intensity values above and below the true intensity values of the signal [46]. Noise adds a level of uncertainty to the accuracy of measurements due to imprecision in measurements of pixel intensity values. Poisson noise or shot noise arises in an image from the fact that intensity values are obtained from measurements of photons arriving at a CCD. These events have an intrinsic uncertainty as they are stochastic quantum events limited by Poisson statistics [47]. Poisson noise cannot be reduced or eliminated, although its effect can be minimised by increasing the number of signal photons counted, the Poisson noise then becomes a smaller fraction of the signal. As the signal increases relative to the noise level, measurements of the signal become increasingly more precise. The precision of quantitative microscopy measurements is therefore limited by the signal-to-noise ratio (SNR) of the digital image.

1.3.2 Image pre-processing

Pre-processing of images occurs at the lowest level of abstraction where both the input and output are intensity images [48]. The necessity of this stage in the image analysis workflow is dependent on the quality of images from acquisition. Where possible it is preferable to improve image acquisition rather than create elaborate pre-processing pipelines. Nonetheless, pre-processing of images can often be useful to suppress information that is not relevant to the particular image analysis.

The aim of pre-processing is to improve the image data by enhancing certain image features, or to correct for defects in acquired images. These defects may arise for example due to imperfect detectors, limitations on microscopy optics, or non-uniform illumination. A useful categorisation of image pre-processing splits methods into four basic types: pixel brightness transformations, geometric transformations, local neighbourhood pre-processing and image restoration [48]. This categorisation is based on the size of the pixel neighbourhood from a single pixel, in pixel brightness transformations, to the whole image in image restoration. In this work pixel brightness transformations and local pre-processing were mainly used so are described further.

Pixel brightness transformations A brightness transform modifies pixel values depending on the properties of the pixel itself. A useful grey-scale transform for contrast enhancement can be achieved using the *histogram equalisation* technique. The aim of this is to distribute brightness levels over the intensity range of the image. Also in this group of brightness transformations are the adaptive histogram enhancement, and contrast limited adaptive histogram enhancement [49]. Pixel brightness transforms involving histogram enhancements are especially useful for increasing contrast in displayed images to better observe image features.

Filtering Local pre-processing operations by filtering use a neighbourhood of surrounding pixels in a kernel to transform the intensity of pixels within an image. This kernel or convolution mask usually comprises a rectangular neighbourhood with an odd number of pixels in rows and columns, thus enabling specification of a central pixel in the kernel. The size and the shape of kernel will depend on the objects in the image to be processed. If for example the objects are large relative to noise as is the case in Figure 1.8 then a small kernel is suitable and will not significantly degrade the large structures.

Filtering methods using a kernel may be used for smoothing, where the principle use of image smoothing kernels is to suppress image noise. This is often achieved by calculating an average brightness value of a neighbourhood. The median filter for example considers the neighboring pixels, sorts them to find the one with the median value and assigns that value to the pixel of interest. A similar operation called the *mean filter* averages the values of neighbourhood pixels and assigns that to the pixel of interest. The Gaussian filter is also useful for smoothing images using a Gaussian kernel. Smoothing methods that can reduce noise while preserving texture patterns and edges in an image are often more useful in improving the outcomes of object detection and segmentation, however are generally computationally more expensive [48]. Examples of edge-preserving smoothing methods that smooth textures whilst retaining sharp edges are: the Bilateral filter, the Guided filter, and Anisotropic diffusion using the heat equation [50].

1.3.3 Segmentation

Image segmentation is a crucial prerequisite step prior to feature extraction and analysis in most image analysis workflows. The goal of segmentation is to partition an image into multiple homogeneous segments, or sets of pixels. This partitioning changes the representation of an image into a more meaningful form. Segmentation is often used to

identify certain distinct structures within an image, for example in cell biology these could be nuclei or cells. This is achieved by way of assigning a label to every pixel within an image such that pixels with the same label share certain characteristics. In the most simplistic case these labels can distinguish between foreground and background pixels [44].

There are many approaches to segmentation which are being constantly developed and improved upon including: clustering methods, compression-based methods, edge detection and region growing methods. Two of the most fundamental and commonly used methods of segmentation are thresholding and the watershed transform.

Thresholding The simplest segmentation process is grey-level thresholding, where an intensity value or threshold is determined to segment objects and background. Thresholding is computationally inexpensive and fast to compute [48]. A threshold can be calculated and applied globally or locally to an image, where local thresholding can account for differences in the intensities of pixels in different regions of the image. A number of algorithms have been developed for automatic selection of a threshold value. Histogram shape analysis is often used to calculate a threshold value. For example the mode method for thresholding bi-modal histograms finds the highest local maxima and then detects the minimum between two maxima. The choice of thresholding method will depend on the specific image data and it is usually worth experimenting with many methods of obtaining a threshold value such as Otsu thresholding [51, 52] or moment preserving thresholding [53] to determine the most appropriate method.

Watershed A watershed transform is a method for segmentation based conceptually on a hydrological watershed. In a hydrological sense a watershed is the line dividing adjacent catchment basins. An image can also be visualised as a topographic surface, where the *z-axis* represents altitudes. Peaks and ridges in the image correspond to high intensity and watershed regions and low intensity regions are catchment basins. Watershed segmentation works to identify and label catchment basins in an image. There are multiple algorithms that have been developed to achieve this [54]. A one dimensional example of watershed segmentation is shown in Figure 1.9

There are two basic approaches to performing a watershed transform. The original computationally intensive method finds a downstream path to a local minimum of image topology for each pixel in the image. The catchment basin is then the set of pixels for which the downstream paths all end up in the same altitude minimum. A second approach considers the flooding of the relief and imagines the progressive filling of the valleys, but to

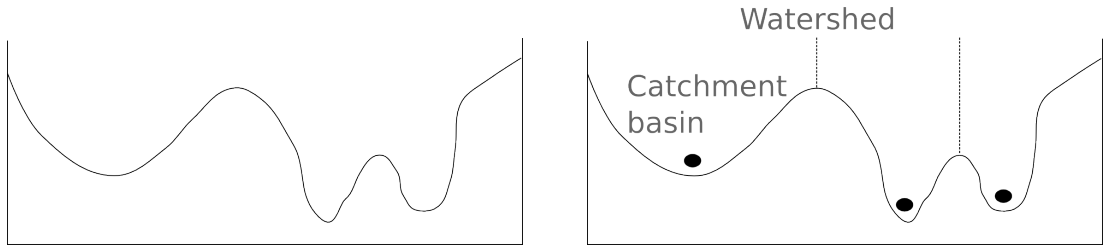

Fig. 1.9 An example of a watershed segmentation in one-dimension. The plot in (a) is a greylevel intensity profile of image data. In image (b) watershed segmentation has been performed showing the local minima catchment basins and the watershed lines.

a point that the flooding does not spill into the next valley. The flooded regions are labelled as specific segmented objects in the relief. Vincent et al., 1991 presented an efficient algorithm to perform this based on sorting pixels in increasing order of value, followed by a flooding with a breadth-first scanning of all pixels in their sorted order [55].

1.3.4 Morphological processing

In image processing morphological processing is a theory and group of techniques for performing transformations on geometrical structures, based on set theory, lattice theory, topology, and random functions [56]. These techniques involve changing the shape of objects in binary or greyscale images using morphological rules from mathematical set theory. Morphological operations on an image are relations of two sets; the image and a structuring element.

Morphological techniques manipulate an image with a kernel or template called a structuring element. The structuring element is positioned at each pixel location within an image and operations performed on the neighbourhood. An output image is generated from the interaction of the structuring element with the image. There are many morphological operations some examples include: dilation, erosion, opening, closing, skeletonisation and distance transforms.

Binary dilation and erosion The fundamental binary morphological operations are dilation and erosion. These operate on binary images comprised of black and white pixels. In binary images one class of pixels is considered foreground and the other background. From the basic morphological operations of dilation and erosion further more complicated operations arise such as opening, closing, and shape decomposition.

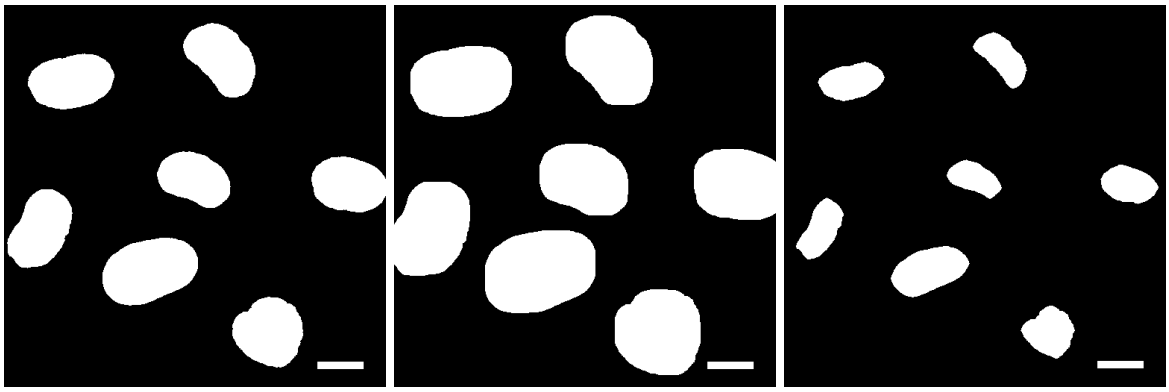

(a) (b) (c)

Fig. 1.10 An example of binary dilation and erosion performed on a thresholded image of nuclei. Image (a) is the thresholded nuclei, image (b) is the result after performing binary dilation on (a), and image (c) is the image after performing binary erosion on (a). Scale bars: 10.0 μm.

A dilation is an isotropic expansion that can be used to fill small holes or join objects together resulting in an increased object size. Conversely an erosion acts is an isotropic shrinking that reduces the size of objects in the image or if sufficiently small removes them. These operations are demonstrated in Figure 1.10. Erosion and dilation are not inverse transforms meaning that if an image is eroded and then dilated the resulting image differs from the original image, instead a simplified version of the original image is obtained with less detail. The process of erosion followed by dilation is called opening and is useful to eliminate specific image details smaller than the structuring element whilst maintaining the general shape of the objects in the image. The process of dilation followed by erosion is known as closing. This is useful to connect objects that are close to each other and fill small holes in objects.

1.3.5 Region description and features

Segmentation and morphological processing are important steps to constructing homogeneous image regions. The extraction of semantic information from these regions is important in understanding the image data. This requires a description of that region, which can be achieved by generating a numeric feature vector. This numeric vector may contain features relevant to shape, location or intensity that describe the region. There are many shape descriptors however there is no generally accepted method of shape

description. In this work a range of shape descriptors have been used, these are described further in subsection 2.3.1 and subsection 3.3.1.

The description of a region is first dependent on its identification, a useful method for region identification is to label each region with a unique integer pixel value. This process is known as labelling or colouring, in which the largest integer number also gives the number of regions in the image. For each labelled region a numeric vector of multiple features can be measured.

Region based shape descriptors may be from simple geometric descriptions such as area, eccentricity, and direction. Other shape descriptors such as statistical moments interpret a greyscale image function as a probability density of a 2D random variable, the properties of which can be described by moments. Shapes can also be described by regional decomposition into simpler shapes, for example by thinning to obtain a skeleton. Pixel intensity features can include the mean pixel value of a region, the maximum, minimum, or standard deviation for example.

1.3.6 Segmentation performance evaluation

The evaluation of image segmentation is an important and often overlooked stage of image analysis. Despite there being many algorithms and significant literature about image segmentation techniques, there is no universal agreement on how to evaluate the performance of image segmentation algorithms [57]. For any given segmentation problem there are multiple possible segmentation algorithms that could be used, finding the optimal method and parameter choice requires a choice of evaluation metrics suitable for the data use.

A helpful starting point to assess the performance of a segmentation algorithm is to review by eye the segmentation contours overlaid on the raw image. This human evaluation is tedious and limits the number of comparisons to a small sample size [58]. Often more thorough quantitative assessment is required to detect subtle differences between segmentation methods. In image processing and computer vision a comparison to a gold standard or reference segmentation data set is a commonly used and instructive method of performance evaluation.

The term gold standard refers to a benchmark that is available under reasonable conditions. The ground truth differs from the gold standard in that it is obtained by a set of measures that is known to be more accurate than the measurements of the system being tested. A ground truth could for example be obtained by performing segmentation

of higher resolution and higher magnification microscopy images. A gold standard on the other hand is not an objectively perfect set, but merely the best available [59]. A gold standard is acquired from the most accurate segmentation procedure, which in this case is generated by manually hand-labelling or semi-automatic labelling images.

Image segmentation is often an ill defined problem, meaning there may not be a single gold standard segmentation. It is therefore useful to compare against multiple perceptually consistent interpretations of an image. These might be for example hand segmented by several experts in the field. Hand segmentation of a gold standard is however a labour intensive process and this is often not feasible [60]. The evaluation of a segmentation algorithms performance and generation of a gold standard data is also complicated by inconsistent data, since images of microscope samples vary in quality and character.

A useful way of evaluating the performance of a automatically segmented image to a gold standard is to use overlap ratio measures. These give a similarity measure between the pixel sets of segmented objects. In chapter 2 the Dice coefficient and Jaccard index have been used extensively for evaluation. In other instances segmentation performance may be evaluated in a binary sense, for example when counting objects it can be useful to use binary classification tests to find the sensitivity and specificity. The sensitivity also called the true positive rate, the recall, or probability of detection in some fields measures the proportion of positives that are correctly identified as such. The specificity also called the true negative rate measures the proportion of negatives that are correctly identified as such [61].

1.4 Key aims

The computerised methods described in this body of work were developed with the intention of advancing our understanding of the cell biology of cardiovascular processes, principally involving haemostasis, thrombosis, and inflammation. A series of image processing methods have been developed to enhance the quantitative output from existing experimental methods in endothelial cell biology. These image processing pipelines aim to generally be capable of:

- extracting quantitative data from endothelial cell microscopy images,
- addressing specific hypotheses arising from endothelial cell biology,
- strengthening the link between general microscopy applications and endothelial cell biology,
- allowing for sophisticated analysis of the extracted data acquired, and
- being easily implementable by our research group.

An initial aim of this work was to further develop an image processing tool that had previously been used to quantify the morphometric features of a large population of the endothelial organelles Weibel-Palade bodies (WPBs) [5]. The state of the art commercially available software applications Acapella, Volocity and PhenoLOGIC (PerkinElmer) were trialled for this purpose but were not adaptable enough to the specific demands of this assay. Specifically, it was not possible to collect the required set of morphometric features, and perform a single cell analysis. The aim of developing a new image analysis pipeline was to accurately segment endothelial cells and their organelles, extract relevant morphometric features, and assign organelles to their cells. This computational tool will enable both whole population level and cellular level analysis, and by studying endothelial cells under different physiological and pathophysiological conditions can be used to gain an understanding of underlying haemostatic processes.

A second aim of this work was to develop an image analysis pipeline forming the basis for a platelet dense-granule disorder diagnostic tool. As part of a proof-of-principle study involving a novel application of super-resolution microscopy the image analysis tool should demonstrate the efficacy of this new imaging technique in diagnosing dense granule disorders. A minimum requirement of this image analysis pipeline was to count the number of dense granules in each platelet, and therefore be used to identify platelets from patients

with a lack of dense granules. The additional quantitative data extracted could also be used to provide further diagnostic indicators. This image analysis approach would provide an unbiased, rapidly acquired data set, forming the basis for a platelet dense-granule disorder diagnostic tool.

A final aim of this work was to enhance the quantitative output from experiments using video sequences to study leukocyte recruitment and extravasation. The generally accepted method of analysis of leukocyte endothelial flow assays has involved manually counting leukocytes that interact by slow rolling or firmly adhering. This manual approach is slow, labour intensive, subjective, and only measures the count. An automated video processing tool aims to perform an objective, faster analysis, which can extract more quantitative information. This would be achieved by accurately detecting leukocytes in each frame and tracking their movement between frames. Information about the trajectories of each leukocyte over an endothelial cell surface can then be acquired and data mining and analysis performed. This extra quantitative data from each experiment would allow for a more thorough analysis of the leukocyte adhesion cascade and leukocyte extravasation than a manual counting approach.

Generally the development of this suite of image processing tools aims to help build a multifaceted picture of the endothelial processes of haemostasis, thrombosis and inflammation. The image processing tools together with the analysis of their outputs will allow for the extraction of a great deal of quantitative information to underpin significant advances in the mechanisms that are at the heart of haemostasis, and of the initiation of inflammation.

1.5 Thesis overview

This thesis describes three image analysis methods developed for endothelial cell biology microscopy. The computational tool discussed in chapter 2 was used to quantify the morphometric features of a large population of endothelial organelles. Segmentation of endothelial cells and assignment of organelles to these cells enables both whole population level and cellular level analysis. The workflow described in chapter 2 has been applied to Weibel-Palade bodies (WPBs) and has also been applied extensively in high-throughput studies of the morphometry of WPB exocytic sites. These WPB 'exocytic sites' reflect the presence on the endothelial plasma membrane of a fusion pore as seen by fluorescent labelling of endothelial cells with vWF internally and externally. Through counting absolute numbers of exocytic sites, the number of exocytic sites per cell and the areas of exocytic sites, these analyses can be used to garner information about von Willebrand Factor (vWF) release that is complementary to and potentially more sensitive than the commonly used enzyme-linked immunosorbent assay (ELISA) method. These studies are proving to be useful in understanding the cellular mechanisms that are employed by endothelial cells in controlling vWF release.

The release of vWF into the blood stream occurs when WPBs fuse with the plasma membrane, whereupon the vWF unfurls to form long platelet-capturing strings. These platelet-decorated strings perform a vital role in primary haemostasis. Within platelets are granules containing bioactive pro-haemostatic molecules. Genetic mutations can reduce the number or capacity of platelet granules affecting their functional ability. Genetic abnormalities such as in Hermansky-Pudlak syndrome (HPS) lead to abnormalities in numbers or capacity of platelet granules resulting in bleeding disorders. Diagnostic imaging of platelets to count their dense-granules has traditionally required electron microscopy, due to the small size of the platelet granules. Dense-granules are approximately 150 nm wide and light microscopy has historically been limited to resolutions of 200 nm to 300 nm depending on the technique. However, developments in super-resolution microscopy now allow sufficient resolution in light microscopy to effectively image platelet granules. An automated image analysis pipeline discussed in chapter 3 provides a quantitative, unbiased, rapidly acquired data set, that forms the basis for a platelet dense-granule disorder diagnostic tool. Using the antigen CD63 as a marker, granules were segmented and assigned to their respective platelets. A proof-of-principle data set with seven healthy controls and three patients with platelet granule abnormalities caused by HPS has shown the effectiveness of this technique.

In chapter 4 a workflow developed for video analysis of leukocyte recruitment and extravasation is discussed. Leukocyte recruitment is part of the innate immune response whereby stages in the leukocyte adhesion cascade: initial capture, rolling, slow rolling and firm adhesion, can be imaged under flow in vitro as a time series. Previously, analysis from such videos has involved manually counting leukocytes that interact by slow rolling or firmly adhering. The method of detecting leukocytes in each frame of the video is described, and how these detected leukocytes are linked over multiple frames to form trajectories. In this chapter an evaluation of both the precision of the detection phase, and accuracy of the tracking phase is described.

Finally chapter 5 contains some concluding discussions of each of the described methods, and additionally discusses their significance and continued usage. Areas of potential future improvement and avenues for further work are mentioned.

Chapter 2

High-throughput morphometric analyses of endothelial organelles

Large volumes of imaging data are generated in high-throughput microscopic surveys of fluorescently labelled endothelial organelles. With such large-scale image data sets an automated computational approach was the only viable method of extracting biologically relevant quantitative features. The extraction of quantitative features then allows for the analysis and interpretation of each experiment.

This chapter describes methods of image segmentation and feature extraction employed for the large-scale quantitative morphometric analyses of endothelial organelles. Additionally, described here are methods of data analysis and interpretation of the resulting large and rich quantitative feature sets. These data analyses require multiple techniques drawn from fields within the broad areas of mathematics, statistics, and computer science including but not limited to: computer programming, data engineering, data mining, machine learning, predictive analytics, and pattern recognition. Data visualisation is also of particular importance to communicate the insights and findings from each high-throughput microscopic survey in a clear, efficient, and meaningful manner.

Image data for use in this chapter was acquired by: Jess McCormack, Kim Harrison-Lavoie, Francesco Ferraro, Mafalda Pinto Baptista Lopes Da Silva and Francesca Patella.

2.1 Introduction

High-throughput microscopic surveys of endothelial organelles from primary human umbilical vein endothelial cells (HUVECs) are a new and highly informative approach to the study of both inflammatory and haemostatic processes, often providing insights that are unobservable in low-throughput approaches. Automated high-throughput approaches are advantageous because:

- the data output is highly quantitative,
- the increased data yield encourages the detection of subtle phenotypic differences,
- automated computational analysis does not suffer from subjective bias,
- experiments are easily extensible,
- computational feature extraction allows for the measurement of multiple features simultaneously,
- the image acquisition is relatively fast, and
- high-throughput techniques can be used to test hypothesis from low-throughput data.

Of particular importance in studies of inflammatory and haemostatic processes are the endothelial storage organelles Weibel-Palade bodies (WPBs). These WPB storage organelles carry a large multimeric protein called von Willebrand Factor (vWF), which plays a critical role in haemostasis and inflammation. To study the biological functions of intracellular WPBs, we can monitor them by fluorescently labelling and then imaging their vWF molecular cargo. Biological function and processes can be inferred by imaging cells and comparing groups under altered cellular conditions. The cellular conditions can be manipulated with secretagogue and drug treatments. Analysis of phenotypic differences between groups offers insight into the function of WPBs, and the machinery involved in their formation, which again impacts on their functioning. Phenotypic changes can include increases or decreases in the numbers of WPB, as well as changes in morphology.

The study of intracellular WPBs was extended to investigate mechanisms of WPB exocytosis by fluorescently labelling and imaging vWF at the external endothelial cell surface. Other proteins marking additional endothelial organelles relevant to haemostasis and inflammation could also be labelled thus including markers for other organelles. Additional labels could be for example the Golgi apparatus, and the endoplasmic reticulum. This chapter does not explicitly describe methods of image processing for these organelles, since they have not yet been extensively tested.

In our high-throughput analysis of endothelial organelle morphometry, image acquisition included: stained cell nuclei, the plasma membrane, as well as intracellular, or extracellular vWF. Staining of cell nuclei and the plasma membrane were used to define cell boundaries, and allow for analysis at the level of individual cells, rather than of a whole monolayer.

Following image acquisition, a custom-made image segmentation protocol was used to threshold organelle pixels within each channel of each image in the data set. Regional properties or features were measured from segmented organelles, and if a plasma membrane stain was present then organelles features were linked to their relevant cell. A range of features were collected allowing population level analysis. Aggregation of data on a per-cell basis, also gave further richness and insight to the data. Data analysis and interpretation of endothelial organelle results can be performed at a population or on a per-cell basis, depending on the biological question being addressed.

Since *a priori* it was not known which measurements would be most useful or interesting, for each survey and each organelle group the same large feature set was collected. Collection of a large number of features provided the most freedom in data analysis. Analyses of large feature sets obtained from images of fluorescently labelled organelles also provided great sensitivity for detecting subtle phenotypes in populations at both a cellular or population level.

To ensure optimal data, careful consideration of the experimental phase and acquisition phases, and image processing approach was required. An optimal experimental workflow was achieved through multiple trials with various antibodies, microscope settings, and image segmentation techniques. This chapter details the developments of image segmentation and analysis methods for high-throughput morphometric analyses of endothelial organelles.

2.2 Image acquisition

Analyses of endothelial organelle morphometry were performed on images acquired from confocal microscopy, with separate channels for different stained cellular organelles. DNA in cell nuclei were always fluorescently stained and then additional specific organelles were stained with immunofluorescence antibody methods.

Human umbilical vein endothelial cells (HUVECs) were imaged on an Opera LX high-content screening (PerkinElmer) confocal microscope. The HUVECs were cultured, fixed and immunostained in 96-well microtitre plates (Nunc MaxiSorp) and imaged using a 40× air objective lens (numerical aperture of 0.6). Various drug and secretagogue treatments as well as changes in growth conditions, or cell activation were applied to different groups on the microtitre plate. The schematic in Figure 2.1 shows a common experimental setup with four treatment groups.

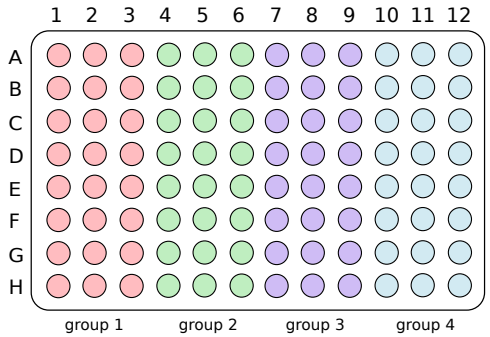

Fig. 2.1 Schematic of a 96-well microtitre plate showing four treatment groups.

In a typical high-throughput survey, nine fields of view were imaged per microtitre well, generating data sets of 864 images and ∼10000 complete endothelial cells in a plate. This provides a large volume of data improving the statistical significance when testing biological hypotheses. For certain studies the volume of data collected was scaled up by imaging more fields of view per well, without requiring more labour intensive laboratory work. At 40× magnification the length of each pixel in the obtained images corresponds to a physical length of 0.1612 µm without camera binning. Morphometric measurements could thus be converted from pixel sizes to dimensions in micrometers.

The Opera confocal microscope has an available excitation wavelength range from 488 nm to 640 nm, with a numerical aperture of 0.6; this gives a minimum resolvable resolution range from 496 nm to 650 nm, calculated with Abbe's equation (see Equation 1.2). The smallest resolvable structures in the images obtained are therefore between 2 and 4 pixels in length, depending on the channel being imaged.

In high-throughput analyses the endothelial organelles that are frequently stained are nuclei, plasma membrane, vWF internal to WPBs, and external vWF (see Figure 2.2). The hypothesis being tested dictates the experimental setup and the specific cell organelles that are immunostained and imaged. Invariably, vWF is immunostained and imaged as a marker for exocytic sites or for internal WPBs. Cell nuclei and additional plasma membrane staining allow for analysis on a per-cell basis. The PerkinElmer Opera high-throughput confocal microscope at the Laboratory for Molecular Cell Biology allows for a maximum of 3 confocal channels to be imaged and an ultraviolet (UV) channel per experiment. Along with the vWF channel, two channels were required for accurate cell segmentation, namely, the nuclei (UV channel) and plasma membrane channel. This leaves one channel that can be used to stain for additional organelles or left empty. The resulting 12-bit image data was combined as a set of images each with multiple channels in a TIFF format.

Staining External or exocytosed vWF were labelled at the cell surface to study mechanisms of exocytosis. Cell nuclei were stained either by Hoechst 33342 or DAPI (4',6-diamidino-2-phenylindole) dyes, which bind to tightly packed AT-rich regions of DNA in chromosomes within the nucleus, making an effective nuclear stain.

Immunostaining of vWF within WPBs involved the use of a primary and a secondary antibody. Over multiple experiments different antibodies have been trialled for this purpose. A Dako rabbit anti-vWF polyclonal primary antibody, is an antiserum with many different antibodies that recognise different portions of the vWF protein, since it is produced using whole vWF isolated from human plasma to immunise the rabbit. In fixed HUVECs, the Dako antibody will bind to the vWF that is in WPBs but will also bind to the unprocessed vWF that is in the endoplasmic reticulum (ER). The ER is a mesh-like structure that extends throughout the interior of the cell, thus ER staining may overlap with the WPB staining and make it difficult to identify WPB against the ER.

To overcome the ER interference problem, a different rabbit polyclonal antibody was used. This was a rabbit antibody from rabbits that were injected with a synthetic 8 amino acid peptide sequence (SSPLSHRS) that is found at the end of the proregion of vWF after it has been processed. This sequence is not available for antibodies to bind to it in the unprocessed ER form of vWF, so as a result, the antibodies will only bind to the vWF in WPBs.

Immunostaining of exocytosed vWF was performed in several steps. Cells cultured in 96 well plates were washed several times in serum-free medium and then rabbit anti-vWF polyclonal antibody (DAKO) was added to either unstimulated, or secretagogue-stimulated

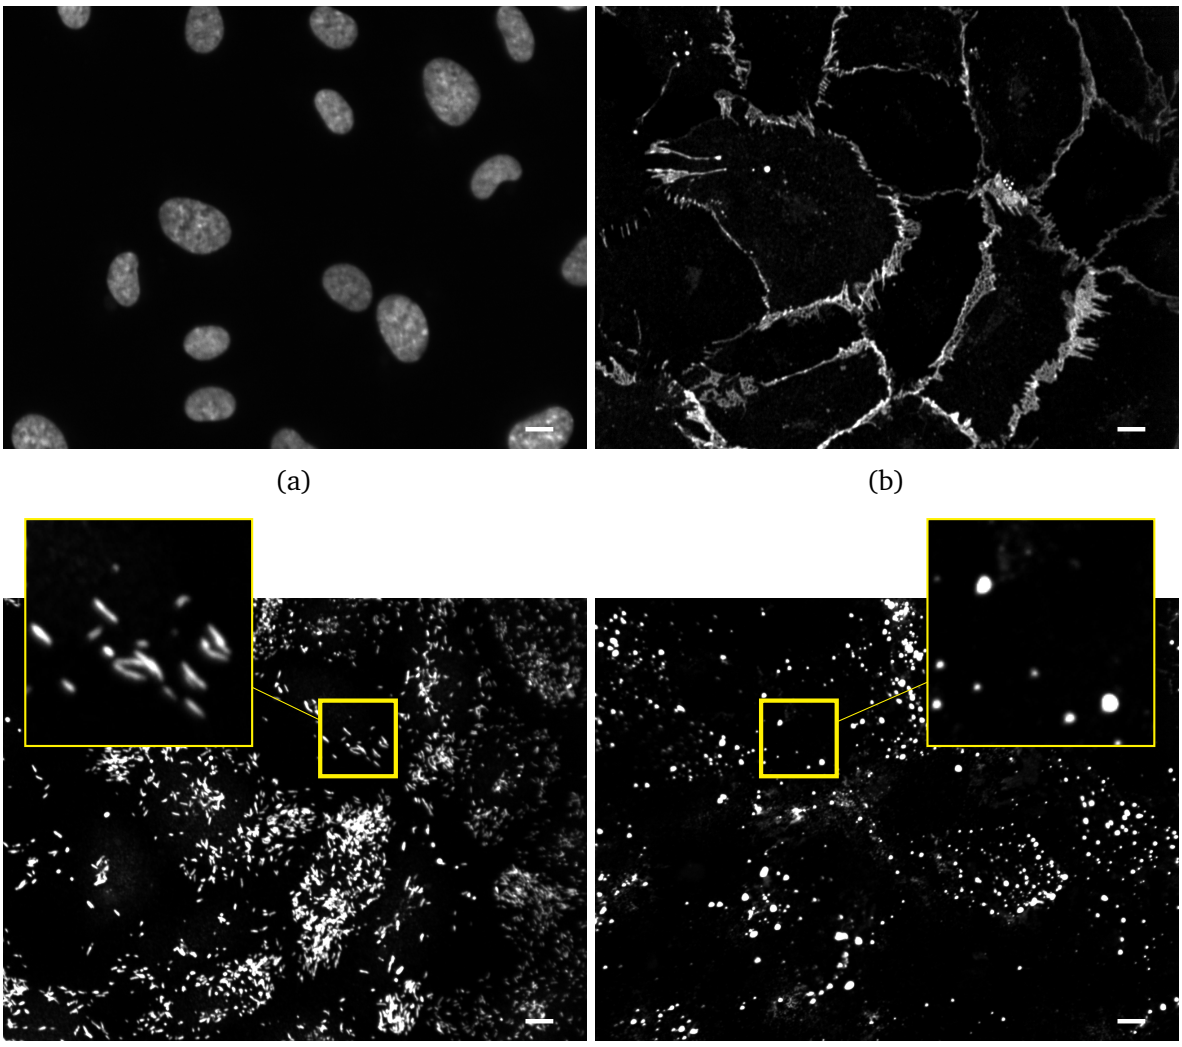

(a) (b)

(c) (d)

Fig. 2.2 Examples of confocal microscopy images acquired in a high-throughput study of endothelial cells, acquired with four labelled organelles and four channels. Image (a) is the nuclei stained channel, image (b) are stained plasma membrane, image (c) is of Weibel-Palade bodies (WPB) stained, and image (d) is stained for exocytosed von Willebrand factor (vWF). Scale bars: 10.0 µm; magnified regions: 25.0×25.0 µm.

cells. Cells were fixed and permeabilised. Following washing, samples were incubated with fluorescently-conjugated secondary antibodies to stain for exocytosed vWF (Alexa 488).

The plasma membrane was stained either by anti-VE-Cadherin antibodies or fluorophore-conjugated wheat germ agglutinin (WGA). VE-Cadherin is a cell-cell adhesion glycoprotein found in endothelial cells [62]. It is important for holding endothelial cells together with their neighbours to form a tight barrier. WGA is a member of the lectin family that binds to N-acetyl-D-glucosamine and sialic acid residues that are found on the surface of cell membranes, and that has been used to stain surface membranes.

2.3 Image processing

An image processing pipeline was developed for the segmentation and extraction of endothelial organelle morphometric features. The pipeline was developed in Python version 2.7 with extensive use of the scikit-image [63] and scikit-learn [64] libraries. A generalised image processing approach was created that is customisable to the specific experimental setup. Functionality was therefore included to allow selection of the appropriate segmentation method depending on the channels imaged and their configuration. If a plasma membrane and nuclei stain were present then cell segmentation could be performed and analysis automatically includes a cell assignment step, thereby facilitating single cell analysis.

The flowchart in Figure 2.3 presents each step in the image processing pipeline for segmentation and feature extraction of four image channels. There are some dependencies between channels, for example the plasma membrane channel is dependent on an input from the nucleus channel as the nucleus acts as a seed for the cell. The WPB channel is dependent on an input from the plasma membrane channel to assign WPB to its constituent cell. Channels are processed in an order relative to these dependencies, such that the nucleus channel is processed first, followed by the plasma membrane channel and then other cell organelles, and finally organelles are assigned to their respective cells. These inter-channel dependencies are facilitated by co-localisation in the Opera confocal microscope achieved with beads for skewcrop and flatfield correction.

The Python code was structured such that it is easily readable and adjustable, with the minimum number of parameters and alteration between experiments. It has provided a generally applicable framework for segmentation of endothelial organelles in various channel combinations and imaging modalities. Adjustable variables and parameters are in the head of the code including the channel configuration, image processing parameters, and images to exclude. Acquired images of insufficient quality may be excluded from analysis as described in subsection 2.3.2. Each organelle segmentation method was written as a separate function, which takes the input image and performs segmentation, returning found contours and a set of features.

Each image in a data set was processed sequentially and on each iteration outputs of morphological results tables and image segmentation contours overlays are saved (see Figure 2.5f, Figure 2.7f, Figure 2.8a, Figure 2.9a and Figure 2.10). On each iteration of the loop, results are appended to the results tables and written to disk, to reduce the computer memory usage. An iterative approach, although slower, is advantageous over

Fig. 2.3 Flowchart of segmentation and feature extraction processes of endothelial organelles; nucleus (A), cells (B), Weibel-Palade body (WPB) (C) and exocytosed von-Willebrand factor (vWF) (D). Dark ellipses in the flowchart are significant intermediary images, rectangles with rounded corners show image processing steps, and plain rectangles show feature extraction steps.

batch processing because it has lower memory requirements, debugging is easier and multiple plates can be run simultaneously. A summary of segmented objects in each image are also printed to the console, for example *Analysing 001001004.tif, image 4 of 864, detected: 22 nuclei, 7 cells, 1794 wpb*. These outputs also act as a quick method for the user to determine whether the segmentation is reasonable, depending on whether realistic outputs are given based on previous results. For example, if 4000 nuclei were segmented in a single image this would suggest there is a problem with the image quality or image processing, since this is an anomalously high value given the cell size and confluency density.

2.3.1 Region properties

The scikit-image library [63] in Python contains a *regionprops* function to measure properties of labelled image regions. Supplying a labelled image and the original intensity image allowed for measurements of morphometric features and pixel intensity measurements. The *regionprops* function was used for measurement of morphometric and pixel intensity features in the analysis of endothelial organelles. Pixel intensity and morphometric properties were measured for every segmented nucleus, cell and WPB or exocytic site in a high-throughput survey. Although the *regionprops* command provides a multitude of possible outputs, only pertinent single value attributes were extracted from the region properties. All the measured attributes along with a description of each one are displayed in Table 2.1. The row, column and field of view for each segmented particle are obtained from the filename of the image, as dictated by the Opera microscope filenaming system. For example a file *003002001.tif*, was imaged from microtitre plate row 3, column 2 and field of view 1. If cellular segmentation was performed then the labelled cell number was also returned.

All measurements in Table 2.1 can be converted into units corresponding to the physical size of the object, by applying a scaling factor. The scaling factor is a multiplier that maps the morphological measurements in pixels to a size in microns, or for the area microns squared. The user is able to adjust the scaling factor relevant for the microscope optical setup.

2.3.2 Data pre-processing

A step was performed prior to performing image processing of a data set to identify and remove unusable images. Images of inadequate quality may arise due to fluorescence

Table 2.1 Terminology and descriptions of endothelial segmented object attributes. Comprised of meta-data to identify the object, morphometric measurements and intensity measurements. The *particle_id*, and *row*, and *col*, and *fov* are obtained from the filename of the input image.

Attribute	Description
particle_id	The id of the cell from which the measurement was taken, consisting of a row number, column number, field of view number and cell number
row	The microtitre row number from which the image was taken
col	The microtitre column number from which the image was taken
fov	The image field of view number
cell	The cell label number of the object
x_centroid	The centroid of the region in the x-axis
y_centroid	The centroid of the region in the y-axis
area	The area of the region
perimeter	The perimeter of the region.
feret	The longest distance between any two points along the region boundary, also known as maximum caliper
equivalent diameter	The diameter of a circle with the same area as the region
convex_area	Area of the convex hull of the region
major_axis_length	The length of the major axis of the ellipse that has the same normalized second central moments as the region
minor_axis_length	The length of the minor axis of the ellipse that has the same normalized second central moments as the region
orientation	Angle between the x-axis and the major axis of the ellipse that has the same second-moments as the region. Ranging from -pi/2 to pi/2 in counter-clockwise direction
solidity	Ratio of area in the region to area of the convex hull image
max_intensity	Pixel value with the greatest intensity in the region
min_intensity	Pixel value with the greatest intensity in the region
mean_intensity	Pixel value with the greatest intensity in the region

artefacts or problems in the microscope acquisition. Fluorescence artefacts can result from bleed-through, or non-specific staining for example [46]. Problems with the microscope acquisition are manifest as very dark images, washed out noisy images, or large amorphous bright regions within an image. Figure 2.4 shows some examples of images that should be removed from further processing. Figure 2.4a contains a large amorphous bright region, while Figure 2.4b was caused by a microscope malfunction. Data pre-processing is useful to check the consistency of data over a large image set, and depending on the acquisition remove any unusable images.

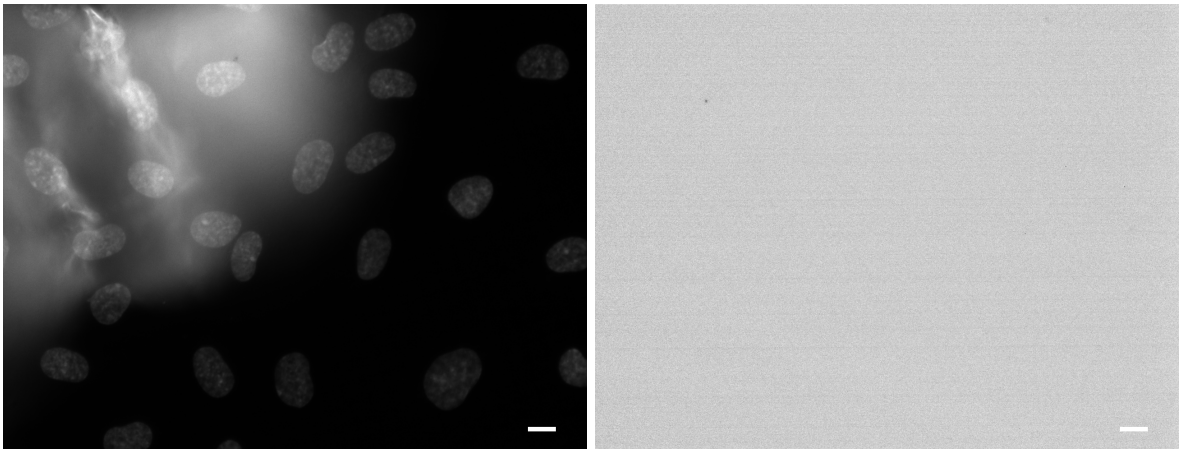

(a) (b)

Fig. 2.4 Examples of images that are unusable and have been excluded from further processing, by a visual inspection of images with high and low mean pixel values. Image (a) could be due to non-flatfield correction or occlusion by debris or dirt, and image (b) is symptomatic of a laser malfunction on the microscope. Scale bars: 10.0 µm.

Problem images in a data set tend to be at the extremes of high or low mean pixel intensity values. A rudimentary but effective approach used to identify images of an insufficient standard was therefore to sort images by their intensity. By sorting images by their mean pixel intensities and first checking the extremes of intensity it was easier for the user to find images of insufficient quality, such as those shown in Figure 2.4.

This data pre-processing quality control step was performed in ImageJ where sorting of image stacks was done using the ImageJ stack sorter plugin[1]. The following steps were carried out to identify images to exclude:

1. open each channel as an independent stack of images,

[1] B. Dougherty, Stack Sorter ImageJ Plugin, http://www.optinav.com/Stack-Sorter.htm, accessed 2015-09-14

2. sort images in each channel according to the image mean pixel value,
3. visually inspect the lowest and highest value mean pixel value images,
4. visually inspect the remaining images by scrolling through the stack,
5. identify and note the filename of images to exclude, and
6. inspect the remaining images in the stack to check for erroneous images.

In this way, images could be excluded from further data processing, improving the reliability and consistency of data. Image exclusion was performed later in Python, where all files in the specified image directory were found and a set object created. The user specifies a list of images to exclude, a second set is created from this list of images to be excluded. The exclusion set is subtracted from the set of all the images in the directory, and the result is sorted, the images could then be removed from processing.

2.3.3 Nucleus segmentation

Segmentation of nuclei was the first stage in morphometric analysis of endothelial cells. Nucleus segmentation forms the basis of cellular segmentation because nuclei are used as seed points or cell markers. Correct identification and segmentation of nuclei was therefore fundamental to cellular segmentation, and consequently the assignment of endothelial organelles to their constituent cells. The consistently high quality staining of the DAPI and Hoechst fluorescent dyes over many experiments, has led to the development of a stable optimised nuclei segmentation routine that does not need adjusting between experiments. Figure 2.3 outlines each step in the segmentation pipeline, and the major stages of nuclei segmentation are demonstrated visually in Figure 2.5.

Threshold (A.0)

The first step in segmentation of nuclei was to apply an intensity threshold to the image to create a binary image (see Figure 2.5b), separating foreground and background objects. A variety of thresholding methods were trialled for this purpose, and Otsu's method was found by eye to give an accurate segmentation contours over a sample of images. Otsu's method [51] is a widely used thresholding technique that calculates a threshold by maximising the variance between two classes of pixels, and was found to be very effective at segmenting objects in the nuclei images.

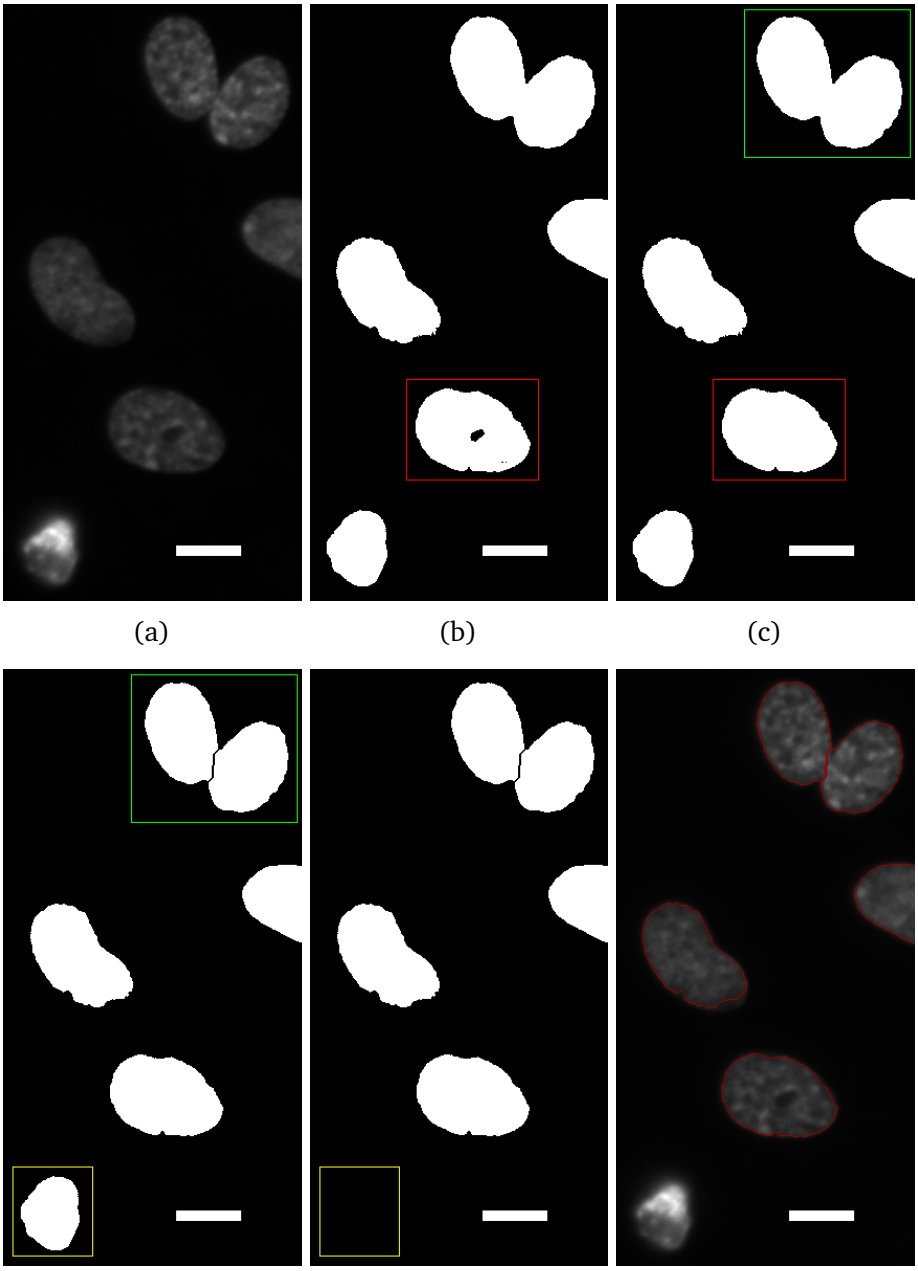

Fig. 2.5 The major sequence of processes involved in segmentation of nuclei. Image (a) is the raw image, (b) after applying Otsu's threshold, (c) following a binary hole fill, (d) depicts the watershed transform, (e) is after a classification step, and finally image (f) shows a quality control image output with segmentation contours. The red box in images (a) and (b) highlights the effect of binary hole filling, the green box in images (c) and (d) shows an example of splitting with the watershed transform, and the yellow box in images (d) and (e) highlights the removal of non-nuclei objects via classification. Scale bars: 10.0 μm.

Binary hole fill (A.1)

The mottled nature of nuclei staining and variation in staining across an image can lead to holes in thresholded foreground objects, as shown by the red box in Figure 2.5c. Holes in binary objects are problematic since they create inaccuracy in morphological measurement, specifically in the measurement of area. In addition holes in segmented objects along the image border lead to errors in image processing. Holes on the border in later processing with the watershed transform could result in incorrect splitting of a region into multiple objects. As such it was necessary to fill small holes in binary objects.

The available hole-filling algorithm in Python has the limitation that it does not fill holes in objects on the image border. The algorithm provided by SciPy[2] requires binary holes to be surrounded completely by a connected border of foreground pixels.

To overcome this, a custom binary hole filler was implemented on the condition that the background area is always larger than the largest foreground connected component. This condition should always hold in correctly acquired nuclei images. The hole filler algorithm extracts the background region by performing a connected component analysis on the image. All other pixels, that are not part of that connected component are set to foreground pixels. The result of applying the binary hole filler is exemplified by the red box in Figure 2.5b and Figure 2.5c.

Watershed transform (A.2)

The watershed transform is a versatile and commonly applied region-based segmentation algorithm [55]. It takes its name from an analogy with hydrology, where watershed lines are the divisions in a landscape splitting an area into catchment basins. The concept of watershed basins can be visualised as the domains of attraction for rain falling over a region. Alternatively, consider a landscape being submerged in a lake: basins will fill with water and at points where water from different basins meet, dams are built [54]. In extending this idea to digital images, we can consider that an image has a topographic relief dictated by its pixel values. Catchment basins in this relief can be found, and the lines dividing these basins are watershed lines. In the case of nucleus segmentation, a watershed transform was used to separate nuclei that are close, overlapping, or touching. The application of the watershed can be observed by the splitting of the objects in Figure 2.5c and Figure 2.5d.

[2]Jones E., Oliphant E., Peterson P., et al. SciPy: Open Source Scientific Tools for Python, 2001-, http://www.scipy.org/ accessed 2016-04-26

There are many technical variations in the implementations of the watershed algorithm, which can be applied depending on the segmentation problem. The watershed transform may be applied to the raw pixel intensity image, or a transformed version of the image. For separation of touching nuclei a watershed transform was performed on a transformed image. A euclidean distance transform was first applied, so that each foreground pixel is given a value, relative to the distance to its nearest boundary. A watershed on a euclidean distance transformed image gives a watershed based entirely on the shape characteristics of the binary objects, where markers are the local maxima of the distance function to the background for separating overlapping objects.

The scikit-image watershed function was used to perform a watershed in this pipeline. It apportions pixels into marked basins using the method described by Vincent et al., 1991 [55], and returns a labelled image.

Binary classification (A.3)

A supervised machine learning binary classification was used to remove incorrectly segmented foreground non-nucleus objects. A classifier was built to predict the class of segmented foreground objects, identifying nucleus or non-nucleus foreground structures. Non-nucleus segmented structures could be, for example: the result of staining artefacts, debris, dead cells or very small thresholded regions. Initial attempts at removing non-nucleus structures by removing connected components less than a specified size were inadequate. Although this size filter was effective for removing the majority of debris and small thresholded regions, it was not effective for removal of staining artefacts and dead cells.

Identifying nuclei from other structures can be approached as a binary classification problem, the objective of which is to separate foreground objects into two classes: nucleus and non-nucleus. A support vector machine (SVM) is commonly used for binary classification [65], and has been used here to generate a model for nuclei classification. Implementation of the classifier used the region properties measured morphological and pixel intensity features for each segmented object. Implementation of the classifier used the LIBSVM library [66], a radial basis function kernel, and practical advice from Hsu et al. in their 2008 paper, 'A Practical Guide to Support Vector Classification' [67].

Labelled data set To apply SVM classification a set of labelled training data was required, where each instance in the training data was assigned a class label. Nucleus objects are here labelled 1 and non-nucleus objects labelled by 0.

Generation of a labelled training set was semi-automated, to reduce tedious manual labour. A random selection of 300 nuclei stained images were taken from three separate experimental data sets, and the images were segmented by Otsu's method, hole filling, and a watershed transform. The resulting binary output image was saved for each of the 300 images, and for each segmented object a set of 14 morphometric and pixel intensity features were extracted. This gave an intensity and morphometric results features data table consisting of 14510 objects, of unknown class.

The class of each object in the features data table was determined in a semi-automated fashion. In the binary images, connected components beneath the resolution limit of the microscope were removed as noise by morphological opening. The larger remaining structures were cross referenced with their raw nucleus stained images, and where a non-nucleus structure had been incorrectly segmented its pixels were set to background. In this way all incorrectly segmented foreground object were removed from the binary images. The binary images were then used as a reference set. For each of the 14510 instances in the results features data, the pixel value at the x and y centroid position for the structure was checked. The class was assigned by reading the pixel value corresponding to the location of the segmented object in the binary reference image set. If the pixel value was that of a foreground pixel then the object was a nucleus and assigned a label 1. If the pixel value was that of a background pixel then the object was a non-nucleus object and was assigned a label of 0. Of the 14510 examples, 9425 were labelled as nucleus objects, and 5085 as non-nucleus objects.

Data pre-processing For effective training of a SVM classifier some data pre-processing steps were performed. Firstly, from the labelled data set only numeric features containing morphometric and pixel intensity information were retained. The categorical attributes specifically row, column, field of view, and particle ID numbers were removed (see Table 2.1).

The remaining features were scaled to standardise their range of values, thereby avoiding features with greater numeric ranges from dominating the function training. There are various methods of feature scaling. Here a standardisation method was used. For a sample (x), the standardised version (x') is calculated using the mean ($\bar{x}$), and standard deviation (σ), as follows:

$$x' = \frac{x - \bar{x}}{\sigma} \tag{2.1}$$

Cross-validation and grid search The classification task here involves separating the data into training and testing sets. The data in the training set contains one predicted value and multiple features. The objective of SVM classification is to use the training data to produce a model, which can accurately predict the target values of the test data given only the test data features. Training vectors are mapped onto a higher dimensional space, and the linear separating hyperplane with the maximal margin in this higher dimensional space found. There are multiple kernel functions that could be used for solving this problem. The radial basis function was used here and is usually a good choice in problems that have many more instances than features [67].

The radial basis function kernel has two parameters: the cost function, C, and gamma γ. These parameters were optimised to improve the accuracy of the classifier in predicting unknown data. A grid search was used for parameter optimisation, where a range of C and γ parameters were specified, and for every possible combination of these parameters a model trained and the predictive accuracy of the model evaluated. A coarse grid search was performed, with C values in the range 10^{-2} to 10^{10} and γ values in the range 10^{-9} to 10^{3}, as shown in Figure 2.6a. This was followed by a finer grid search, to identify a more precise C and γ value. The fine grid search used C values in the range 2^{-5} to 2^{15} and γ values in the range 2^{-15} to 2^{3}, as shown in Figure 2.6b.

A data set consisting of 14510 example labelled data points were shuffled and split into a training and validation group of 8706 examples, and a test group of 5804 examples. A 5-fold cross-validation was performed, for each C and γ value combination. In 5-fold cross-validation the training and validation group was randomly shuffled and split into 5 groups. Sequentially, each subset was tested using the classifier trained on the remaining 4 subsets. The mean predictive accuracy and standard deviation across the 5 groups was then reported. In this way, each example in the whole training set is predicted once so the cross-validation accuracy is the percentage of data which are correctly classified. Cross-validation is important to prevent over-fitting of a model.

The most accurate model was found searching the finer base 2 C and γ parameter grid search. A C value of 8.0, and γ value of 0.03125, were found to be optimal parameter choices.

Implementation Binary classification with a SVM was applied after an initial segmentation performed via thresholding, binary hole filling and a watershed transform. Features were then extracted for each segmented object. This array of features was pre-processed, by dropping non-numeric features, and scaling the remaining feature values. Feature scaling

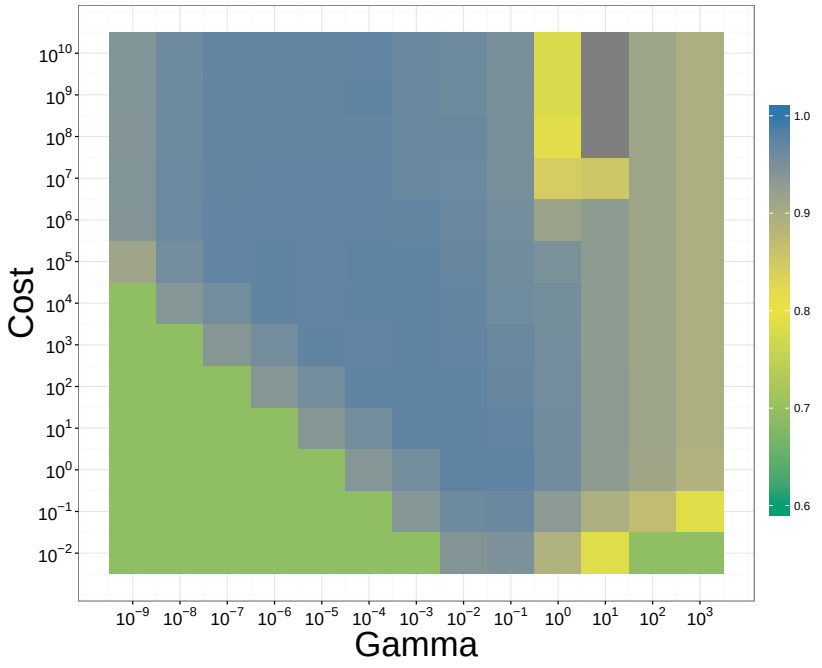

(a)

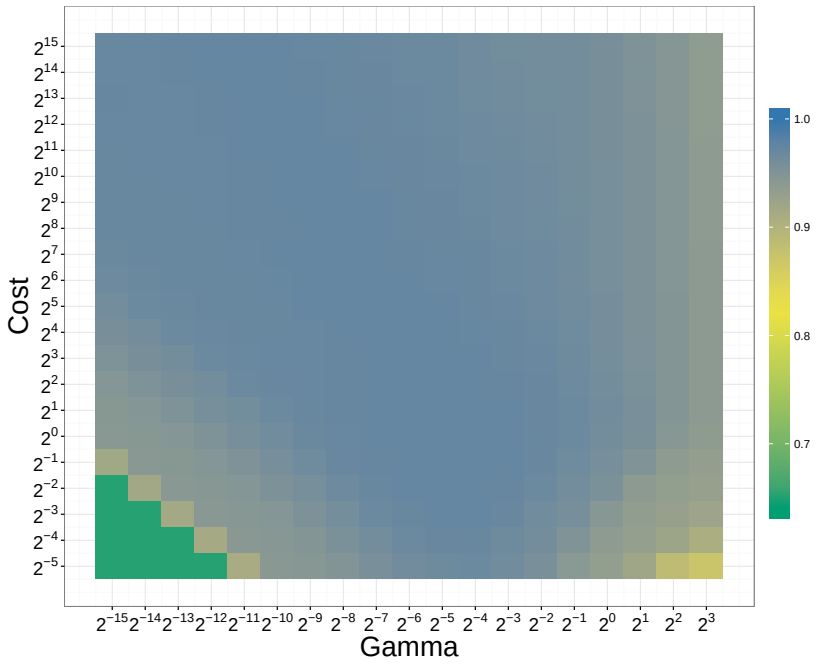

(b)

Fig. 2.6 Heatmaps depicting on a gradient intensity scale the accuracy of support vector machine classification with models trained with various combinations of C and γ parameters. A 5-fold cross-validation is used for each parameter combination so an accuracy and standard deviation of the parameter combination is returned. The plot in (a) is a coarse grid search over a base 10 logarithmic range, and in plot (b) the search was refined over a base 2 logarithmic range.

was performed using identical feature scaling values as was applied to scale the features in the training data. The SVM classifier was then used to predict the classification of all objects by their morphometric and pixel intensity features.

Objects that were predicted to be non-nucleus were removed by setting their pixels to background pixels in a labelled image. For the remaining objects, a binary image was created and the image relabelled. A final feature extraction step was performed on the labelled image, having removed predicted non-nucleus objects.

2.3.4 Cell segmentation

The segmentation of individual endothelial cells opens up a deeper analytical study of phenomena at a cellular level. In population analysis it is not possible to tell how individual cells respond to secretagogue or drug treatments. After having performed cell segmentation all organelles can be assigned to their respective cells. To do this each segmented organelle was assigned an identifier, relating to the cell from which it was derived.

Stages in cell segmentation are outlined in Figure 2.3, and the major intermediary stages are shown visually in Figure 2.7. The general approach to cell segmentation involved pre-processing steps to enhance contrast and reduce noise (see Figure 2.7b), followed by separation of cells in the confluent monolayer. A range of techniques were trialled to separate touching cells in the confluent endothelial monolayer. Methods were trialled based on a survey of automated segmentations applied to optical imaging of mammalian cells [68]. This survey of methods categorises segmentation methods as being spatially blind or spatially guided, where spatially blind techniques are essentially threshold based. As expected, spatially blind methods such as thresholding were not reliable for finding cell boundaries. These methods were not adept at handling the intensity variation along cell boundaries, variation in staining, or overlapping cells. Furthermore spatially blind techniques also do not use information provided by nuclei seed points.

Spatially guided techniques produced more accurate cell segmentation results, using nuclei as cell markers. Of the spatially guided segmentation methods region growing, evolving generalised voronoi diagram [69], random walk, level set methods and marker-based watersheds were trialled. Region growing and level sets were found to be too computationally expensive for the scale of the high-throughput surveys, and the results were inconsistent with different staining protocols. The remaining three techniques; evolving generalised voronoi diagram method, random walk segmentation and marker based watershed produced comparable results. The watershed method was used because it

was found to give the greatest accuracy according to overlap ratio measures with the gold standard (see subsection 2.4.2). The watershed has the added bonus of being efficient and implementable in Python through scikit-image [63].

To perform the watershed, the pre-processed image was inverted and the binary nuclei image added by performing an element-wise minimum with the binary nuclei image. A marker based watershed transform was used to find cell boundaries according to the nuclei markers. A cell was segmented for every nucleus, but if the cell was on the image border it has all its pixels set to background. Cells on the border region were assigned a pixel value of zero, since they are not complete cells, making single cell analysis is impossible (see Figure 2.7e). Despite these cells being excluded from single cell analysis the data can still be used for population level analysis.

Endothelial cells in microtitre plate wells form a confluent flat monolayer, where cell size, shape and orientation vary. Cells were demarcated with a plasma membrane stain, as can be seen in Figure 2.2b. The plasma membrane stain was used to segment the cell boundaries, in conjunction with seed points from cell nuclei. Since nuclei are used as markers for cells, the accuracy of nuclei segmentation is of fundamental important for ensuring accuracy of cell segmentation.

The segmentation method presented here assumes that each segmented nucleus is contained within exactly one cell. In other words, each cell can only contain one nucleus. This is almost always the case, although binucleation can occur and cells undergoing mitosis can be imaged at any stage in division. The proportion of cells undergoing mitosis and the number of binucleated cells is variable, depending on the experimental protocol.

Contrast-limited adaptive histogram equalisation (B.0)

An initial pre-processing step was to perform contrast enhancement, using a contrast limited adaptive histogram equalisation (CLAHE) [49]. Histogram equalisation improves the contrast in an image globally, by effectively spreading out the most frequently used intensity values. The adaptive method of histogram equalisation, performs histogram equalisation on local regions of the image. It is therefore suitable for improving the local contrast and enhancing the edge definition in each region of an image.

The adaptive method of histogram equalisation has the drawback that in homogeneous image regions with only small intensity variation it has a tendency to over-amplify noise. By limiting the contrast enhancement in adaptive histogram equalisation this over-amplification can be avoided. The contrast amplification is calculated for a pixel from the slope of the transformation function, which is proportional to the slope of the neigh-

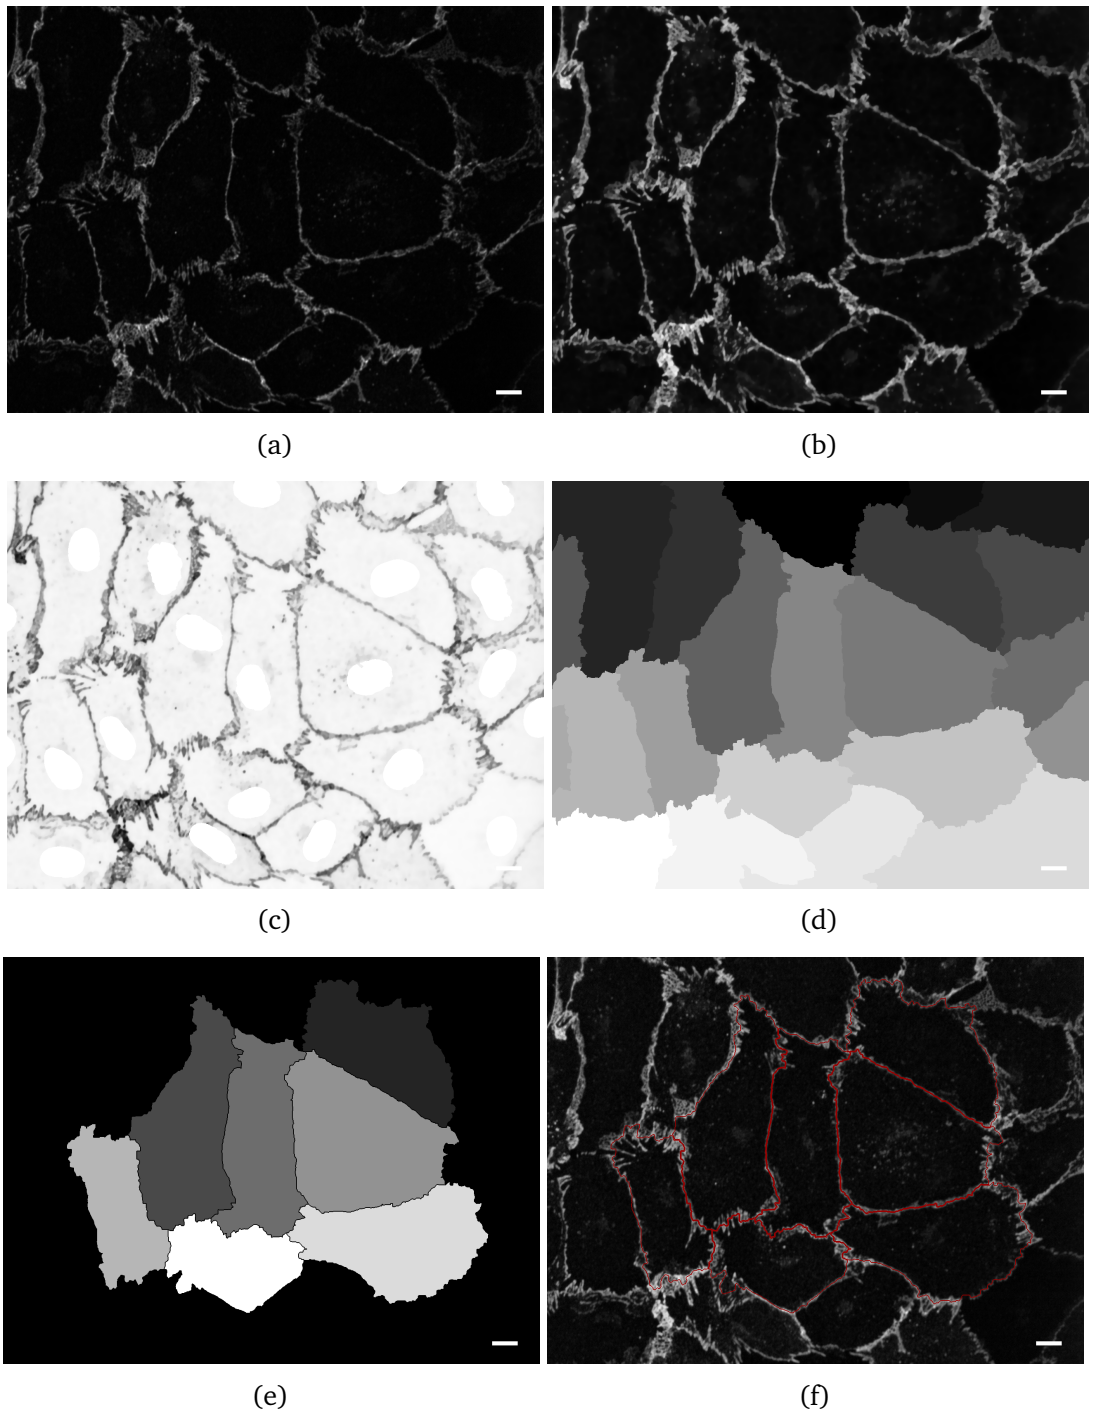

Fig. 2.7 Processes involved in segmentation of cells. Image (a) is the raw acquired image, (b) after applying contrast limited adaptive histogram equalisation and denoising to (a), (c) following inversion and addition of nuclei to (b), (d) depicts the watershed transform applied to (c), and (e) is after clearing borders and relabelling (d), and finally image (f) is a quality control image with segmentation contours. Scale bars: 10.0 µm.

bourhood cumulative distribution function. CLAHE limits the amplification by clipping the histogram at a predefined value before computing the cumulative distribution function. This limits the slope of the cumulative distribution function and therefore of the transformation function.

CLAHE was used in this instance to improve contrast and enhance junctions between cells. The CLAHE technique works well in images where out of focus regions may be lighter or darker than the rest of the image. The kernel size and clipping limit parameters were chosen based on experimentation over multiple images from both available staining techniques.

Noise reduction (B.1)

The plasma membrane stained image contains Poisson noise from the imaging sensor, as well as noise from small punctae within the cell. A noise reduction step was included to reduce this. For this purpose, a bilateral denoising filter and total variation filter were investigated. These algorithms were chosen because they produce posterised images with flat domains whilst maintaining edges. This is appropriate for cell segmentation where there are large homogeneous regions within cells, and well defined edges.

A bilateral filter is an edge-preserving and noise reducing filter. It averages pixels based on their spatial closeness and radiometric similarity. On the other hand, the aim of total variation denoising is to minimise the total variation of the image, which can be roughly described as the integral of the norm of the image gradient.

Total variation denoising with Chambolle's algorithm [70] was found to be more effective for noise reduction in plasma membrane stained images. A weight parameter was chosen based on experimentation over multiple images from several data sets. The outcome of applying CLAHE and a total variation denoising can be seen in Figure 2.7b.

Invert and element-wise minima (B.2 and B.3)

The greyscale image was inverted, and nuclei were superimposed onto the plasma membrane stained image. This step proceeding denoising prepares the image for the watershed transform. The image is inverted so that cell interiors have higher pixel values than junction areas, and the watershed therefore runs from highest to lowest intensity regions using nuclei added as seed points.

Watershed transform and clear borders (B.4 and B.5)

Segmentation of cells by the watershed transform was found to give a computationally efficient and reliable output. The watershed transform used nuclei as seed points. To do this the denoised image is inverted and binary nuclei markers added, as is displayed in Figure 2.7c. The nuclei markers form the seed points for the watershed transform, which is shown in Figure 2.7d. The scikit-image watershed function returns a labelled image (see Figure 2.7e). Connected components on the image border were removed, and the cell relabelled, so background pixels are 0 and then cells are assigned consecutive higher pixel values, as shown in Figure 2.7e.

2.3.5 Weibel-Palade body segmentation

Segmentation of WPBs was achieved with a method similar to that described in the literature and used previously in our research group [5, 71]. The exact method of segmentation has been adapted to the varying demands of different staining protocols and antibody combinations. In earlier trials, non-specific staining of the endoplasmic reticulum caused uneven background illumination in the acquired images. A rolling ball background subtraction [72] pre-processing step improved the signal-to-noise ratio in these images. As both the staining and imaging protocols have improved, the fidelity of the images has increased, and this pre-processing step is therefore unnecessary.

Adaptive threshold (C.0)

The segmentation method needed to accurately find contours of both strongly and weakly fluorescent WPBs, as are shown in Figure 2.8b. Additionally, closely apposed WPBs needed to be identified as separate objects. A global thresholding approach is unable to meet these criteria, so a local adaptive thresholding was applied. In this case, the threshold value is the weighted mean for the local neighborhood of a pixel, subtracted from a user defined constant [73]. An experimentally determined subtraction constant and local neighbourhood or block size were used, and the segmentation verified by a biologist.

Remove small objects (C.1)

Before extracting morphological features, segmented objects beneath the resolution limit of the optical system were removed. Large areas over $10\,\mu m^2$, where multiple WPBs were clumped closely together were also removed. These clusters, such as that displayed in the

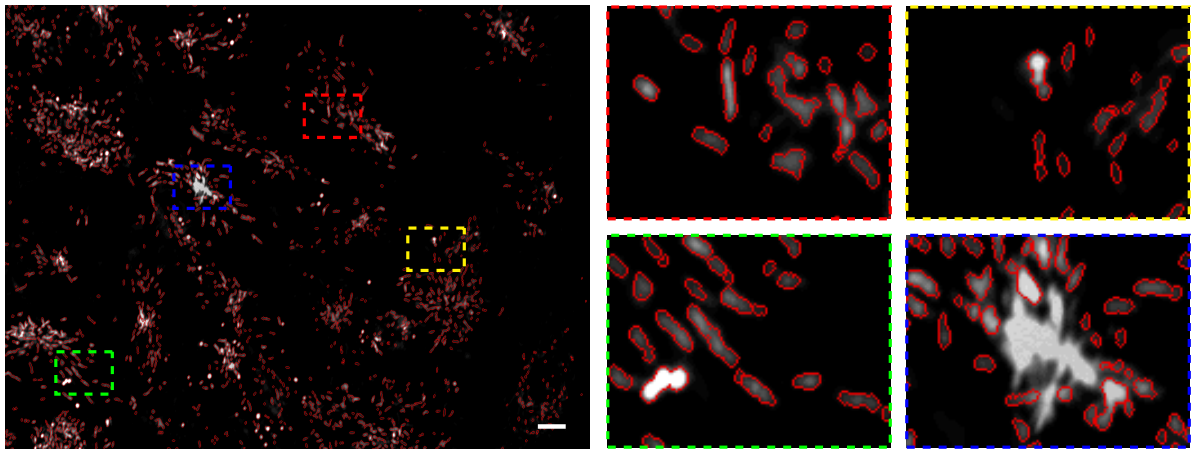

(a) (b)

Fig. 2.8 Fluorescent Weibel-Palade bodies are shown in image (a) with overlaid segmentation contours by adaptive thresholding. Image (b) shows magnified regions of the image, in the red, yellow and green boxes the segmentation of both highly and weakly fluorescent organelles can be seen. The blue box shows the removal of a cluster of WPBs that were inseparable and therefore removed from analysis. Scale bar: 10.0 µm; magnified regions: $20.0 \times 15.0\,\mu m$.

blue box in Figure 2.8b were unfortunately inseparable at the resolution of the optical system.

2.3.6 Exocytic sites segmentation

The WPB cargo vWF is amenable to antibody labelling when exocytosed. Exocytosed vWF is arrested at the cell surface adjacent to the exocytic sites, and appears as round blobs (see Figure 2.9a).

Gaussian blur (D.0)

Image noise was reduced in the vWF stained image using a Gaussian blurring pre-processing step. A sigma value was chosen for the Gaussian blur that does not impact the image resolution. This step was necessary to reduce the sensitivity of the later applied watershed transform. Reducing the pixel variation in the small exocytic sites decreases over splitting by the watershed transform. In segmented vWF exocytic sites, which typically have a small area, the watershed tended to over-split objects. Therefore a smoothing with a Gaussian blur was employed to reduce the sensitivity of the watershed.

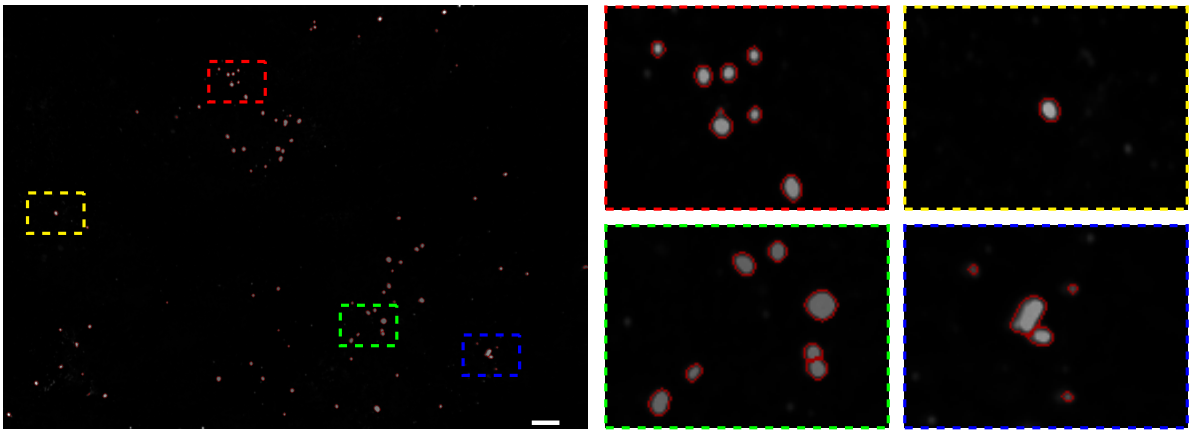

(a) (b)

Fig. 2.9 Image (a) shows segmentation contours overlaid on to an example von-Willebrand factor stained exocytic sites image. Image (b) contains four magnified regions of image (a), where the red and yellow boxes, show a cluster of exocytic sites, and a single exocytic site respectively. The green and blue boxed images demonstrate instances where the watershed transform has split adjacent exocytic sites. Scale bar: 10.0 μm; magnified regions: 20.0×15.0 μm.

Threshold (D.1)

A variety of thresholding techniques were trialled for segmentation of vWF exocytic sites, and by looking over a sample of images from different treatment groups, the most effective method selected. A Moment-preserving threshold [53] was used for this purpose. The low numbers of exocytic sites and uneven spatial distribution created some low contrast acquired images, leading to large variation in threshold value for each image. In some images with few or no exocytic sites a very low threshold would be set and noise segmented, whilst in other images a high threshold would under represent true exocytic sites. To overcome this image by image variation in threshold value, the threshold value was calculated over all unstimulated images in a data set.

Watershed (D.2)

Adjacent exocytic sites segmented as a single object were split using a marker-based watershed flooding algorithm. In each image the *peak_local_max* function was used to return the coordinates of local peaks in an image. This operation dilates the original image and merges neighboring local maxima closer than the size of the dilation. Locations where the original image were equal to the dilated image were returned as local maxima. These local maxima were used as seed points for the marker-based watershed. Examples of

watershed splitting of adjacent exocytic sites can be seen in the green and blue boxes of Figure 2.9b.

Remove small objects (D.3)

The final step in segmentation of vWF exocytic sites was to remove segmented objects beneath the resolution limit of the optical system. The resolution was calculated by Abbe's equation to be five pixels (see Equation 1.2). Morphometric and intensity measurements could then be extracted.

2.3.7 Additional functions

For the specific purpose of morphometric analysis of endothelial organelles, some functionality was required that is not provided natively by Python or any Python libraries. Here additional functions used in image analysis are detailed.

Feret diameter

A feature identified of particular importance that was not available natively in the *regionprops* function was a measure of the maximum caliper diameter (also known as the Feret diameter). This is defined as the distance between the most separated points of the segmented region boundary. This has previously been identified as a descriptive feature for morphometric analyses of WPBs [5], and is as such listed in Table 2.1. A function was written to calculate the Feret diameter of each region: see listing Listing 2.1.

Listing 2.1 Function for calculation of the Feret diameter

```
1 def calculateFeret(coordinates):
2     feret = np.nanmax(squareform(pdist(coordinates)))
3     feret = feret + (((2*((0.5)**2))**(0.5))*2)
4     return feret
```

The function takes a set of coordinates and returns the maximum euclidean distance between those points, where in this case points are pixels. The maximum Feret diameter must be between pixels on the region edge, to reduce the number of operations the region properties function was first used to return a list of pixel coordinates on the boundary of a region. This coordinate list of boundary pixels was then passed to the *calculateFeret* function in Listing 2.1. The distance between each pixel and every other pixel was calculated, and

the maximum distance returned (see line 2). A scaling factor was applied in line 3 to return a sub-pixel resolution, to correspond with Feret diameter measurements in ImageJ [74].

Cell assignment

Analysis of morphometry of endothelial organelles is a valuable outcome from high-throughput microscopic studies of endothelial cells. The assignment of organelles was performed automatically on plates where a plasma membrane channel was present. The assignment of organelles to cells was achieved with a function, as in Listing 2.2.

Listing 2.2 Function to determine in which cell a segmented organelle is located

```
1 def assignCell(label_image, intensity_image, features):
2     properties = measure.regionprops(label_image, intensity_image)
3     cell = pd.Series([prop.max_intensity for prop in properties])
4     features['cell'] = cell
5     cols = features.columns.tolist()
6     cols.insert(4, cols.pop())
7     features = features[cols]
8     features[0] = features[0].map(str) + features['cell'].map("{:03}".format).map(
          str)
9     return features
```

The *assignCell* function takes three inputs: a label image, in which segmented organelles are assigned a unique pixel value or label, an intensity image, which is the *cell_labels* image (shown in Figure 2.7e), and lastly a features table. The features table contains the morphological measurements for each segmented organelle in the label image. In lines 2 and 3 of Listing 2.2 the scikit-image *regionprops* function takes the coordinates of each segmented and labelled organelle, then measures in the *cell_labels* intensity image the maximum pixel intensity of those coordinates. A Pandas [75] series was returned with the pixel intensity for each segmented organelle in the labelled cell image. Cells on the image edge return a pixel value of zero, so can be included or excluded in further analysis.

Lines 4 to 7 of Listing 2.2 inserts the found pixel values as a new column into the morphological features table. Each cell in the data set was given a unique identifier in line 8, which adds the cell number to the particle_id. In data where a plasma membrane stain was present and cell segmentation was performed a particle_id was returned constructed from the row number, column number, field of view number and cell number. Grouping data by these particle_id strings allows for analysis on a cell-by-cell basis. Finally, line 9 of the function returns the features table, which can be appended to the file on disk.

Synthetic coordinates

The coordinate data of segmented organelles can be used to study their spatial distributions, and how different secretagogues and drug treatments may effect the spatial distribution. This was of particular interest when looking at vWF exocytic sites.

There are many ways in which coordinate data can be utilised to investigate spatial distributions. A useful approach was in testing for complete spatial randomness (CSR), this is a process whereby point events occur within an area in a random fashion. Several approaches were trialled to investigate whether organelles were distributed in a random fashion within cells. The quadrant method is a simple approach that subdivides each cell into congruent rectangular sub-cells and the number of points within each cell counted and the distribution analysed. This method has several major drawbacks; it requires quadrants to have equal area, where the partition size impacts the results and there is no obvious way to chose the quadrant size.

A better approach calculates the euclidean distances between points and their nearest neighbors, requiring no artificial partitioning scheme. The Clark-Evans test [76] was then tried as a means to test CSR. In each cell the density of points was calculated and assuming a Poisson distribution used to determine a mean nearest neighbour distance for the cell. For each cell the ratio of the predicted nearest neighbour distance to the real mean nearest neighbour distance was calculated. Values ~ 0 indicate clustered points, ~ 2 are regular ordered points and values ~ 1 are randomly distributed. The Clark-Evans test was however not reliable, since it assumes that cells are infinite or very large. The effects at the region boundary for cells was significant and there was no easy method to compensate for edge effects as with regular shapes.

Finally, an approach for testing the degree of spatial randomness was arrived at using synthetic coordinates. These synthetic coordinates generated artificially could then be used to compare to the real data points. The generation of synthetic coordinates was performed using the Python function shown in Listing 2.3, this takes as input the cell labels image, an organelle features table and a number. The number is preset by the user, to determine how many coordinate pairs to generate, 100 or 1000 pairs were typically used.

Listing 2.3 Function to generate a set of synthetic coordinates for each organelle

```
1 def syntheticCoordinates(cell_labels, features, number):
2     properties = measure.regionprops(cell_labels)
3     coordinates =[prop.coords for prop in properties]
4     coords = np.empty([features.shape[0], 2]); coords.fill(np.nan)
5     synthetic_coordinates = np.zeros([features.shape[0], 2*number])
```

```
6     for i in range(number):
7         for index, row in features.iterrows():
8             cell_number = int(features['cell'][index])
9             if (cell_number > 0):
10                random_coordinates = coordinates[cell_number-1][np.random.randint
                      (0,coordinates[cell_number-1].shape[0],1)]
11                coords[index,0] = random_coordinates[0,1]
12                coords[index,1] = random_coordinates[0,0]
13        synthetic_coordinates[:,i*2] = coords[:,0]*pixel_dimension
14        synthetic_coordinates[:,i*2+1] = coords[:,1]*pixel_dimension
15    return synthetic_coordinates
```

All the coordinates in each cell of the labelled image are retrieved in lines 2 and 3 of Listing 2.3. Two empty arrays are created in lines 4 and 5 to output the coordinate sets into. A nested for-loop for the number of features and number of coordinate pairs to generate then generates random coordinates. Each row in the features table represents a segmented organelle, the cell number is retrieved for that organelle and if it is greater than 0, it is a non-edge cell so a synthetic coordinate is generated. This is done by sampling a random coordinate pair from the list of cell coordinates obtained in line 3. The *syntheticCoodinates* function returns a data table with many synthetic coordinate pairs for each real coordinate pair.

Colour composite image

A useful output from the image processing pipeline was a colour composite image displaying all image channels and segmentation contours. This was achieved by assigning each image channel a colour and overlaying the various colours in a colour composite image.

The Listing 2.4 details the python code written to creating a colour composite image. The function takes as an input a dictionary of the image channels present on the plate and the segmentation contours. Each image channel is assigned a unique colour or combination within the RGB image, in lines 6 to 10. In lines 11 and 12 the image intensity levels are stretched to an intensity range calculated as the 98 percentile.

Listing 2.4 Function to create a colour composite segmentation overlay image

```
1 def create_composite(image_dict, contours):
2     contours = sum(contours_dict.values())*255
3     composite_dim = tiffStack.shape[1:]
4     composite = np.zeros((composite_dim[0], composite_dim[1], 3), dtype=np.float)
5     for key, value in sorted(image_dict.iteritems()):
```

```
6           if key == 'nucleus_image': composite[:,:,0] = value + contours
7           if key == 'cell_image': composite[:,:,1] = value + contours
8           if key == 'wpb_image': composite[:,:,2] = composite[:,:,2] + value/2+
                contours; composite[:,:,1] = composite[:,:,1] + value/2 +contours
9           if key == 'golgi_image': composite[:,:,2] = value+contours
10          if key == 'exitSites_image': composite[:,:,0] = value/2+contours;
                composite[:,:,1] = value/2+contours;
11      p2, p98 = np.percentile(composite, (2, 98))
12      composite_rescale = exposure.rescale_intensity(composite, in_range=(p2, p98))
13      return bytescale(composite_rescale)
```

An example colour image output created using the *create_composite* function in Listing 2.4 is shown in Figure 2.10. This is an especially useful means for the biologist to visually check their results and any abnormal cells or segmentation results, so is generated for each image in a data set.

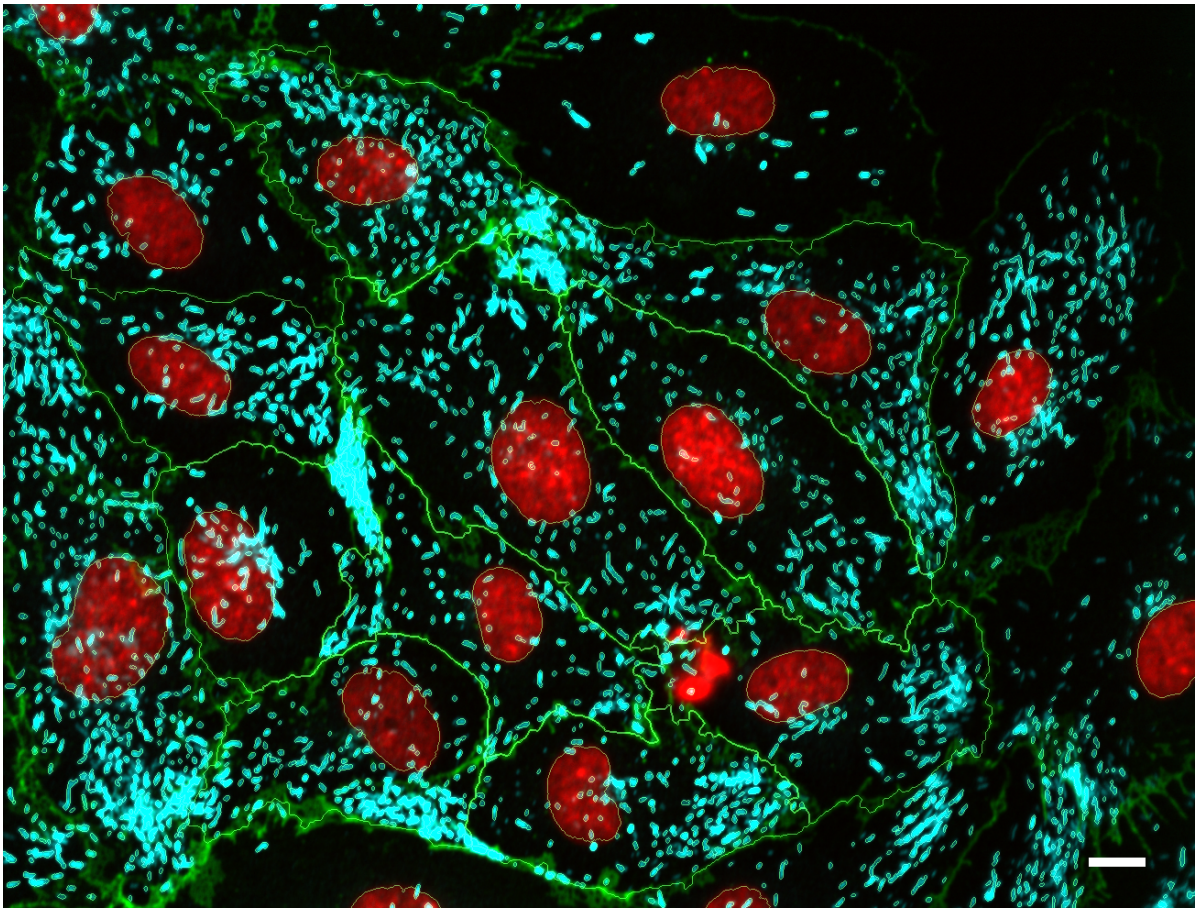

Fig. 2.10 An example colour composite image, including the following channels: red channel for cell nuclei, green for plasma membrane staining, and cyan colour for vWF staining. Segmentation contours are overlaid for all segmented objects in the image, except segmentation contours of cells on the edge of the image have been excluded. Edge cells may not be completely within the field of view so are not used in data analysis of cellular populations. Scale bar: 10.0 μm.

2.4 Performance evaluation

To evaluate the segmentation performance of nuclei, cell boundaries, vWF exocytic sites and WPBs in endothelial cells, methods and metrics have been chosen based on the type of image data and how the data was used. Segmentation performance may be evaluated in a binary or a continuous manner. For example determining whether or not a particle was detected is a binary problem, whilst the accuracy in dividing spatial regions is a continuous problem. Spatial segmentation can be evaluated by calculating the ratio of overlap between a segmentation result with a gold standard segmentation.

A secondary validation of segmentation methods is established by the consistency of biological results achieved. For example, the consistency with other methods of experimental methods such as low-throughput microscopy or enzyme-linked immunosorbent assays (ELISA) may be assessed.

2.4.1 Nucleus segmentation evaluation

The segmentation of nuclei is crucial to accurate cell segmentation, where each nucleus is used as a seed point defining a cell. The presence or absence of a nucleus is therefore more important than the accuracy of contours of the segmented object, as such the performance of nucleus segmentation has been evaluated as a binary classification problem. A cursory visual inspection of segmented nuclei images indicates that the contours are generally accurate to the nucleus edge.

A basic thresholding approach to nucleus segmentation achieved a reasonable nucleus segmentation accuracy, suitable for the majority of nuclei. However, this segmentation method tended to be inappropriate in certain instances, for example: where cells were undergoing division, at removing debris or other fluorescent non-nucleus objects, segmentation of cells on the image edge, and segmentation of dead cells. The frequency of these occurrences depended on the condition of the cells undergoing analysis. Generally these faults were a small fraction of the total segmented nuclei objects. The consequences of incorrectly segmented nuclei on cell segmentation was large enough to warrant additional steps to improve segmentation. The SVM classifier and watershed transform were used to reduce the number of such errors, as described in subsection 2.3.3.

A testing set of 5804 labelled gold standard segmented objects from nuclei images were used to asses the performance in a binary sense of nuclei classification. The optimal C and γ values (established in subsection 2.3.3) were used to perform SVM classification on the 5804 example test set, and the predictive ability evaluated. The predictive ability was

evaluated by establishing the number of true positives, false positives, true negatives, and false negatives between the gold standard labelled data and the automated analysis, as is shown in Table 2.2.

Table 2.2 Confusion matrix of support vector machine classifier predictive ability applied to 5804 sample testing set, when using a radial basis function kernel and optimal C value of 8.0, and a γ value of 0.03125.

		True	
		p	n
Predicted	p'	1996	22
	n'	145	3641

In this case a true positive is a nucleus that is correctly predicted by the SVM model to be a nucleus. Conversely a true negative is a non-nucleus object that has been correctly classified as a non-nucleus object. False positives are non-nucleus objects falsely identified as nuclei. Finally false negatives are nuclei falsely classified as non-nucleus objects.

An automated correspondence and matching step written as a Python script was used to label each object in the automated segmentation and gold standard segmentation. Using the centroid coordinates of each automatically segmented object, the pixel value at the corresponding location in the gold standard image was measured. If this value was one then the object was a true positive and if zero then it was a false positive. All matched objects were removed from the gold standard and the remaining objects were labelled as false negatives.

The values in Table 2.2 were used to calculate the precision metrics in Table 2.3.

Table 2.3 Performance metrics of support vector machine nuclei classifier with optimal C and γ values.

	Nuleus	Non-nucleus	Average
Precision	0.99	0.93	0.97
Recall	0.96	0.99	0.97
F1 score	0.98	0.96	0.97

2.4.2 Cell segmentation evaluation

Evaluation of the performance of cell segmentation considers the ratio of overlap between cell contours in a hand-labelled gold standard, compared to an automated segmentation

method. Accuracy of cell segmentation impacts the assignment of organelles to their respective cell, and therefore is essential to ensure valid interpretation of data when performing analysis at a single cell level.

In analysis of cellular populations of endothelial organelles there are two principal sources of error. The first is the errors in detection of nuclei, which are used as seed points for cells, as described in subsection 2.4.1. The second is inaccuracy in cell boundary segmentation contours achieved thorough automated image segmentation. The two errors are compounded such that the total error in cell segmentation is a combination of the error in nucleus detection and the error in cell segmentation contours.

The degree of error in segmentation of cell boundary contours was evaluated by comparison of the similarity between gold standard cell segmentation and automatically segmented cell boundaries. Two overlap ratio coefficients were used, namely the Dice coefficient [77], and Jaccard index [78]. These metrics measure the similarity and diversity between two finite data sets respectively, where in this case the sets are sets of pixels. The Jaccard index is defined as the size of the intersection divided by the size of the union of the two sets,

$$J(A,B) = \frac{|A \cup B|}{|A \cap B|}, \tag{2.2}$$

where A and B are two sets. Correspondingly, the Dice coefficient for two sample sets is given by,

$$D(A,B) = 2\frac{|A \cap B|}{|A| + |B|}. \tag{2.3}$$

To evaluate segmentation performance a data set was prepared containing data where two separate plasma membrane staining methods had been applied. Wheat germ agglutinin (WGA), henceforth referred to as staining I, provided a general cell membrane stain. An alternative second antibody based staining method was provided by VE-cadherin, henceforth referred to as staining II, which specifically stains cell-cell junctions. The WGA staining method is shown in Figure 2.11a and the VE-cadherin antibody based method in Figure 2.11b. These staining methods produced images with different characteristics and unique challenges for image segmentation. A generally applicable segmentation protocol was required that would be effective on both staining modalities.

A sample of 517 cells were hand annotated, demarcating cell boundaries, to give a labelled gold standard set for each cell in the data set. An example hand segmented cell is shown in Figure 2.12a. Segmentation contours were then generated from various image processing pipelines and segmentation algorithms. For each cell, the Dice and Jaccard metrics were calculated through comparison of the automatic segmentation contours to

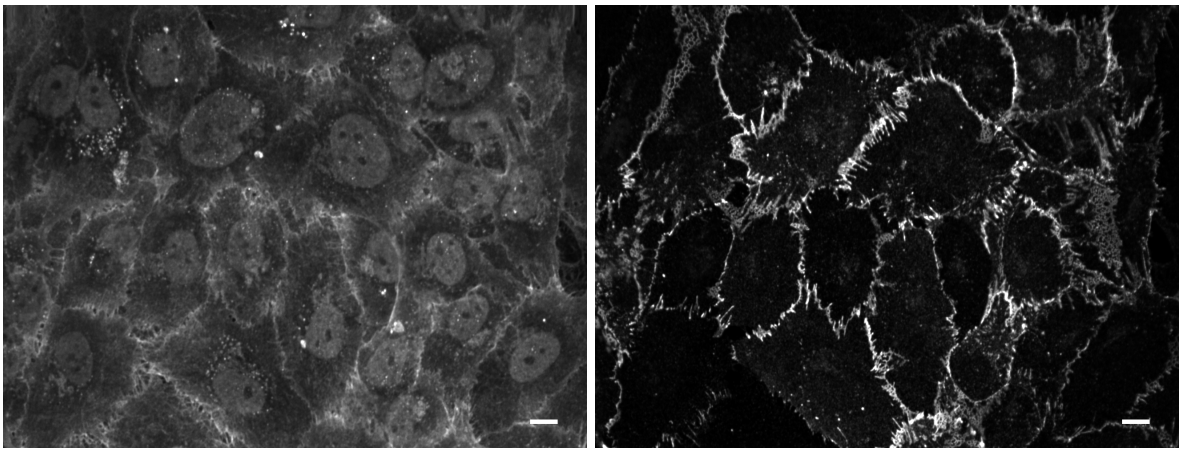

(a) (b)

Fig. 2.11 Images of plasma membrane staining acquired from two different staining protocols. Image (a) shows a wheat germ agglutinin (WGA) cell membrane stain (staining I), and image (b) shows a VE-cadherin cell-cell junction stain (staining II). Scale bars: 10.0 μm.

the hand annotated gold standard. An example of automated segmentation contours of a cell is shown in Figure 2.12b. The union and intersection of these two sets could then be used in the calculation of the Dice and Jaccard metrics. The union and intersection are displayed in Figure 2.12c.

The Dice coefficient and Jaccard index were calculated in this way for all cells in the 517 cell data set, and the mean and error on the mean reported (see Figure 2.13).

Many approaches were trialled for cell segmentation, as discussed in subsection 2.3.4. After a cursory viewing of the segmentation contours some methods were soon abandoned, as they were found to be inappropriate for the specific segmentation challenge. Methods that showed promise were further analysed.

An initial method to estimate cell boundaries was to apply a Voronoi decomposition of the image, using nuclei as seed points. This method does not require staining of the plasma membrane, but instead the Voronoi diagram tessellates an image using the nuclei centroids as seed points. Cells are constructed from equidistant lines between nuclei. In this way a Voronoi performs an estimation of cell boundaries. A more precise segmentation could be achieved with staining and imaging the cell plasma membrane. The mean Dice coefficient and Jaccard index for segmentation approaches are displayed in Figure 2.13. A host of pre-processing and segmentation methods were trialled for cell segmentation, as discussed in subsection 2.3.3. Many of these methods were found to be inappropriate even with a cursory look at the resulting segmentation contours. The two most promising methods

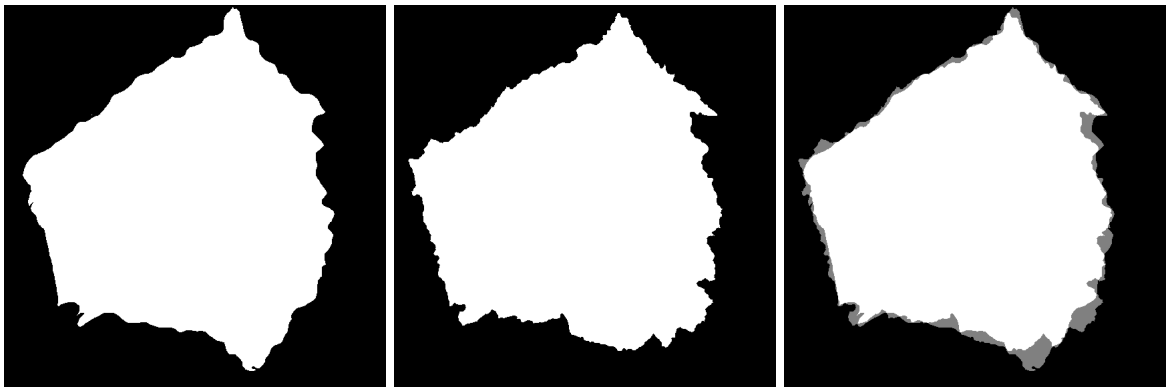

(a) (b) (c)

Fig. 2.12 An example cell manually segmented to give a gold standard in (a), and the corresponding contours for the cell generated by an automated segmentation method in (b). Image (c) shows the union and intersection of the cell, where the intersection are the white foreground pixels and the union are the grey pixels in the background.

identified were the evolving generalised Voronoi diagram (EGVD) [69] and a marker based watershed [54].

A script written in Python calculated the Dice and Jaccard metrics of each cell in the data set, from which the mean and the error on the mean could be calculated for each segmentation approach. For each cell within the 517 cell data set, the associated hand labelled and automated segmentation contours were identified and the Dice coefficient and Jaccard index calculated. The Voronoi and watershed methods were performed natively in Python, whilst segmentation contours for the EGVD method were generated in the StemCell3D software[3].

The best performing method of cell segmentation (Watershed II) used a contrast limited adaptive histogram equalisation (CLAHE), followed by a denoising step and finally a watershed transform. The performance of this method was consistent across antibody I and antibody II. Generally methods involving a plasma membrane stain performed better than the Voronoi. Both the watershed and EGVD approaches achieved good results, with the Watershed II method giving the best results. It has the additional benefit of using functions all native to Python and being computationally more efficient than the EGVD.

[3]W. Yu, H. Srivats, S. Shvetha, S. C. Chia, V. Pascal, A. Sohail and H.K. Lee, StemCell3D software, http://imaging.imb.a-star.edu.sg/17/stemcell3d.html, accessed 2015-11-28

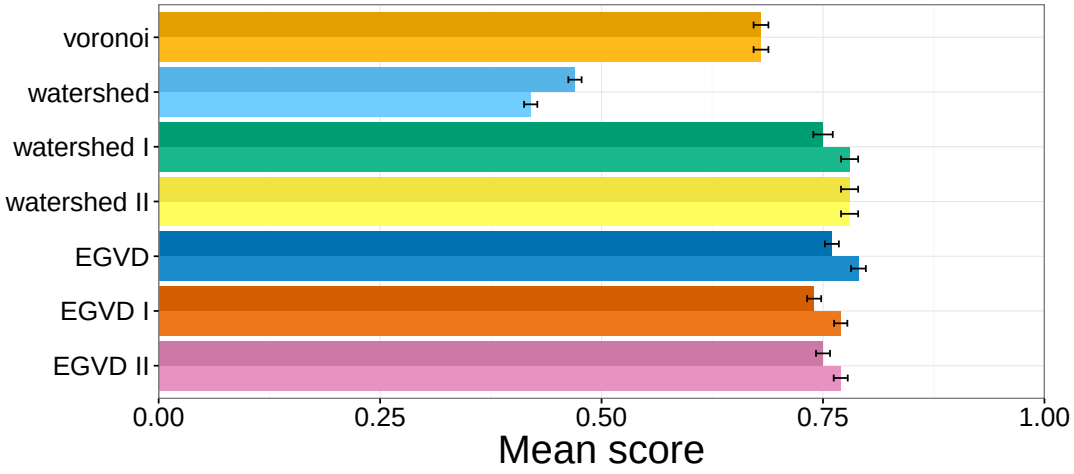

(a)

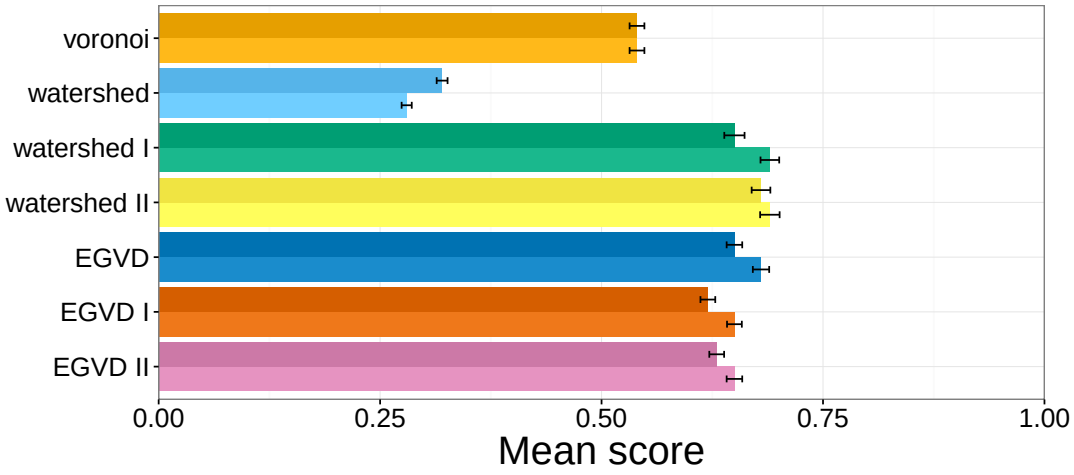

(b)

Fig. 2.13 Chart displaying the mean Dice coefficient (a) and Jaccard index (b), calculated with reference to a gold standard hand labelled set comprised of 517 cells. Three principle segmentation algorithms were evaluated: the Voronoi, evolving generalized voronoi diagram (EGVD), and watershed. These were applied in combination with one or both of contrast limited adaptive histogram equalisation (CLAHE), and Chambolle's denoising. The watershed I and EGVD I methods use only CLAHE, whilst watershed II and EGVD II use CLAHE and denoising. Methods were evaluated over two staining methods, staining I (wheat germ agglutinin), the darker colour upper bar in each pair, and staining II (VE-cadherin), the lighter colour lower bar in each pair. Error bars represent the standard error on the mean.

2.4.3 Weibel-Palade body segmentation evaluation

The two principle data uses from WPB segmentation were, first, counting numbers of WPBs, and measuring the Feret diameter of WPBs. The accuracy in detection of WPBs is important to accurately determine the number of WPBs. While the accuracy of WPB segmentation contours is important in the measurement of Feret diameter and other morphometric features. The frequency and density of WPBs imaged with our optical setup makes establishing a gold standard untenable, especially in areas of clustered WPBs, and for WPBs with staining variation along their length. As such, WPB segmentation has not been evaluated with reference to a gold standard set.

Since the WPB segmentation model described in subsection 2.3.5 was built upon legacy methods used and previously published by our research group [5, 71], an explicit evaluation of the WPB segmentation is not provided here. The methods have however been previously verified through comparison of results at high-throughput with data obtained in low-throughput, but with higher-magnification and higher-resolution optical systems. The numbers of WPBs have also been corroborated with levels of vWF expression measured via alternative methods such as ELISA.

2.4.4 Exocytic sites segmentation evaluation

The exocytic vWF sites labelled at the cell surface typically appear as round fluorescent blobs. Segmentation of these objects was achieved by applying a Gaussian blurring, a moment-preserving threshold [53] and a watershed transform, as described in subsection 2.3.6.

Analysis and interpretation of vWF exocytic site data has been focused primarily on their numbers at a population and a single cell level, with a secondary motivation to study their areas. Evaluation of the segmentation performance has therefore mainly concentrated on the ability of the segmentation method to correctly identify the exocytic sites. This performance evaluation does not supply a means to evaluate the accuracy of the area of measurement, because the sites are generally very small (diameter between $0.5\,\mu m$ to $3.0\,\mu m$ [79]) accurately demarcating the area at each site would be inaccurate and laborious.

The segmentation routine's ability to identify exocytic sites was confirmed over a sample of nine images. These nine images were obtained from three separate cell treatment groups, commonly used in experimentation. The treatment groups differ by levels of cell activation and therefore exocytic site frequency. A good segmentation performance in all groups is

necessary. The segmentation was evaluated across these treatment groups and combined to give a total or average score.

The image set contained a random sample of three images of unstimulated cells; three random images from histamine stimulated cells and three randomly selected images from PMA stimulated cells. The exocytic sites in these nine images were manually annotated to form a gold standard comprised of 720 vWF exocytic sites.

The automated segmentation method described in subsection 2.3.6 was applied to each of the nine images. Based on a comparison with the gold standard, manually labelled images foreground objects and automatically segmented objects were categorised as either *true positives*, *false positives* or *false negatives*. A particle in the automatically segmented image was labelled as a *true positive* if any of its pixels overlapped with a foreground pixel in the hand labelled image. Conversely, if the pixels of a segmented object in the automatically segmented image were all background pixels in the hand labelled image, this was deemed a *false positive*. Finally, *false negatives* were hand labelled objects in the gold standard image that did not correspond to any foreground pixels in the automatically segmented image.

For each of the nine images, the number of *true positives*, *false positives* and *false negatives* were counted, and three metrics were used to evaluate the performance, as shown in Table 2.4. The precision was calculated as the fraction of retrieved instances that are relevant, or the ratio of the sum of *true positives* to the number of *true positives* and *false positives*. The recall was calculated as the fraction of relevant instances that were retrieved, or the ratio of *true positives* to the sum of *true positives* and *false negatives*. The balanced F1 score combines the harmonic mean of the precision and recall to give a general evaluation metric.

Table 2.4 Evaluation of exocytic sites segmentation protocol over three treatment groups and nine images, with a hand-labelled gold standard of 720 exocytic sites.

Metric	Unstimulated			Histamine			PMA			**Total**
	1	2	3	1	2	3	1	2	3	
True positive	7	4	9	56	72	97	179	154	105	**683**
False positive	1	0	1	32	8	8	4	4	39	**97**
False negative	0	0	0	1	2	1	13	15	5	**37**
Precision	0.88	1.00	0.90	0.64	0.90	0.92	0.98	0.97	0.73	**0.88**
Recall	1.00	1.00	1.00	0.98	0.97	0.99	0.93	0.91	0.95	**0.95**
F1	0.93	1.00	0.95	0.77	0.94	0.96	0.95	0.94	0.83	**0.91**

Of the 720 hand labelled exocytic sites across the nine images, 683 were correctly identified by the segmentation method giving a recall or accuracy of 0.95. The higher number of *false positives* reduced the precision score to 0.88, and yielded an overall F1 score of 0.91.

2.5 Results

Briefly presented here are results from a high-throughput study of vWF exocytic sites, included to give an indication of the experimental scale, and methods of data visualisation. This experiment included four groups of HUVECs: an unstimulated control group, a group stimulated with 100 ng ml^{-1} thrombin for 10 min, a group stimulated with phorbol 12-myristate 13-acetate (PMA) 100 ng ml^{-1} for 10 min, and a group stimulated with 100 ng ml^{-1} histamine for 10 min. Samples were immunostained for internal vWF and VE-cadherin. Following washing, samples were incubated with fluorescently-conjugated secondary antibodies to stain for exocytosed vWF, VE-cadherin, and Hoescht to visualise the nucleus. Cells were then imaged and image data were pre-processed to remove any unusable images, following this the image processing workflow was run to segment organelles and extract morphometric features.

In this data set 13110 exocytic sites were segmented; 108 from the control cell group, 3639 from thrombin stimulated group, 3068 from the histamine stimulated group, and 6295 from the PMA stimulated group. The results shown in Figure 2.14a reflect the relative exocytic strengths of the secretagogues thrombin, histamine and PMA. The mean number of exocytic sites per cell is is a more useful indicator than the total number of exocytic sites for each treatment group because it is relative to the number of cells in each group, which may vary due to the cell confluency. An additional result that has led to further insight is the difference in exocytic site area shown in Figure 2.14b, potentially suggesting different vWF exocytic mechanisms between secretagogues.

Single cell analysis also facilitates visualisation of distributions of exocytic site numbers per cell as shown in Figure 2.14c. These normalised data could indicate a difference in cellular response between secretagogue treatments. Histamine tends to affect more cells but have a reduced secretory response, whereas thrombin and PMA effect fewer cells but cause a larger exocytic response. This is demonstrated by the increase in the number of cells with five or more exocytic sites in the thrombin and PMA groups as compared to the histamine group.

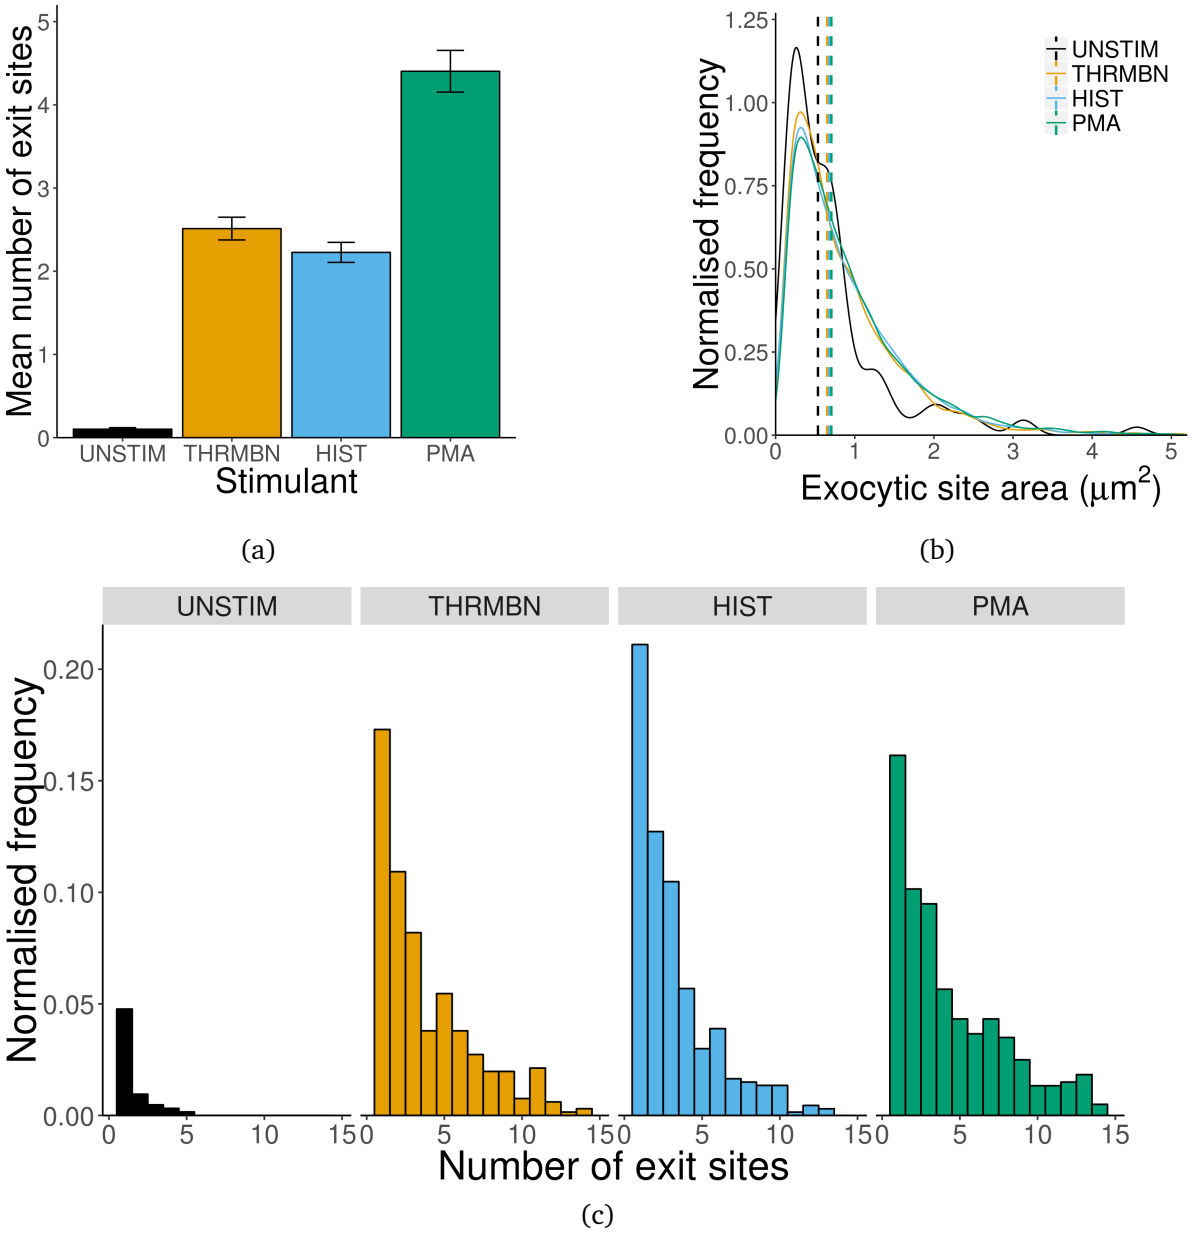

Fig. 2.14 Exocytic sites example data from four cell groups: unistmulated, thrombin, histamine stimulated and phorbol 12-myristate 13-acetate (PMA) stimulated. Plot (a) is a bar chart showing the mean number of exocytic sites per cell in the different endothelial cell groups, where error bars are calculated as the standard error on the mean. Plot (b) is a density plot to show the distribution of exocytic site areas amongst the four groups, with the median exocytic site area plotted as a dashed line. Plot (c) is a normalised distribution of the number of exocytic sites per cell for each of the four treatment groups.

2.6 Cell classification study

The experimental, image processing, and data analysis methods were assessed by investigating how well two cell phenotypes could be created and identified in cell populations. The two cell types used were: small interfering RNA (siRNA) knockdown cells and control endothelial cells. Endothelial cells were obtained from pooled donors, and naturally have variability in both the numbers of WPBs present and the morphometric attributes of those WPBs. The two cell types, control and siRNA knockdown, will hereby be referred to as cell type 0 and cell type 1. Where, cell type 0 is an untreated control cell type, whilst cell type 1 are cells which have been treated with siRNA. The siRNA mediates the reduction of vWF protein levels, reducing the number of WPBs and decreasing the length of residual WPBs. The quantum size, as measured by the distance between length clusters, remains constant upon progressive reduction of vWF cell content [5]. In general it is therefore expected that the WPBs in cell type 1 should be shorter and fewer than those in cell type 0.

A labelled population of siRNA knockdown cells, and a labelled population of control cells, were used to train supervised learning prediction models. Three machine learning routines were applied in the R programming language [80] version 3.2 to classify cells as cell type 0 or 1. The accuracy of the classification prediction could be assessed by comparing the ratio of cell types to known volumetric ratios of the cell types. Classification of cell phenotypes in this study relies on accurate segmentation of WPB contours and cell boundary contours. Segmentation was performed on images acquired with staining of nuclei, WPBs, and a plasma membrane boundary marker. The segmentation of these images is described in subsection 2.3.3, subsection 2.3.4 and subsection 2.3.5. For each cell, a set of features was measured and aggregated, describing the morphometric characteristics of the cell nucleus, cell plasma membrane, and WPBs within the cell.

2.6.1 Data pre-processing

Information about the microtitre plate layout and data structure in this cell classification study is given in Table 2.5. There were five groups in this experiment, with a proportion of the data in group A of known cell type 0, and another part in group E of the data of known cell type 1. The remaining three data groups were comprised of varying proportions of cell types 0 and 1, where volumetric ratios were used to control the ratio of each cell type.

Morphometric features extracted from segmented cells of known cell types in treatment groups A and E, were used as a labelled supervised machine learning training data set. Treatments A and E were split into data for training, and testing the model. Three machine

Table 2.5 The experimental groups, consisting of groups mixed ratios of cell types is used in a cell classification experiment. Columns refer to those on the microtitre plate, and data type refers to those for machine learning.

Group	Columns	Cell type		Data type	
		0	1	Training	Testing
A	1-2	1.00	0.00	✓	✓
B	3-4	0.75	0.25	✗	✓
C	5-6	0.50	0.50	✗	✓
D	7-8	0.25	0.75	✗	✓
E	9-10	0.00	1.00	✓	✓

learning algorithms were then applied to build predictive models with the training data set. The models were applied to classify each cell in the testing sets A, B, C, D, and E as type 0 or type 1 cells, based on their extracted features. The three supervised machine learning approaches trialled were support vector machines (SVM), random decision forests, and classification and regression trees.

Further to this, a synthetic data set was constructed using morphometric features from cells randomly selected from groups A and E, split into five groups representing the ratios given in Table 2.5. The predictive ability of the classifier on the synthetic data is instructive as to the error on the real data. Since there is potential error introduced by mixing volumes of two cell types, the difference between the synthetic data and the real data gives an indication of the accuracy mixing.

Feature selection and scaling Classifier prediction of endothelial cell type requires a set of cellular features on which to build a model and predict with. Each cell has a wealth of data that can be exploited to form cellular features, including pixel intensity and morphometric data about the cell, WPB morphometry and pixel intensity, and nuclei pixel intensity and morphometric data. The cellular features were created by merging WPB morphometric data and cell morphometric data, by the *particle_id*, and then aggregating the results. A set of 10 numeric features were created, pertaining to the cell and its constituent WPBs. The degree of correlation between these features is shown by the size of circles in Figure 2.15a, the colour also indicates whether negatively or positively correlated.

To standardise the range of independent variables in the data, and to improve the result of the machine learning algorithms, the data were normalized. This ensured that each feature contributed approximately proportionally to the final outcome. All features were

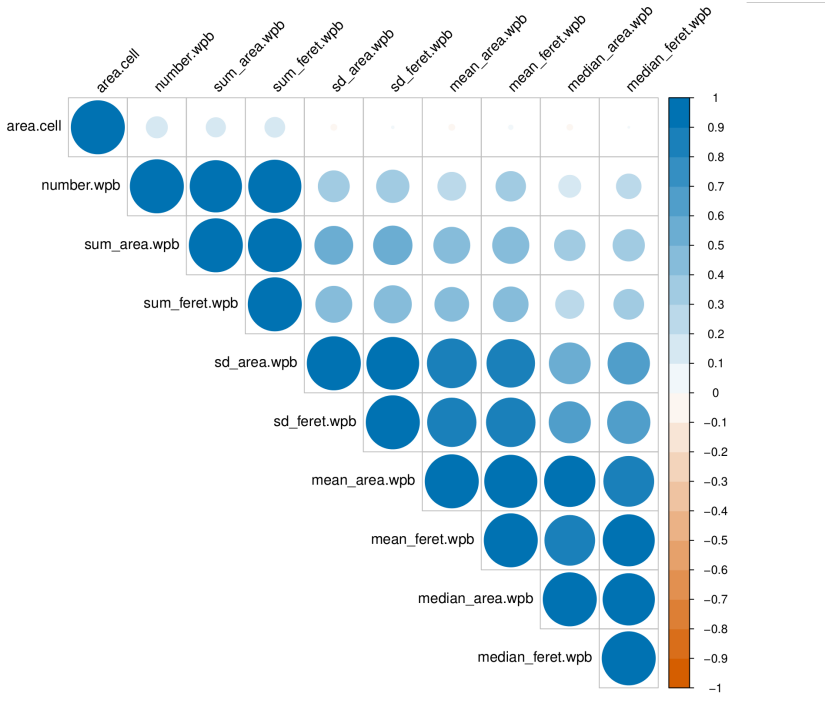

(a)

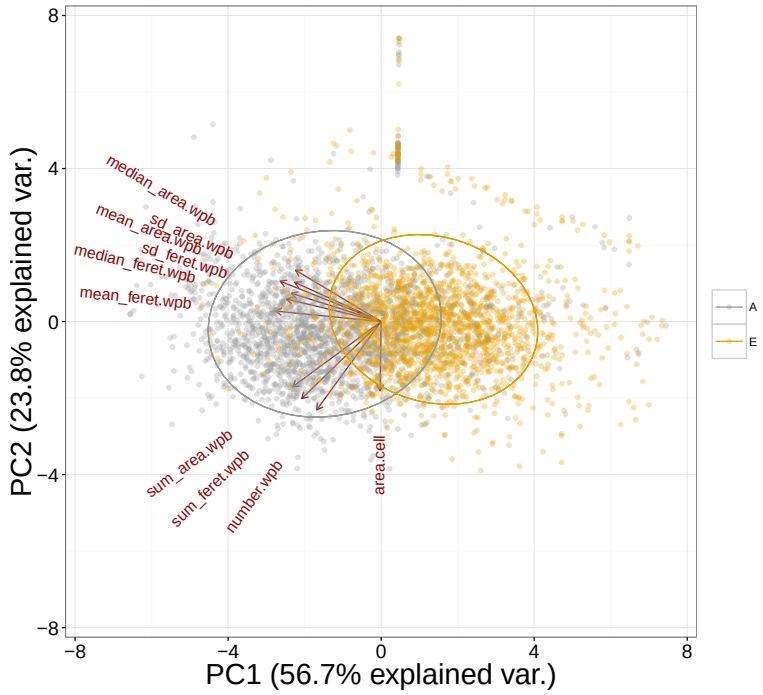

(b)

Fig. 2.15 Plot (a) is a correlation matrix showing the degree of correlation between cellular features, where large circles are more highly correlated, blue colours are positively correlated and red colours negative. The biplot in (b) plots the two greatest variance principle components (PC1 and PC2) of groups A and E. A normal contour line ellipse is drawn for showing the 68 % probability for each group.

scaled by a linear transformation on the original values, to prevent attributes in greater numeric ranges dominating those in smaller numeric ranges. Features were scaled in the range 0 to 1 by the formula given as:

$$X' = \frac{X - X_{min}}{X_{max} - X_{min}}, \tag{2.4}$$

where X is the original features set, X' the scaled feature set, X_{min} is the minimum value within the X set, and likewise X_{max} is the maximum value within X. Normalization by Equation 2.4 was applied separately to each feature set.

Principal component analysis The aim of principal component analysis (PCA) is to reduce the dimensionality of data by combining features into principal components, where principal components are linear combinations of correlated or uncorrelated features. This can improve algorithm performance and allows high dimensional data to be visualised in lower dimensional spaces. The PCA method is one of two commonly used methods to determine effective linear combinations of features, the other being multiple discriminant analysis (MDA) [81]. In PCA an orthogonal transformation ensures that components are linearly uncorrelated, and principal component orthogonality is ensured because components are eigenvectors of the symmetric covariance matrix.

The number of principal components generated by PCA is always less than or equal to the number of original features. Principal components are ordered by their variance, where the first principle component accounts for the greatest variance, and each successive component accounts for less.

A PCA was performed on the cellular features of labelled cell types 0 and 1, where features were first log transformed. Variables were log transformed to reduce the scale of each feature in order to make the distributions more normally distributed and to stabilise the variances. In this case a log transform was used, since all othe features were positive real numbers. The underlying structure of the data could then be explored by performing a PCA and plotting a biplot.

Figure 2.15 shows a biplot of the first two principle components, which together account for 80.6 % of the variance in the data. The two labelled data groups A and E in Table 2.5 were plotted representing cell types 0 and 1 respectively. The overlapping clusters in the biplot, indicate that our cell types may not be easily separable, which would suggest that the predictive ability of the machine learning classification may be limited. The vertical strip of points positively offset from the x-axis are from cells that do not contain any WPBs;

hence the variability in the direction of the *area.cell* vector but not in any vectors associated with WPB morphology characteristics.

Cross-validation A k-fold cross-validation technique was applied prior to each machine learning approach to gauge how well the predictive model performs. This step was included to limit problems like over-fitting the data and give insight into how the predictive model generalises to independent data sets.

In k-fold cross-validation the data is randomly split into k equally sized subsamples. Of the k subsamples, a single subsample is retained as the testing data for testing the model, and the remaining $k-1$ subsamples are used as training data. The cross-validation process is then repeated k times (the folds), with each of the k subsamples used exactly once as the testing data. The k results from the folds can then be averaged to produce a single heuristic estimation of the performance. In this way all observations are used for training once and testing at least once.

To perform cross-validation of the labelled groups A and E (Table 2.5), and of the 3842 cells a sample of 1200, approximately a third of the total labelled data was extracted for 5-fold cross-validation. The extracted data were shuffled and partitioned into five complimentary subsets, and a predictive model was built on each data subset. In each case, the remaining data that were not used to generate the model were used to test the performance of the model. A 5-fold cross-validation was used, such that five predictive models were built and each one's performance tested on different data partitions, reducing the variability.

2.6.2 Machine learning methods

Three commonly used machine learning techniques were trialled for cell classification, one based on hyperplane separation, one on decision trees, and one utilising an ensemble of decision trees. A brief introduction is given to each technique, along with the results of 5-fold cross-validation.

Support vector machine

The first supervised machine learning algorithm trialled for classification of cell types in this experiment was a support vector machine (SVM). A SVM is a discriminative classifier, formally defined by a separating hyperplane. The algorithm uses labelled training data to output an optimal hyperplane which categorises unseen examples. In essence, SVMs

look for the optimal separating hyperplane between two classes, by maximising the margin between the classes [65]. The goal in training a SVM is to find the separating hyperplane with the largest margin, generally, the larger the margin the better the generalisation of the classifier.

SVM requires two parameters, the cost and gamma functions. The cost function controls the cost of misclassification on the training data. A small cost value reduces the number of ignored data points and leads to a soft margin, while a large cost value is stricter, but can potentially overfit the data. The gamma function is the parameter of the Gaussian radial basis function.

To perform SVM classification of the data, the known cell types in treatment A and E were appropriately labelled and combined to from a table for cross-validation. 5-fold cross-validation was performed after shuffling the data and partitioning into five groups, each of which was used exactly once for training. SVM was performed in the R programming language [80] version 3.2, with parameter optimisation and SVM functions provided by the e1071 package [82]. For each data partition C and γ functions were optimised by a grid search over the parameter 2^{-6} to 2^{6} for γ and 2^{0} to 2^{10} for the cost function. A model was trained using the optimal parameters and a radial basis function kernel. For the SVM method optimal C and γ parameter values were 32, and 0.0625, respectively.

5-fold cross-validation The 5-fold cross-validation result of applying the SVM with a radial basis function kernel and optimised C and γ parameters by a grid search is given by the confusion matrix in Table 2.6.

Table 2.6 Confusion matrix of support vector machine 5-fold cross-validation on cell classification.

		True	
		p	n
Predicted	p'	1435	233
	n'	417	1759

The cross-validation gave a sensitivity of 0.78 and a specificity of 0.88. The higher specificity or true negative rate suggests that the SVM classifier is more successful at classifying type 0 cells than type 1. The overall accuracy calculated as the average of the specificity and the sensitivity was 0.83 for the SVM method

Classification and regression trees

A second machine learning method used to construct a prediction model was using classification and regression trees (CART). These are non-parametric decision tree learning techniques that produce either classification or regression trees, depending on whether the dependent variable is categorical or numeric, respectively. The goal of decision tree methods is to create a model that predicts the value of a target variable by learning simple decision rules inferred from the data features [83].

Algorithms for constructing decision trees usually work from the top down, by choosing a variable at each step that best splits the feature set. CART are built by first finding the single feature that best splits the data into two groups. The data is separated and the process applied separately to each subgroup, until no improvement can be made or the subgroups reach a minimum size. The second stage of the procedure is to prune the tree back using cross-validation techniques.

At each stage in a decision tree different algorithms have different metrics for determining the optimal split [84]. Generally the splitting criterion is determined by the impurity function, measuring the homogeneity of the target variable within the data subset. The CART algorithm uses the Gini impurity as a measure of how often a randomly chosen element from the set would be incorrectly labelled, if it was randomly labeled according to the distribution of labels in the subset. Gini impurity can be computed by summing the probability of each item being chosen multiplied by the probability of a mistake in categorising that item.

CART based methods were implemented in R version 3.2 using the recursive partitioning and regression trees (rpart) package [85]. The CART model was built in rpart using a two stage procedure of tree building and pruning. The resulting models are represented as binary trees [81].

In this instance the classification and regression tree (shown in Figure 2.16) selected the standard deviation of Feret diameter of WPBs within the cell as the most important splitting criteria.

5-fold cross-validation The 5-fold cross-validation result of applying the CART algorithm is given in the confusion matrix in Table 2.7. The 5-fold cross-validation technique confusion matrix for the CART method is as follows:

The cross-validation gave a sensitivity of 0.77 and a specificity of 0.85. The higher specificity or true negative rate suggests that the CART classifier is more successful at

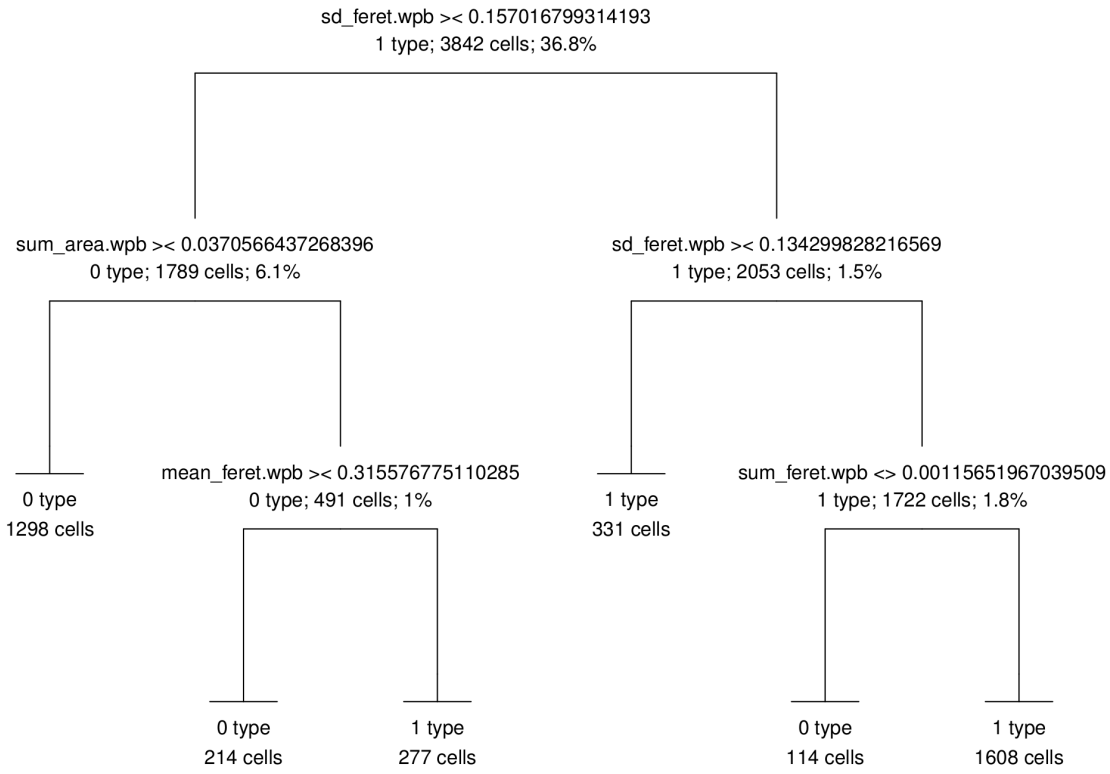

Fig. 2.16 Hierarchical classification tree for cell classification. Each node gives the splitting criterion and value, where <> means that cases with lower values go left and >< mean cases with lower values go right.

Table 2.7 Confusion matrix of classification and regression tree method 5-fold cross-validation on cell classification.

		True	
		p	n
Predicted	p'	1438	296
	n'	414	1696

classifying type 0 cells than type 1. The overall accuracy calculated as the average of the specificity and the sensitivity was 0.82 for the CART method.

Random forests

The third learning method trialled for cell classification is an extension of decision tree cell classification called the random forest. This ensemble based method uses an average over multiple trees to improve the predictive ability of the classifier and limit over-fitting of the data. Rather than growing a single tree with many branches, a random forest builds an ensemble of shallower trees, where different weights are assigned to each feature [86]. Many permutations of trees are generated by using different subsets of the data. This leads to different classification results in each tree depending on the bias of the data subset.

The key idea of an ensemble random forest decision tree implementation is that errors in shallow decision trees will be washed out when aggregated leading to a more accurate prediction. Figure 2.17a shows how the residual mean squared error rates of the random forest algorithm are reduced as the number of trees increase. For this cell classification problem, the random forest algorithm was implemented in the R programming language [80] version 3.2 with 500 trees and 3 variables considered at each split.

The random forest algorithm has several advantages over the SVM and CART algorithms. It is robust to over-fitting because the model is generated stochastically, so its generalisation abilities are better. It may be run quickly on large data sets, and trees can be generated in parallel.

The importance of each variable in the random forest method is computed from permuting out-of-bag data. The out of bag error is a cross-validation step performed to get an unbiased estimate of the test error rate. The out of bag error is calculated by leaving out about one third of the trees. Each left out case is then used in the construction of the k^{th} tree to get a classification. At the end of the run, j is taken to be the class that got most of the votes every time case n was out of bag. The proportion of times that j is not equal to the true class of n, averaged over all cases is the out-of-bag error estimate. The Gini index is used to calculate the node impurity. The random forest, and classification and regression tree algorithms, both identified the standard deviation of Feret diameter to be the most important indicator of cell type.

5-fold cross-validation The aggregated confusion matrix from 5-fold cross-validation of the random forest method is shown in Table 2.8.

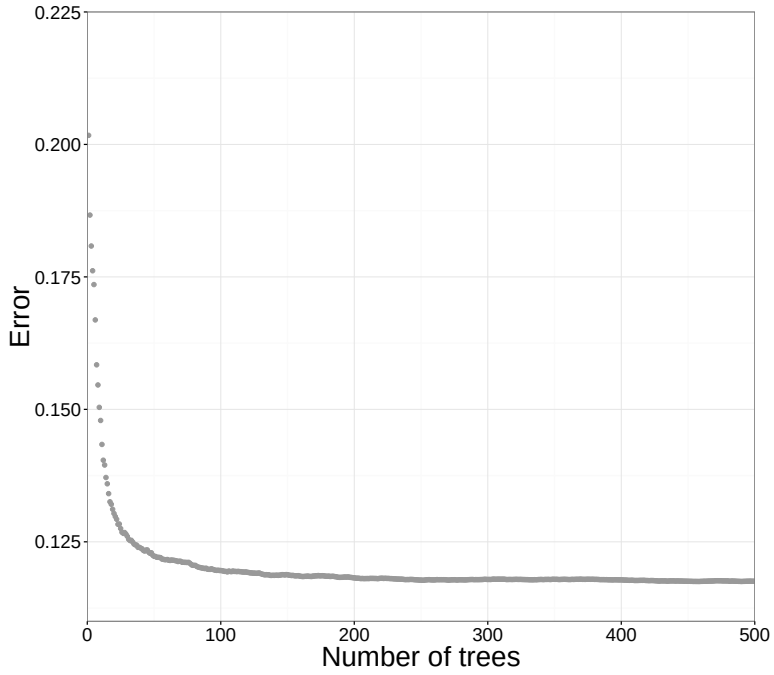

(a)

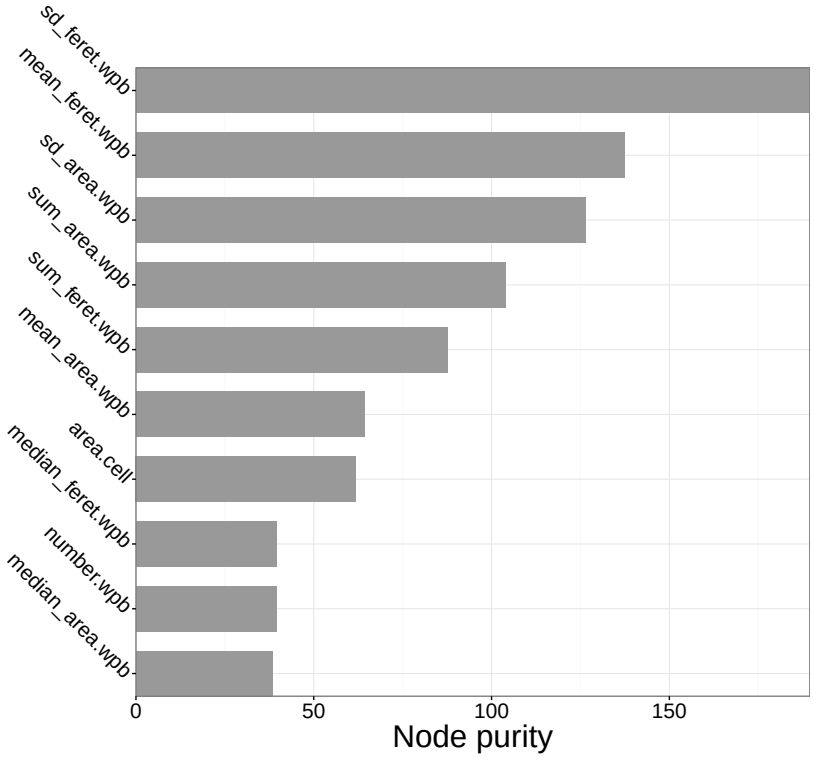

(b)

Fig. 2.17 Plot (a) shows the mean squared residuals (MSE) error rate plotted against the number of trees generated by the random forest algorithm. Plot (b) is a dot plot showing the importance of each variable, calculated as the Gini index on splitting on the variable, averaged over all trees.

Table 2.8 Confusion matrix of random forest method 5-fold cross-validation on cell classification.

		True	
		p	n
Predicted	p'	1460	250
	n'	392	1742

The cross-validation gave a sensitivity of 0.79 and a specificity of 0.85. The higher specificity (or true negative rate) suggests that the random forest classifier is more successful at classifying type 0 cells than type 1. The overall accuracy, calculated as the average of the specificity and the sensitivity, was 0.83 for the random forest method.

2.6.3 Results

The machine learning methods SVM, CART and random forest were used to classify all cellular data samples. In addition to the 5-fold cross-validation results on the labelled data, the accuracy of each method was evaluated by comparing the predicted ratios of cell types to the measured volumetric ratios (as given in Table 2.5). For each cell and each machine learning method, a predictive result between 0 and 1 was output, and this was rounded to either 0 or type 1 to form a binary classification (cutoff at 0.51). For each group in Table 2.5 the predicted ratio of cell types was calculated.

In addition to the volumetric ratio of cell types 0 and 1 created experimentally, a second synthetic proportional cell data set was created. This artificial data set was constructed from the labelled cell types in groups A and E. Randomly sampled data from the two groups were used to generate a set of five groups with proportions equal to those created volumetrically. The synthetic data set consisted of 2400 example cells, where each of groups A to E had 480 cells. Cells were of type 0, and type 1, in fractions corresponding to the volumetric mixed ratios. The SVM, CART and random forest machine learning prediction methods were then applied to these synthetic data to see whether variation can be attributed to biological causes or the machine learning algorithm.

The results from applying the three machine learning routines to the real and synthetic data types are given in Table 2.9 and visualised in Figure 2.18. All three machine learning methods were able to detect variations in ratios of cell types. The methods tended to be biased towards predictions to classification of cell type 1, which might be explained by the natural variability in WPB numbers creating a subpopulation of untreated cells that do not contain any WPBs. A subset of cell type 0 with very few WPB will be indistinguishable from

cell type 1. Generally, predictions from the synthetic data were closer to true ratios, than to the real volumetric ratio data. This could reflect the error in volumetric measurement (see Table 2.9).

Table 2.9 Results of machine learning algorithms applied to acquired biological data and synthetic data

Group	Columns	Ratio	Real			Synthetic		
			SVM	CART	RF	SVM	CART	RF
A	1-2	0.00	0.22	0.25	0.01	0.19	0.23	0.00
B	3-4	0.25	0.40	0.44	0.39	0.36	0.37	0.25
C	5-6	0.50	0.51	0.53	0.50	0.55	0.56	0.49
D	7-8	0.75	0.69	0.70	0.70	0.70	0.72	0.74
E	9-10	1.00	0.89	0.89	0.98	0.91	0.92	0.99
Error			0.55	0.63	0.22	0.49	0.52	0.03

A measure of the error of each method was calculated by summing the cumulative error in each ratio predication for groups A to E. This was performed for both the real and synthetic data in Table 2.9. By this metric, the random forest method performed the best at cell classification in this study.

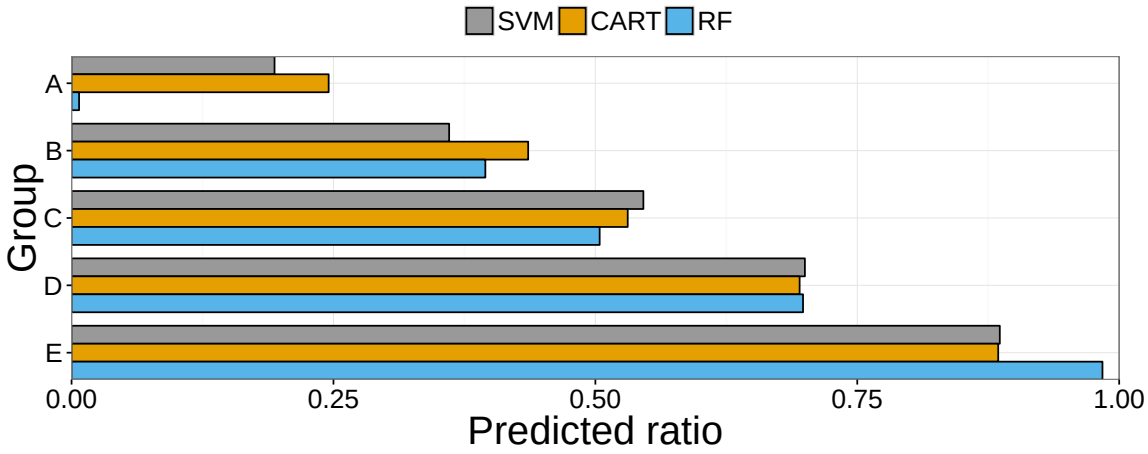

(a)

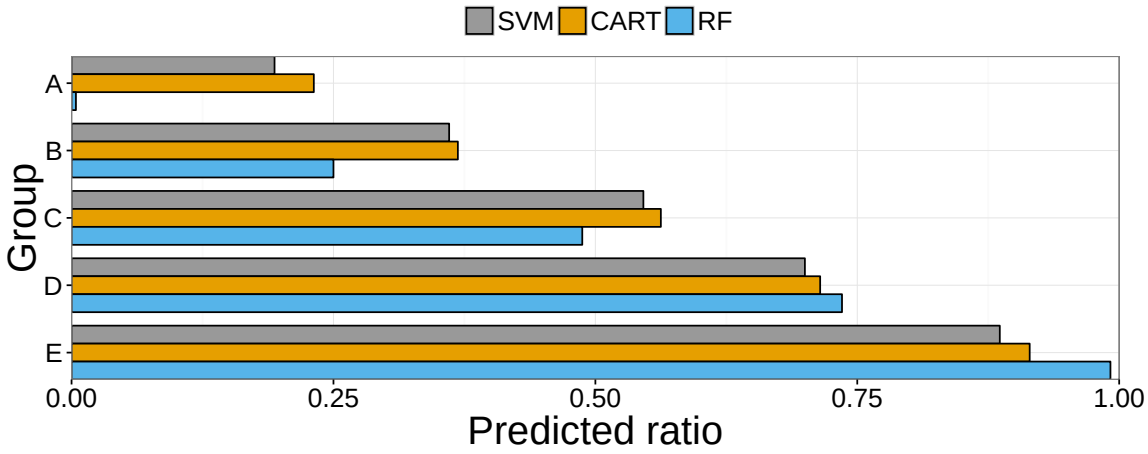

(b)

Fig. 2.18 Results of three machine learning algorithms applied to predict the ratio of two endothelial cell types for five ratio groups: A=0.00, B=0.25, C=0.50, D=0.75 and E=1.00. The prediction results from the three methods for the real volumetric ratios are shown in (a) and prediction results on the synthetic data ratios are shown in (b). The machine learning techniques that were used are support vector machines (SVM), classification and regression trees (CART), and random forests (RF).

2.7 Summary

This chapter presented a high-throughput morphometric analysis of endothelial organelles is introduced in section 2.1, with a description of equipment and image acquisition methods given in section 2.2. In section 2.3 methods were presented to segment and extract morphometric features from microscopy images of endothelial cells and their prominent organelles. A set of quantitative evaluation metrics to asses the performance of these segmentation methods is described in section 2.4.

The segmentation method was able to detect nuclei with a high degree of accuracy, where the precision, recall and F1 scores for detection of nuclei within images were each 0.97 (see Table 2.3). Cell segmentation presented a more challenging problem since accurate contour detection was required, and multiple plasma membrane stains were used. The most accurate method of cell segmentation gave a Dice score of 0.78 and a Jaccard score of 0.69 when evaluated over 517 hand labelled gold standard cells. Finally an evaluation of the performance of detection of vWF exocytic sites was performed using a gold standard set of 720 hand labelled vWF exocytic sites from nine images. The automatic segmentation method was able to correctly detect 683 of the 720 sites, giving a recall of 0.95, a precision score to 0.88, and an F1 score of 0.91.

Results from a high-throughput study of vWF exocytic sites are presented in section 2.5, this simplified case is included to demonstrate the kinds of experiments that have been conducted, and how the data can be visualised. Many plots could have been shown for this purpose from a large back catalogue of analysed experiments.

In section 2.6 a comprehensive experiment constructed to assess the experimental process, segmentation, feature extraction and, data analysis is described. The accuracy of the experimental, image processing and analytical approaches were investigated when used for cellular analysis. A data set of endothelial cells containing known volumetric ratios of two cell types was generated. Each cell in the data set was segmented and a set of cellular features extracted. The features were normalized to reduce the variability between features (see Equation 2.4). Two groups in the data were of known cell type: untreated and siRNA treated. These two groups formed labelled data groups, which were used as training data in supervised machine learning. Three machine learning models: support vector machines (SVM), classification and regression trees (CART), and a random forest method were built and tested, to predict the ratio of cell types in each group, and this was compared to the known ratios. For the real and synthetic data it was found that the random forest method was able to predict cell types with a high degree of accuracy (see

Table 2.9), suggesting that the developed experimental and image analysis approach is robust enough to provide reliable biological results.

Chapter 3

Image analysis for platelet dense-granule deficiency diagnostics

Platelets contain alpha-granules, dense-granules and lysosomes that secrete bioactive molecules important for many platelet functions. Depletion or deficiencies in these granules are associated with bleeding disorders. A commonly applied approach to diagnose patients with bleeding disorders caused by dense-granule defects is to image a sample of their platelets with an electron microscope and count the number of dense-granules. This approach requires specialised equipment and techniques, it is laborious, is difficult to automate in a clinical setting, and is highly operator-dependent. Developments in super-resolution microscopy have provided advantageous alternative imaging techniques that can be used to image platelet dense-granules. Imaging of platelets and their dense-granules via super-resolution microscopy provides the basis of a novel diagnostic tool to identify patients with bleeding disorders.

This chapter describes the development of image analysis techniques aiming to automatically count, and measure the morphometry of platelet dense-granules in platelet images obtained from super-resolution microscopy. A proof-of-principle data set was acquired and an image processing workflow constructed to assess the viability of super-resolution microscopy imaging for diagnosing dense-granule deficiency disorders. This chapter is based on work contributing to the published article Westmoreland et al., 2016: 'Super-resolution microscopy as a potential approach to platelet granule disorder diagnosis' [37]. Image data for use in this chapter was acquired by David Westmoreland.

3.1 Introduction

Platelets are cytoplasmic fragments without nuclei, derived from megakaryocytes. They are the smallest cells present within blood and are crucial to haemostasis and thrombosis. They contain an open canalicular system, a dense tubular system, and secretory organelles. Platelet secretory organelles have been classified into four distinct groups based on their content and appearance by electon microscopy; dense-granules, α-granules, multivesicular bodies, and lysosomes [87]. The absence of dense-granules in platelets significantly slows the rate of effective haemostasis at the site of an injury [88]. Dense-granules contain serotonin, histamine, ADP, ATP, GDP, GTP, magnesium pyrophosphate [89], polyphosphate [90], and a dense calcium core.

In case of vascular injury platelets rapidly adhere to the subendothelium at the site of vascular injury and secrete their granules. This provides a localised and concentrated delivery of effector molecules at sites of vascular injury [91]. Secretory granules release bioactive compounds which mediate platelet functions such as aggregation. On release of their contents a platelet granule recruits other platelets to form a platelet plug or clot.

Haemostasis is dependent on platelet granule integrity; deficiencies in these organelles can result in excessive bleeding [92]. Hermansky-Pudlak Syndrome (HPS) and Chediak-Higashi syndrome (CHS) are genetic disorders that can be diagnosed based on a lack of platelet dense-granules. A host of symptoms can manifest in patients with HPS and CHS, including; eye and skin albinism, visual impairment, bleeding disorders, inflammatory bowel disease, and lung fibrosis [93]. Bleeding disorders arise because a lack of dense-granules increases the time for effective clot formation at the site of an injury to occur [88].

Techniques for diagnosing platelet granule disorders have been limited to either counting dense-granules in electron microscopy images, measuring levels of ADP/ATP, or assessing aggregation response to agonists. Presently the most reliable method for diagnostics involves counting granules in electron microscopy images. This requires specialised equipment, is highly laborious, and even amongst experts there can be disagreement about how to identify these dense-granules. Figure 3.1 shows an example of platelets imaged via electron microscopy. Dense-granules can be observed in Figure 3.1a as black spots. Similarly Figure 3.2 is an example of platelet images acquired for this study via super-resolution microscopy, where CD63 was used to characterise dense-granules. The relative absence of CD63 seen in magenta can be observed in HPS patients in Figure 3.2b, Figure 3.2c, and Figure 3.2d, as compared to platelets from a healthy volunteer in Figure 3.2a. Dense-granules in platelets have a diameter of approximately 150 nm to 400 nm. Before the rise

of sub-diffraction limit imaging the only viable way to image platelet dense-granules was via electron microscopy [94].

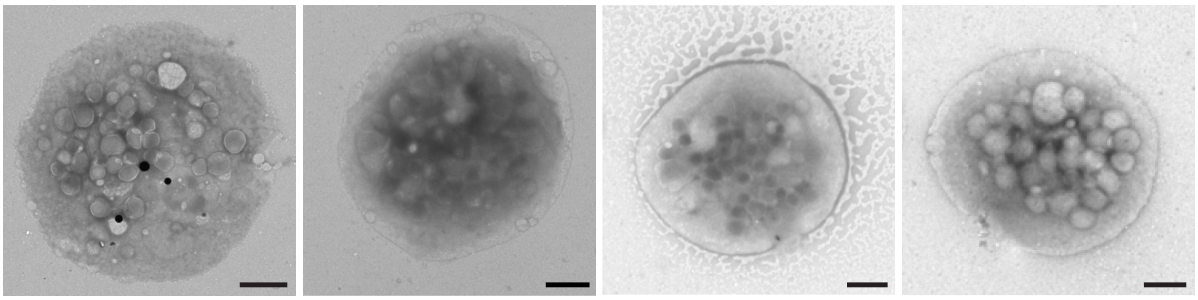

(a) (b) (c) (d)

Fig. 3.1 Whole mount electron microscopy images of platelets. Image (a) is an example of a single platelet from one of 7 healthy volunteers. Images (b, c and d) are examples obtained from each of HPS patients one, two and three, respectively. Note the lack of dense-granules (black spots) seen in the HPS patients (b, c and d) compared to the example from the healthy volunteer in image (a). Scale bars: 1.0μm. Adapted from Westmoreland et al., 2016 [37].

In this study super-resolution microscopy techniques have been used to image CD63 characterising platelet dense-granules. The antigen CD63 is an itinerant integral membrane protein that traffics between post-Golgi organelles, including dense-granules [95]. It has also been found to be altered in a patient with dense-granule deficiency [37]. Super-resolution microscopy techniques increase the attainable resolution from light microscopy by circumventing the diffraction barrier, such that dense-granules characterised by CD63 can be imaged via light microscopy. A multitude of techniques have been developed to overcome the diffraction limit.

Structured illumination microscopy (SIM) is a super-resolution technique that exploits Moiré effect to overcome the diffraction limit [96]. SIM illuminates a sample with patterned light and measures the fringes in the Moiré pattern (see subsection 1.2.3). SIM imaging of platelets allows for a new immunohistochemical approach to diagnosis of platelet granule disorders, since SIM can discriminate positive marker specific α-granules that are 200 nm to 400 nm across and dense-granules approximately 150 nm within a 2 μm to 3 μm platelet.

This new platelet dense-granule imaging approach required a robust image analysis and quantification protocol to differentiate SIM platelet images of healthy control patients to patients with dense-granule abnormalities caused by Hermansky-Pudlak syndrome (HPS). Using CD63 as a marker an image analysis workflow has been developed (see section 3.3)

and applied to a proof-of-principle data set with seven healthy controls and three patients with HPS.

3.2 Image acquisition

Whole blood was collected from each of the three patients with HPS and a further seven healthy volunteers. The three HPS patients had a similar clinical phenotype, whilst the healthy volunteers had no known history of bleeding disorders. All HPS patients were noted to have oculocutaneous albinism from birth and a history of excessive bleeding after trauma.

Each whole blood sample was collected into a solution of acid citrate dextrose and centrifuged. Platelet-rich plasma was collected, left for 30 min, diluted in a Tyrode's buffer and fixed with a formaldehyde in Phosphate-buffered saline (PBS) at a final concentration of 10 % for 10 min before centrifugation in coated coverslips. Permeabilisation with 0.2 % TX-100 in PBS was followed by incubation with primary, then secondary antibodies conjugated to Alexa Fluor dyes or Cy5 and mounted. Imaging used an inverted widefield fluorescence microscope (IX71, Olympus) modified for super-resolution SIM.

Each super-resolution image was reconstructed from a sequence of nine images acquired under excitation with nine different sinusoidal illumination patterns [97]. An additional image for comparison was created summing all nine raw images. Out of focus light in each image was suppressed by multiplication of the zero and first order SIM passbands by Gaussian and complementary inverted Gaussian functions [98]. The two-colour images were acquired sequentially under excitation of the sample with laser light at 488 nm (CD63) and 561 nm (tubulin). Image z-stacks were obtained by axially translating the specimen in 0.2 µm steps using a piezoelectric translation stage (NanoScanZ, Prior Scientific).

Examples of acquired SIM images from the three HPS patients and a healthy control can be seen in Figure 3.2, where the absence of dense-granules characterised by CD63 can be observed in Figure 3.2b, Figure 3.2c, and Figure 3.2d, as compared to the granules present in healthy control platelets in Figure 3.2a.

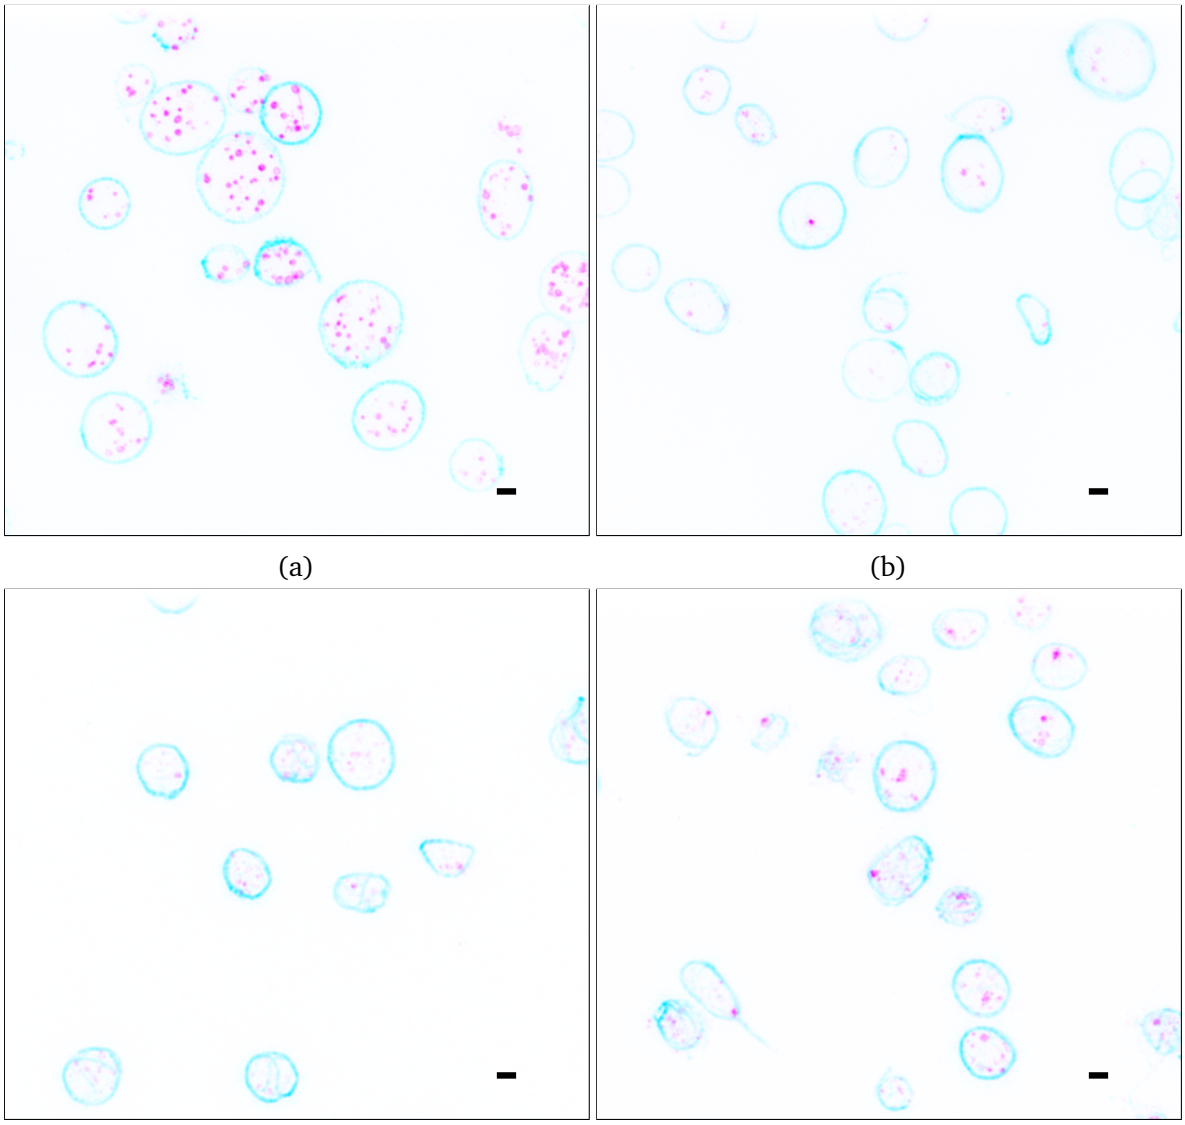

(a) (b)

(c) (d)

Fig. 3.2 Structured illumination microscopy (SIM) example images of platelets from one of seven healthy volunteers in image (a). Image (b, c and d) were obtained from Hermansky-Pudlak (HPS) patients 1, 2 and 3, respectively. The two-color channels were acquired sequentially under excitation of the sample with laser light at 488 nm (CD63) and 561 nm (tubulin); z-stacks are displayed as maximum intensity projections and in the lookup table has been inverted in these images for clarity, where tubulin is cyan and CD63 is magenta coloured. Scale bars: 1.0 µm.

3.3 Image processing

An image processing pipeline was developed to segment CD63 positive-granules and platelets in SIM images. The image processing pipeline in this proof-of-principle study was developed as a macro in ImageJ [99] version 1.49b. The use of ImageJ was particularly convenient for prototyping image processes, because each image filter applied was shown visually. In this pipeline platelets were identified by their tubulin labelled structures, where tubulin forms a layer just under the limiting membrane providing a convenient marker of platelet perimeters. Dense-granules were characterised by the presence of CD63, which is a promiscuous protein, and is likely found in all granules and in lysosomes. Therefore not all CD63-positive structures are dense-granules, however there is a usable difference in CD63 between HPS patient platelets and platelets from healthy volunteers that can be used in diagnosis. The approach to segmentation of platelets is described in subsection 3.3.2, and the segmentation of CD63 positive structures is described in subsection 3.3.3. An additional analysis of the radial distribution of CD63 positive-granules is described in section 3.4.

In each image, morphometric features were extracted from segmented platelets and each CD63 positive structure. Importantly, the number of CD63 positive-granules per platelet was also counted, from this the mean number of CD63 positive-granules per group could be calculated. The mean number of dense-granules provided the primary indicator for diagnosis of HPS. Accurate segmentation of both CD63 positive-granules and their constituent platelets was required to get a measure of the number of granules per platelet.

Each of the 194 images in the acquired proof-of-principle data set was processed sequentially and a series of outputs generated. These included images showing image segmentation contour overlays and morphological results tables. An overlay of segmentation contours allows the segmentation result to be checked visually to ensure it is accurate. The flowchart in Figure 3.3 presents each step in the image processing pipeline for segmentation and feature extraction of CD63 positive structures within platelets.

3.3.1 Region properties

The ImageJ measure command was used to measure morphometric and pixel intensity features for each segmented foreground area. By redirecting from the binary mask image to a pixel intensity image measurements of pixel intensities could be taken. Pixel intensity and morphometric properties were measured for every segmented CD63 positive structure and segmented tubulin labelled platelet. All the measured attributes along with a description for each are displayed in Table 3.1.

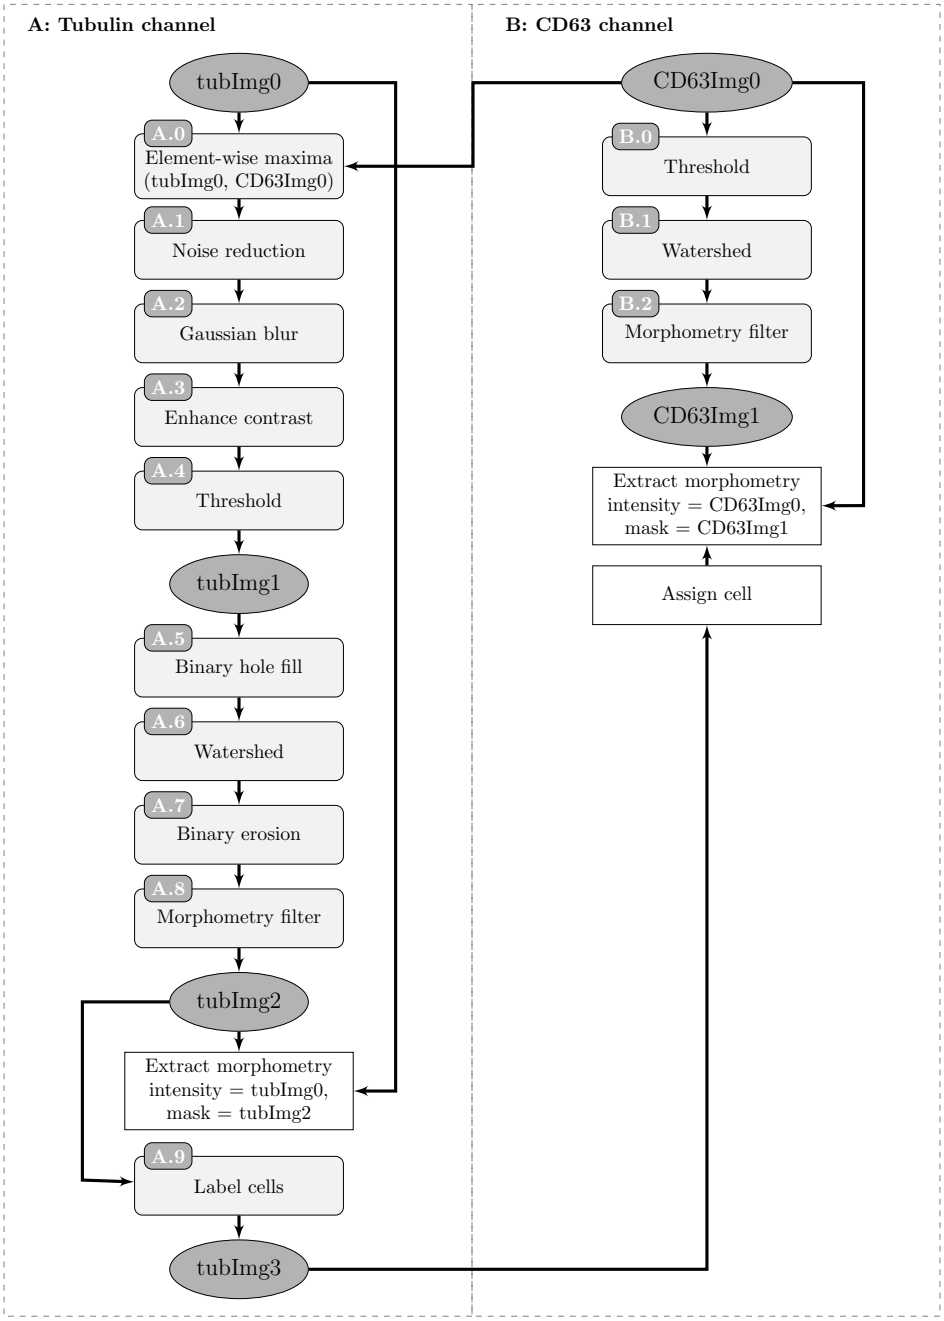

Fig. 3.3 Flowchart of segmentation and feature extraction processes of platelets and CD63 positive-structures. Dark ellipses in the flowchart are significant intermediary images, rectangles with rounded corners show image processing steps, and plain rectangles show feature extraction steps.

Table 3.1 Terminology and descriptions of segmented object attributes in platelet image processing. Comprised of meta-data to identify the object, morphometric measurements and intensity measurements.

Attribute	Description
label	The image name and slice number.
area	The area of selection in square pixels or calibrated square units.
mean	The average pixel value, calculated as the sum divided by the number of pixels.
standard deviation	The standard deviation of the pixel values.
min	The minimum pixel value.
max	The maximum pixel values.
centroid	The average of the x and y coordinates for pixels in the selection.
perimeter	The length of the boundary line of the selection.
major	The major axis length of the fitted ellipse.
minor	The minor axis length of the fitted ellipse.
angle	The angle between the primary axis and a line parallel to the x-axis of the image.
circularity	A metric for how circular the selection is, calculated by: $circ. = 4\pi \times \frac{[area]}{[perimeter]^2}$.
Feret	The longest distance between any two points along the selection boundary.
angle	The angle of the Feret's diameter from the horizontal is the Feret angle.
minimum Feret	The minimum caliper diameter is the minmum Feret.
Feret coordinates	The starting coordinates of the Feret's diametera.
integrated density	The sum of the values of the pixels in the selection or image.
raw integrated density	Raw integrated density returns the product of area and mean grey values.
median	The median value of the pixels in the image or selection.
aspect ratio	The aspect ratio of the selections fitted ellipse, calculated as the major axis divided by the minor axis.
roundness	$roundness = 4\pi \times \frac{[area]}{[majoraxis]^2}$ of the fitted ellipse is the inverse of the aspect ratio.
solidity	$solidity = \frac{[area]}{[convexarea]}$ where the convex area is the area inside the convex hull of the selection.

3.3.2 Platelet segmentation

The segmentation of platelets was the first step in the image processing pipeline for platelet dense-granule deficiency diagnostics. Only features from CD63 labelled structures within segmented platelets could be extracted so ensuring accurate segmentation of platelets was important in this study. Segmented CD63 positive structures can then be assigned to their respective platelets. This was achieved by assigning each platelet in an image a unique pixel value identifier, as seen in Figure 3.4i. The pixel value at the coordinates of the CD63 segmented particle on the labelled image was measured to assign it to an appropriate platelet.

Element-wise maxima (A.0)

Segmentation and labelling of platelets was performed using a composite image combining the acquired tubulin image with the CD63 channel image. This was done because CD63 is localised within platelets, so foreground regions segmented could be maximised by combining the two channels. This improves the segmentation result especially at the platelet edge. The element-wise maxima of the tubulin and CD63 images provides a more distinct platelet marker by choosing the maximum pixel value at any location from both images. To this combined image a series of pre-processing steps were performed to enhance features of the image.

Pre-processing (A.1, A.2 and A.3)

A *rolling ball* background subtraction removes smooth continuous background illumination from the image [72]. This was implemented using the *sliding paraboloid* algorithm, which is a development on the *rolling ball* algorithm also available in ImageJ. This can be conceptualised by imagining that the greyscale image has a height dimension, corresponding to the pixel value at every point in the image. A ball is rolled over the image and the hull of the volume reachable is the background. Using the sliding paraboloid algorithm replaces the ball with a sliding paraboloid of the same curvature at its apex as a ball of a given radius.

The *rolling ball* was followed by a Gaussian blurring with a sigma value of three pixels (not impacting the image resolution), and finally a contrast enhancement. This stage of the image processing pipeline, following the addition of the CD63 channel and pre-processing steps is shown in Figure 3.4b.

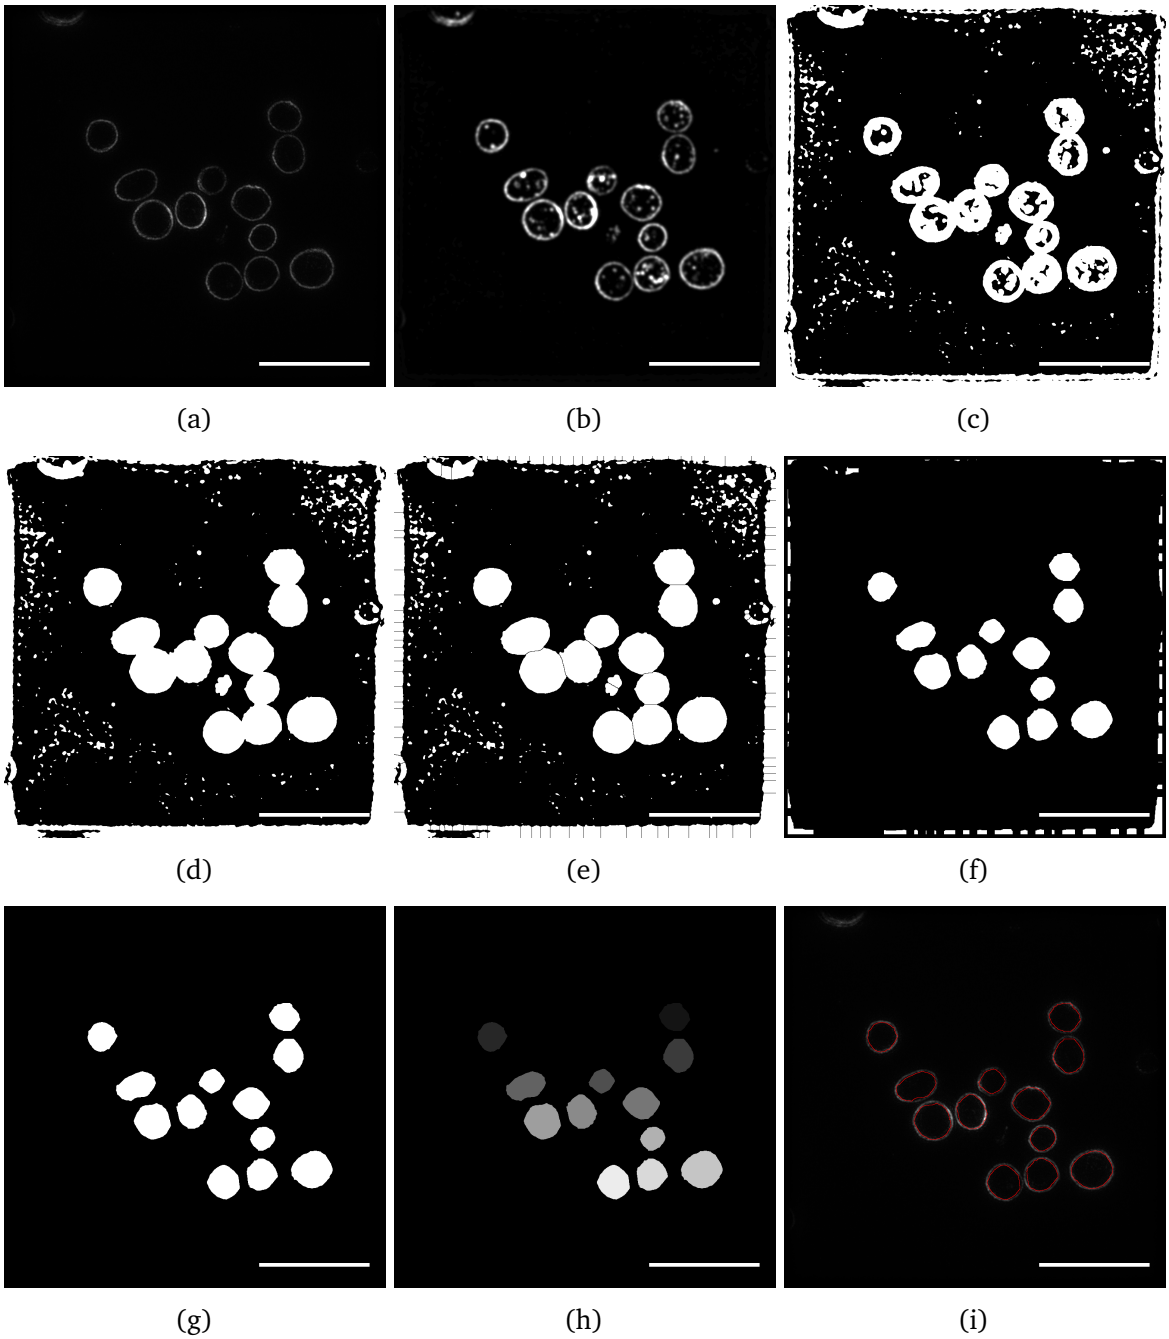

Fig. 3.4 This figure shows the major sequence of image processes involved in segmentation of platelets: image (a) is the raw image, and image (b) after the addition of the CD63 channel and following image pre-processing by de-noising, a Gaussian blur and contrast enhancement. Image (c) depicts the thresholding step using the minimum error threshold. Image (d) shows the binary hole fill, and image (e) following a watershed transform. Image (f) was obtained after a morphological filter, (g) after a binary erosion. Image (h) shows a labelled image and finally image (i) shows the image output with segmentation contours. Scale bars: 10.0 μm.

Threshold (A.4)

Segmentation of platelets was achieved using a minimum error threshold [52], where the three patient data threshold values were set according to control data that were acquired on the same day. The minimum error threshold sets a low threshold value to increase the probability of full segmentation of the platelet structure. Thresholding a full platelet tubulin ring was essential to retain the platelet in this image processing pipeline. A minimum error threshold was likely to capture the full platelet structure, however has the side effect of segmenting some background noise. Figure 3.4c shows an image following a minimum error threshold, in which segmented platelets are present as well as noise segmented as small white blobs.

Binary hole fill and watershed (A.5 and A.6)

The segmented structures in the binary image were filled with a binary hole filling algorithm, as shown in Figure 3.4d. A watershed transform was then performed to separate any touching platelets [55], this is demonstrated in Figure 3.4e.

Binary erosion (A.7)

The binary image shown in Figure 3.4d was eroded by 10 pixels in order to shrink the area of each platelet to within the tubulin ring. An erosion removes pixels from the edges of objects, and was used to avoid any contribution of CD63 found at the platelet surface. An erosion was necessary as the minimum error threshold sets a low threshold value where the segmented platelet structures tended to extend beyond the tubulin ring. To compensate for this binary erosion was used to reduce the size of segmented platelets. The effect of a binary erosion can be seen in Figure 3.4f, where the platelet size is reduced and much of the image noise also removed as a consequence of the binary erosion.

Morphological filter (A.8)

The final stage in segmentation of platelets was to remove foreground objects with an area of less than $2.0\,\mu m^2$ or circularity less than 0.7. The reported diameter of platelets varies from $2\,\mu m^2$ to $5\,\mu m^2$ and in this set of individuals our tubulin ring measurements gave areas ranging from $2.3\,\mu m^2$ to $3.1\,\mu m^2$. The circularity constraint removed debris or non-platelet objects from the image since the minimum circularity of platelets has been reported to be 0.7 [100]. The final segmented platelets overlay contours image is shown

in Figure 3.4g, with the segmentation overlaid on the original image in Figure 3.4i. A labelled image used to assign CD63 positive structures to their respective platelet is shown in Figure 3.4h.

3.3.3 CD63 positive-granule segmentation

Segmentation of CD63 positive structures marking granules was the second step in the image processing pipeline for platelet dense-granule deficiency diagnostics. The image processing steps for CD63 segmentation and assignment of granules to platelets are shown in Figure 3.3, where labelled CD63 granules appear as small round blobs as in Figure 3.5.

Threshold (B.0)

The segmentation of CD63 positive-granules was performed according to a threshold value obtained from Moment-preserving thresholding [53]. The threshold was calculated based on all images in each of the seven control CD63 image data sets. The three patient data threshold values were then set according to their corresponding control data that was processed on the same day. An example of a CD63 channel image following thresholding is shown in Figure 3.5b.

Watershed and remove small objects (B.1 and B.2)

A watershed transform was then performed to separate any touching CD63 positive-granules [55], this is demonstrated in Figure 3.5c. Areas smaller than $0.01\,\mu m^2$ were removed, these are less than the minimum resolution of the microscope. A set of morphometric features was measured for each identified CD63 granule according to Table 3.1. Each measured feature sample is assigned to a platelet based on the x and y coordinates of the object and the corresponding pixel value in the labelled platelet image.

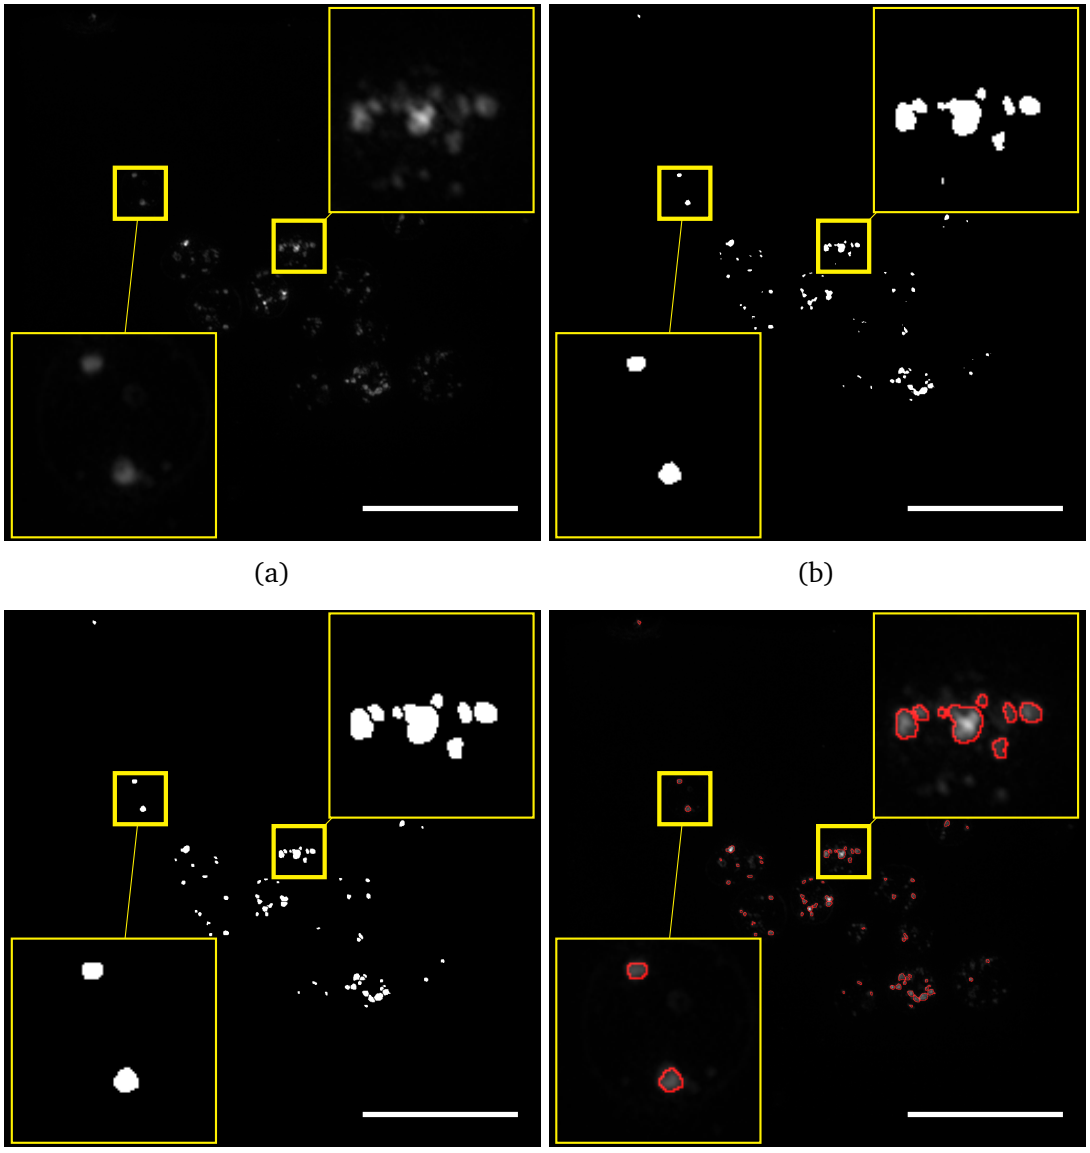

(a) (b)

(c) (d)

Fig. 3.5 The major sequence of processes involved in segmentation of CD63-positive granules: image (a) is the raw CD63 channel image, image (b) is after Moment-preserving thresholding, image (c) is following a watershed transform, and finally image (d) shows segmentation contours on the original image. The yellow circles are magnified regions to demonstrate the segmentation process better. Scale bars: 10.0 µm; magnified regions: 3.3×3.3 µm.

3.4 CD63 radial distribution analysis

A distinctive characteristic of cells defective in machinery that forms lysosome-related secretory organelles such as platelet granules is that defects in targeting of proteins to those granules can lead to an abnormal accumulation on the cell surface [101]. In the case of HPS patients we discovered that there is an increase in CD63 on their platelet surface. A radial distribution analysis was used to investigate and quantify the extent to which there was an increased accumulation of CD63 at the cell surface in HPS patients. The radial profile expresses the pixel intensity binned by their distance from a central point. Thus giving a set of pairs of measurements of intensity and radius that can be plotted to give a radial profile.

The radial profiles of HPS patient platelets was compared to radial profiles of platelets obtained from volunteers without bleeding disorders. Examples of CD63 distribution in patients with HPS and healthy volunteers can be seen in Figure 3.6, the pathological increase in accumulation of CD63 at the cell surface can be seen by comparing the platelet of a healthy volunteer in Figure 3.6c to the platelet from a HPS patient in Figure 3.6f.

The ImageJ radial profile plot plugin was used in this analysis to measure the radial profile of each platelet[1]. The plugin creates concentric rings from a central point and in each ring calculates the normalized integrated density. The normalized integrated density is the sum of pixel values in that ring divided by the number of pixels in the ring. The radial profile measures the radial distance from the center against the normalized integrated density. This is illustrated by the concentric rings in Figure 3.6c and Figure 3.6f.

The radial profile analysis was performed following morphometric analysis of the platelets data set. The coordinates of the centroid of each found platelet in the platelets data set was used as the center point for each radial profile. It was assumed that each platelet was a perfect circle in this analysis, where the maximum radius, (r), was calculated from the area, (A), as: $r = \sqrt{A/\pi}$. Since the area of each platelet was known the maximum radial distance was calculated, the radial profile plugin was applied at each platelet centroid with an appropriate area circle. The white circles surrounding platelets in Figure 3.6 show the platelets for which the radial profile analysis was performed in these example images.

From the total population of 2812 segmented and measured platelets, a subset of 1549 of highly circular platelets (with circularity >0.85) were selected for radial analysis. Radial integrated density measurements were taken for each platelet in the CD63 and the tubulin

[1] P. Baggethun, Radial Profile Plot ImageJ Plugin, http://imagej.nih.gov/ij/plugins/radial-profile.html, accessed 2015-07-15

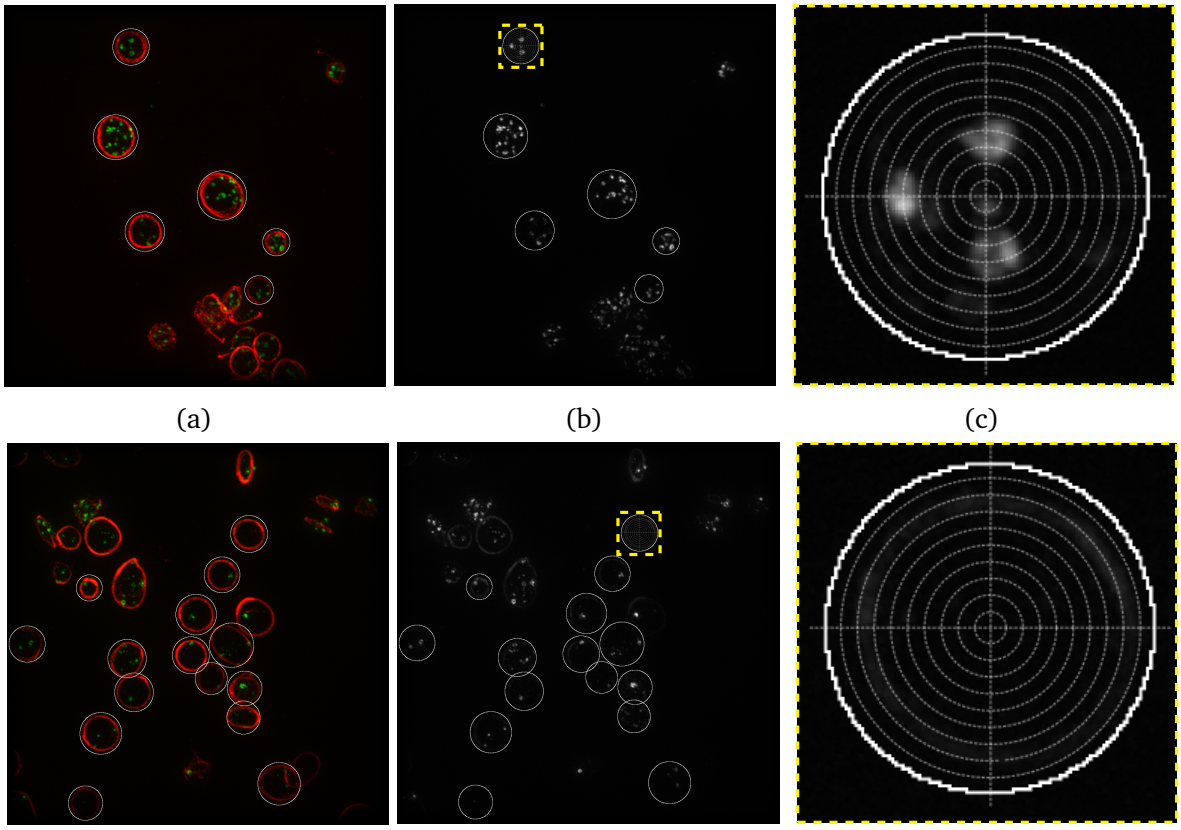

(a) (b) (c)

(d) (e) (f)

Fig. 3.6 Examples showing the CD63 at the platelet surface in Hermansky-Pudlak syndrome (HPS) platelets. A representative image with tubulin (red) and CD63 (green) staining is shown from the controls in image (a), the CD63 stain of this is shown in image (b), and a $9.0 \times 9.0\,\mu m$ magnified region showing a representative platelet with concentric rings is shown in image (c). Similarly a representative image from a Hermansky-Pudlak syndrome (HPS) patient is shown in image (d), the CD63 stain of this is shown in image (e), and a $9.0 \times 9.0\,\mu m$ magnified region showing a representative platelet with concentric rings is shown in image (f). The accumulation of CD63 at the platelet surface can be seen in the patient images. In these images all platelet with white rings indicate the outer ring in the concentric radial profile analysis, only platelets with a white ring were segmented and analysed in the radial profile analysis.

channels. The maximum integrated density concentric ring in the tubulin channel for each platelet was used to determine the location of the tubulin ring. This was used to define the plasma membrane, and measurements beyond this were excluded.

The data from each platelet was normalized to allow aggregation into groups independent of size and image intensity. Both the radial distance and pixel intensity values were normalized by dividing by the maximum, giving a scale from 0 to 1. This allowed each HPS patient and volunteer to be compared on normalized scales to see differences in CD63 radial distribution.

3.5 Results

Platelet samples from three HPS patients and seven healthy volunteers were obtained and both the electron microscopy and SIM methods used to differentiate between the patients and the volunteers. Results given here were obtained from SIM microscopy for which the image processing workflow described in section 3.3 was developed.

In each of the 194 images in this proof-of-principle dataset segmentation of platelets and CD63-positive structures was performed. The number of CD63-positive objects was counted for each platelet, and a host of additional features were also collected. The number of CD63-positive structures can be used to distinguish patients with dense-granule deficiencies from controls.

Figure 3.7 shows the ability of the SIM method to distinguish between platelets of healthy volunteers and HPS patients. The cumulative frequency distribution plot shows the three HPS patients shifted left of the healthy controls. All of the HPS patients had fewer CD63 dense-granules than the healthy volunteers. The mean number of granules per platelet provides a striking diagnostic indicator, across the controls the mean number of CD63 positive structures per platelet is 6.8±0.5, compared to 2.4±0.5 in patients with HPS.

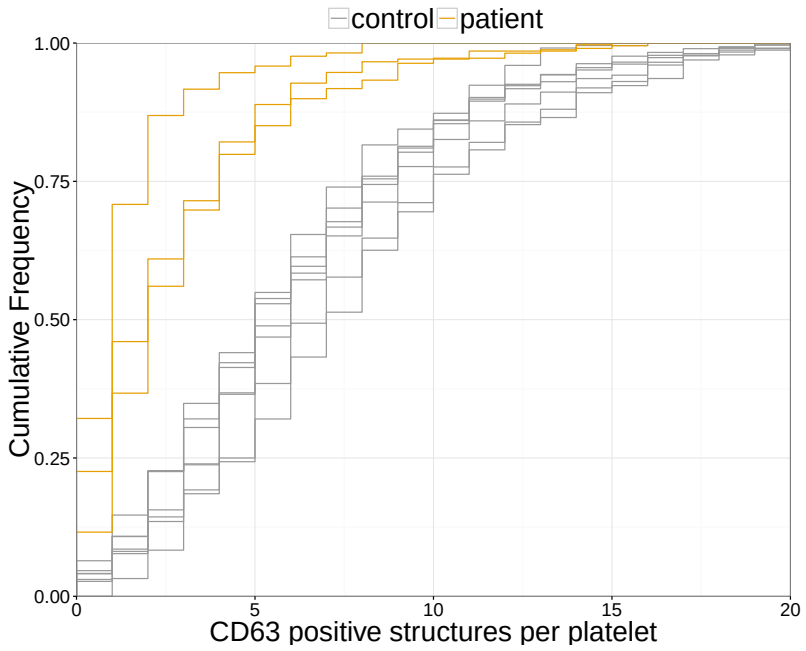

Fig. 3.7 Empirical cumulative distribution function of the number of CD63-positive structures per platelet for controls (healthy volunteers) and patients.

Results from the CD63 radial distribution analysis described in section 3.4 are displayed in Figure 3.8. There is a striking difference in the distribution of CD63 in HPS patients and healthy controls. In HPS patients there is a significant increase in CD63 located at the plasma membrane, this novel result as observed by eye could was confirmed using a CD63 radial distribution analysis.

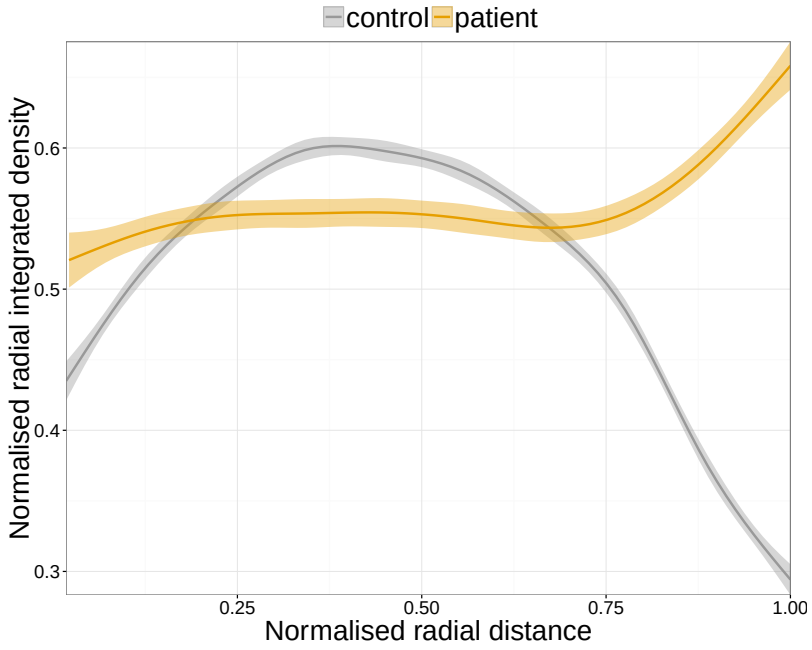

Fig. 3.8 Normalized integrated density radial profile of control group platelets and Hermansky-Pudlak syndrome patient platelets with 95 % confidence interval. The radial distance is measured from the platelet centroid, where 0 is the centroid and 1 the maximum of the tubulin ring.

3.6 Summary

This chapter described the use of image analysis in the application of super-resolution microscopy as a potential new approach to diagnosing platelet dense-granule deficiency disorders. In section 3.1 the work is introduced and in section 3.2 the sample preparation and image acquisition described.

In this study the SIM immunohistochemical approach was used in parallel with whole-mount electron microscopy for comparison. The itinerant integral membrane protein CD63, that traffics between post-Golgi organelles was used to characterise platelet granules. CD63 positive objects were co-stained with tubulin to demarcate platelet perimeters. An image processing workflow described in section 3.3 was used to segment granules and count the number of granules per platelet.

A super-resolution optical imaging approach was found to be effective and rapid in differentiating between healthy controls and patients with a platelet bleeding disorder. The automated unbiased SIM workflow requires less than 5 % of the time that subjective manual interpretation of electron microscopy images requires [37]. A fully automated image analysis workflow provides reproducible results and could be applied to multiple data sets as an automatic thresholding value set by the Moments algorithm in ImageJ [99] defines the CD63 structures. The methodology presented here could be applied at relatively low cost to a high-throughput microscopy platform with automated data analysis that would present the results to a clinician for interpretation.

This radial analysis described in section 3.4 is a novel approach to quantifying CD63 distribution in platelets. The combination of CD63-positive granule counting and radial analysis provides two independent fully quantitative measures of the HPS disease. In the future these quantitative analytical methods could be applied to other dense-granule deficiency disorders.

Chapter 4

Automated detection and tracking of leukocytes *in vitro*

To investigate the role of the endothelium, and in particular the role of machinery stored in the WPB, live imaging of leukocyte recruitment under flow *in vitro* was used. Interaction between the endothelial monolayer and leukocytes were recorded with a phase contrast microscope and high frame rate microscope camera. Manual quantification of interaction events in the acquired video sequences is very time consuming, potentially inaccurate, and laborious.

This chapter describes a quantitative computational method for leukocyte cell detection and tracking using a Haar-like features object detection framework. The method is capable of automatically identifying and accurately tracking multiple leukocyte trajectories *in vitro* adhesion assays.

Image data for use in this chapter was acquired by: Jess McCormack and Francesca Patella.

4.1 Introduction

Mechanisms of leukocyte recruitment can be studied in leukocyte endothelium flow adhesion assays *in vitro*. In these assays physiological levels of shear stress are maintained and leukocyte interaction with the endothelium recorded. Recruitment of leukocytes requires leukocyte interaction with vascular endothelium and consists of multiple steps in a consecutive adhesion cascade (see subsection 1.1.4 and Figure 1.2). The mechanisms of adhesion can be inferred from analysis of leukocyte endothelial interaction in flow assays, where the endothelial monolayer is in differing activation states.

To observe mechanisms of leukocyte recruitment across human endothelial cells an *in vitro* fluid flow chamber was used (see Figure 4.1). A microscope capable of phase contrast microscopy connected to a flow assay system recorded video frames of leukocytes in culture medium, flowing through the chamber, over a confluent monolayer of HUVECs. Physiological levels of shear stress were maintained by controlling the flow rate (see subsection 4.1.1). Hypotheses could then be tested by comparing the effect of endothelial cell conditions that inhibit or promote leukocyte recruitment as compared to a control group.

Phase contrast microscopy is advantageous for the study of live cells and their motility. It provides sufficient contrast for cell edge detection without exogenous dyes, using only moderate levels of light. Resulting images are free from staining artefacts and photodamage, as are common in fluorescence microscopy [102]. Phase contrast microscopy translates tiny variations in the phase of incident light from a specimen into perceptible changes in light amplitude. Phase differences in light occur due to the relative refraction index of the medium through which the light has travelled (see subsection 1.2.4).

Quantitative analysis of leukocyte flow assays aims to evaluate the degree of interaction of the monocyte-like transformed line THP-1 cells with the endothelium. Experiments *in vivo* often use a measure called the *rolling flux*, which is the number of leukocytes crossing an imaginary line across the observed vessel [104]. A traditional approach to quantification *in vitro* involves counting the number of interactions occurring between the leukocytes and endothelium, where an interaction would be deemed an event in which a leukocyte slows significantly or stops. This method of analysis is time consuming, requires a subjective judgement and is prone to observer bias. It is proposed that a computational approach to quantify these rolling assays, would yield much more quantitative data about interaction events, leading to biological insight into the adhesion cascade and mechanisms of leukocyte rolling.

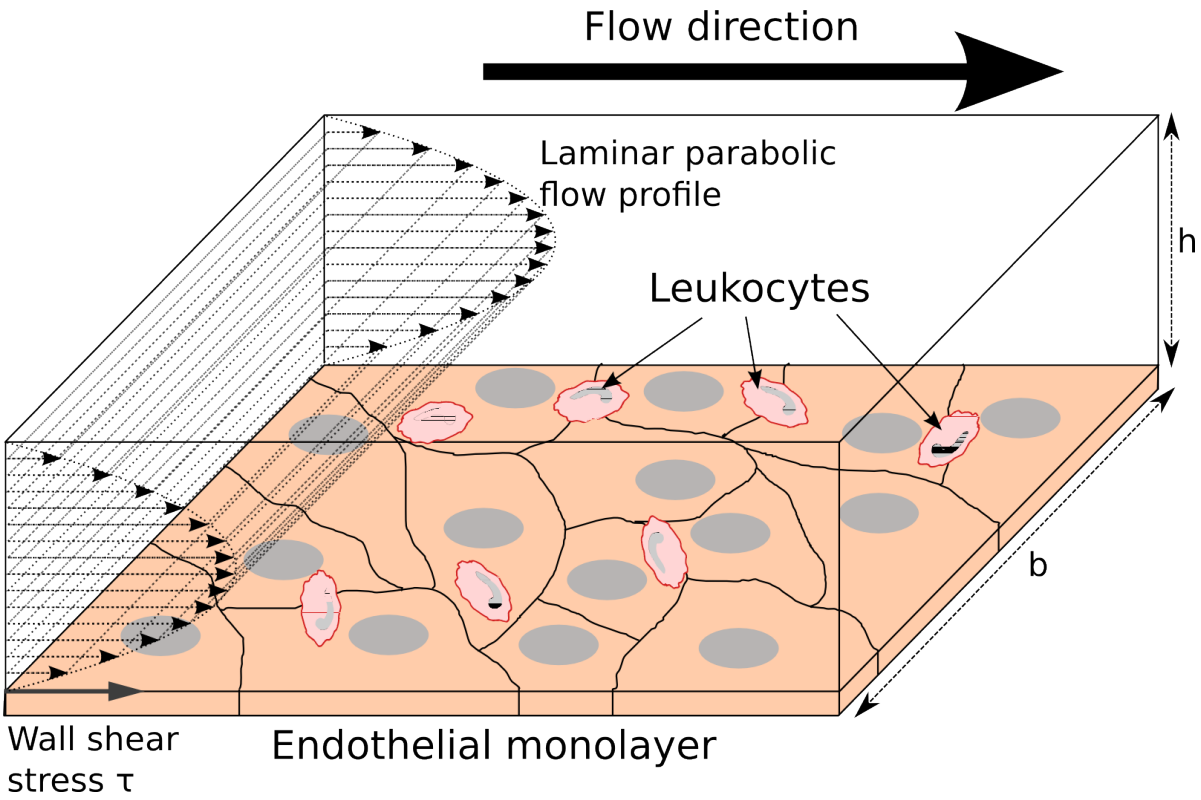

Fig. 4.1 An illustration of a leukocyte flow assay with a fluid flow chamber connected to a pump to push through a culture medium containing leukocytes over a confluent endothelial monolayer. A laminar flow profile of the culture medium creates a wall shear stress at the endothelial surface, this shear stress is maintained at a physiologically relevant value. Adapted from Jossen et al., 2014 [103].

Existing computational approaches to tracking of leukocytes have mainly been focused on the *in vivo* modality. There is the additional challenge *in vivo* of removing the respiratory motion during the video [105, 104]. Techniques exist for quasi-automated tracking of leukocytes [106, 107] *in vitro*, where leukocytes are identified manually by the operator.

For our experimental setup a solution involving operator identification is impractical due to the quantity of leukocytes to be identified over the duration of the video. The method described here has been developed for fully automated and unbiased leukocyte detection and tracking without operator identification and minimal parameter adjustments between experiments.

4.1.1 Shear stress

The shear stress under which these *in vitro* experiments are carried out was set to approximate the conditions under which these vein endothelial cells would operate. The rate of shear at the endothelial monolayer is proportional to the flow rate in these experiments, and the system maintains a suitable flow rate to generate an appropriate shear stress. Calculation of the shear stress at the endothelial monolayer, known as the wall shear stress is important in the interpretation of the acquired experimental data.

Shear stress is a force along a plane parallel to the imposed stress causing deformation by slippage of a material. Endothelial cells are subjected to shear stress from blood flow and a mechanical shear stress is converted into intracellular signals that affect cellular functions. In a flow assay a physiologically relevant shear stress was maintained over the endothelial monolayer in the μ-slide chamber by controlling the flow rate (Φ) in the system. To calculate the resulting shear stress on the cells, the mathematical model assumes a Newtonian fluid in laminar flow between infinitely wide parallel plates, the wall shear stress (τ) is calculated as:

$$\tau = \frac{6\eta}{h^2 b}\Phi, \tag{4.1}$$

where, (η) the dynamical viscosity, (h) the height of the channel and (b) the channel width [108]. A viscosity value of 1×10^{-3} Pa appropriate for water at 22 ℃ is here approximated for the viscosity of the cell culture medium containing leukocytes, and a shear stress of 0.07 Pa maintained.

4.1.2 Leukocyte interaction velocity

Fluid flow through a tube or pipe can generally be laminar or turbulent, where the flow profile is largely dependant on the flow rate and viscosity of fluid. Turbulent flow is a chaotic flow regime characterised by the presence of vortices or eddies. Generally, turbulent flow occurs at higher flow rates and is difficult to model. On the other hand, laminar flow in a pipe has a predictable parabolic distribution of velocities across a cylindrical cross-section. In laminar flow the velocity of parallel fluid layers have a constant velocity relative to neighbouring layers [109]. This is manifest in a fluid flow chamber as a profile increasing from stationary at the walls towards a maximum at the cross-sectional centre (see Figure 4.1). Laminar flow was assumed within the parallel-plate flow chamber used in the leukocyte interaction assays, given the leukocyte medium viscosity and flow rate of the system.

Distinguishing leukocytes retarded by interaction with the endothelial monolayer was complicated by the differential velocity profile in laminar flow. The velocity of a leukocyte is dependent not only on its position in the flow profile but also molecular interactions occurring with the endothelial monolayer. In the experimental setup used with the microscope acquiring from an above view, it was not possible to determine the height of the leukocyte in the flow profile. There was insufficient information to distinguish free flowing and interacting leukocytes at all heights in the flow profile.

A minimum velocity or critical velocity of non-interacting leukocytes can be calculated, on the assumption that most interaction occurs for leukocytes flowing close to the vessel wall. The leukocytes in faster flow-streams if interacting will be retarded and pulled towards the vessel wall. In the outer flow-streams a critical velocity has been calculated to distinguish interacting leukocytes from free-flowing leukocytes.

Calculation of the leukocyte critical velocity requires in the first place an approximation of the mean flow velocity ($\bar{v}$), of the leukocyte medium in the parallel-plate chamber,

$$\bar{v} = \frac{v_{max}}{2 - \epsilon}, \tag{4.2}$$

this is based on some known parameters and the ratio $\epsilon = D_L/D_C$. Where (D_L) is the leukocyte diameter and (D_C) the chamber diameter. The chamber diameter was given as 0.4 mm and a value of 15 µm was used as the average diameter of a leukocyte [110]. The value (v_{max}) is the maximum leukocyte velocity or the velocity of leukocytes in the centre fastest laminar stream.

A value of (v_{max}) was approximated by measurement of the distance travelled during the camera exposure of the fastest leukocytes in a video. This can be seen in a video frame by the length of the leukocyte trail, this is the distance travelled by a leukocyte during the camera exposure time. From the distance travelled and the camera exposure time of 30 ms a maximum velocity can be calculated.

The critical velocity (v_{crit}) can then be estimated as,

$$v_{crit} = \bar{v} \cdot \epsilon \cdot (2 - \epsilon), \tag{4.3}$$

using the mean flow velocity ($\bar{v}$) and the ratio (ϵ) in this experimentally derived equation [111]. Any cell with a velocity below the critical velocity is likely to be retarded by an adhesion interaction with the vessel wall. For quantitative analysis of leukocyte-endothelium interactive flow assays, the distribution of leukocyte velocities can be compared across acquisitions in differing endothelial cell conditions, an estimate of the number of interacting leukocytes as derived from Equation 4.3 was also instructive.

4.2 Video acquisition

HUVECs were seeded into μ-slides $I^{0.4}$ (Ibidi, Munich, Germany) and treated with IL-4 (Ibidi) 24 h before experimentation. The slide was placed on the microscope stage of an Axiovert 135 (Carl Zeiss, Welwyn Garden City, UK) maintained at 37 °C and imaged using a 20× lens. The μ-slide was connected to a syringe pump system (Harvard Aparatus, Holliston, MA, USA) to draw fluid through the chamber to give a wall shear stress of 0.07 Pa (0.7 $dyne\ cm^{-2}$). Perfusion media (HBSS + Ca^{2+} + Mg^{2+} + 0.2 % BSA) was drawn through the chamber for a few minutes to flush thorough any cellular debris and ensure that the HUVEC monolayer was intact. Control or genetically or pharmacologically altered HUVECs were treated with perfusion media with or without an activating stimulant or drug treatment, for 5 min to 10 min under flow. $5 \times 10^5\,ml^{-1}$ THP-1 cells (grown in RPMI medium 1640 (Gibco BRL) plus $4.5\,g\,l^{-1}$ D-glucose, $1.5\,g\,l^{-1}$ sodium bicarbonate, 1 mmol sodium pyruvate 10 mmol hepes, $300\,mg\,l^{-1}$ L-glutamine, 0.05 mmol 2-mercaptoethanol and 10 % FCS) were added to the media with or without a secretagogue or drug treatment, to stimulate or mock stimulate HUVECs. For each condition typically a 3 min to 6 min movie was recorded to observe any THP-1 adhesive interactions with the monolayer. Videos were captured using a QImaging Scientific CMOS (sCMOS) camera.

4.3 Video processing

Analysis of leukocyte adhesion cascade image sequences aims to:

1. accurately detect the leukocytes in each frame of the video,
2. link corresponding leukocytes over multiple frames to form trajectories, and
3. quantify information about those trajectories.

A variety of approaches were trialled for accurate detection, linking and quantification of leukocytes in flow assays.

Manual identification of leukocytes in each frame was impractical due to the large numbers of frames and leukocytes to be identified and the level of reproducibility required. Accurate and robust identification by automated threshold based segmentation is limited for phase contrast images, although increasingly common in fluorescence microscopy [112]

Two prominent artefacts introduced by phase contrast microscopy are halos and shade-offs. A halo is a bright region surrounding a specimen, and the shade-off is an effect gradually reducing the intensity profile towards the centre of a large specimen. Halos were observed in this study as the result of diffracted light passing through the phase ring as well as the non-phase areas and interacting at the image plane [113]. The halo adds artificial structure to the specimen. In our setting, bright halos appear around the leukocytes with an intensity and width that depend on the local thickness of the cell. Shade-offs equalise the intensity of the inner and outer regions of large leukocytes to the same value [39]. These two artefacts complicated the task of segmentation and hampered the use of thresholding techniques. The bright ring surrounding the adhering leukocyte in the lower left image of Figure 4.2b is an example of a halo, the shade-off can also be observed in this image. Phase contrast artefacts and the elongated shape of high velocity leukocytes renders blob detection via the Laplacian of the Gaussian (LoG), the difference of Gaussian (doG) and determinant of Hessian (doH) as unfeasible.

Finding segmentation contours was not necessary in this study, since the exposure time of the camera and flow rate is such that a leukocyte cell appears in a video frame as an elongated streak, due to the displacement over the camera exposure duration. The length of the elongated streak is proportional to its velocity. Morphological information about the cell cannot be reliably extracted due to this optical deformation occurring. As such, the objective of the image processing approach was to detect the leukocytes and determine their position in each video frame. The velocity and acceleration can then be calculated

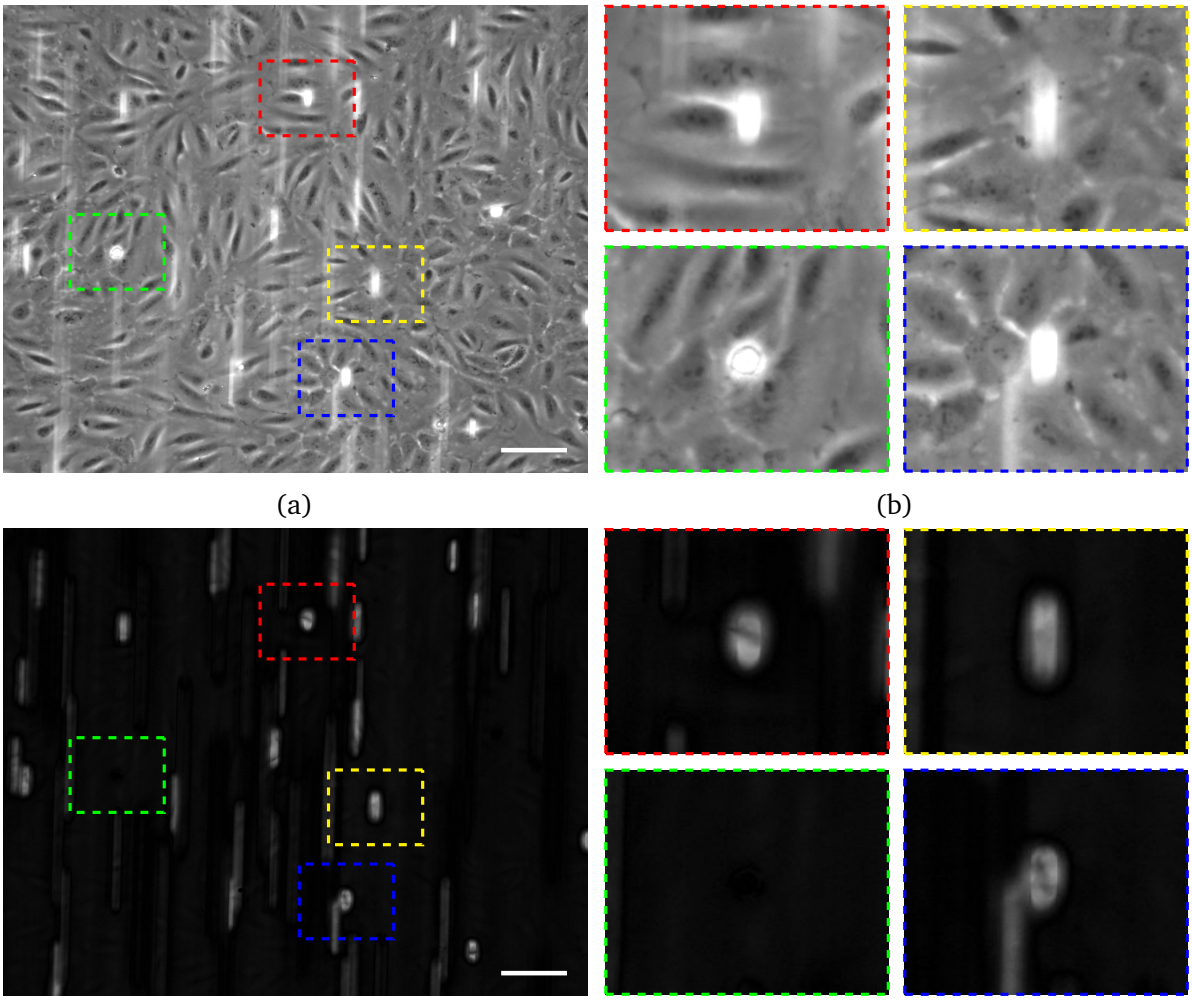

Fig. 4.2 Image (a) is a video frame obtained from a phase contrast microscope image sequence of leukocytes in an endothelial interactive adhesion assay, with four cells identified by coloured boxes in image (b). The magnified regions ($142 \times 114\,\mu$m) in (b) show leukocytes rolling over the endothelial monolayer, with four cells identified in red, yellow, green and blue boxes. An adhering leukocyte is seen in the green box. Image (c) shows the same frame as in (a) but after performing a contrast limited adaptive histogram equalisation (CLAHE) pre-processing step and background removal via frame differencing. In image (d) the corresponding magnified regions ($142 \times 114\,\mu$m) are shown. Note the absence of the adhering leukocyte in the green box; because it has been stationary for more than n frames it forms part of the background model. Scale bars: $100\,\mu$m.

from inter-frame positional changes. Detection can be viewed as a classification problem in which the task was to tell the presence or absence of a specific object in an image.

An approach to object detection that has been successful in many modalities is to use a Haar-features object detection classification cascade [114]. This is a machine learning based approach where a cascade function is trained from a lot of positive and negative images, and then used to detect objects in other images. Leukocyte detection and tracking using Haar-features object detection was performed using version 2.7 of the Python programming language, with version 2.4.10 of the OpenCV (Open Source Computer Vision) programming library. The video processing pipeline comprises three stages:

1. a background removal pre-processing step,
2. detecting leukocytes in each frame using the classifier, and
3. linking leukocytes between frames to form trajectories.

Prior to video analysis of a data set a classifier must also be trained to detect leukocytes; a single classifier can be used for analysis of multiple videos.

4.3.1 Pre-processing

A pre-processing step was performed on each frame of the video to compensate for illumination changes that may occur during the video, and any spatial illumination variation in the frame. Contrast limited adaptive histogram equalisation (CLAHE) enhances the image, and is effective even in image regions that are darker or lighter than most of the image [49]. The CLAHE algorithm pre-processing step was applied to each frame in the video sequence to enhance the contrast between leukocytes and their background. This step helps to ensure sufficient contrast difference between leukocytes and the background for detection of leukocytes.

Histogram equalisation is often used as an image enhancement method that alters the dynamic range and contrast of an image by spreading out the most frequent intensity values. This is achieved by using a cumulative distribution function as the mapping function. Intensity levels are changed such that the peaks of the histogram are stretched and the troughs are compressed [115]. Adaptive histogram equalisation differs from ordinary histogram equalisation in the respect that the adaptive method computes multiple histograms each corresponding to a different image region. The CLAHE algorithm works by dividing the image into smaller tiles, and applying a contrast limited local histogram

equalisation. Contrast limiting equalisation minimises the amplification of noise in regions of uniform low intensity. If a histogram bin is above the specified contrast limit, pixels in that region are clipped and distributed uniformly to other bins before applying histogram equalisation. Following equalisation, artefacts in tile borders are removed by bilinear interpolation [116].

4.3.2 Background removal

Identifying moving objects from a video sequence is a fundamental and critical task in video surveillance, traffic monitoring and analysis, human detection and tracking, and gesture recognition in human-machine interfaces [117]. In videos that have a constant or slowly changing background a common approach to aid in identifying moving objects is using frame differencing (see Figure 4.2c). Moving objects can be identified as regions within a frame that differ significantly from a background model, where the background model maybe static and predefined or slowly changing.

In leukocyte endothelial flow assays the background model typically changes only slowly during the recording. Fluid flow over the endothelial monolayer creates a shear stress at the endothelial surface. To minimise this shear stress endothelial cells adapt, by aligning and migrating in the direction of flow. The realignment of endothelial cells under flow is a two step process, involving the reorganisation of the cell cytoskeleton. Firstly, the endothelial cells contract and round up losing their original orientation occurring within minutes of flow being imposed. Secondly over longer time periods the cells spread and migrate in the direction of flow. Cell contraction and rounding is less pronounced in confluent endothelial cells [118].

The rounding and contraction of the endothelial background in THP-1 interaction assay recordings complicates background removal, since a background removal algorithm is required that adjusts to longer term background changes and is robust to changes in illumination. An initial attempt at removing the background was to create a mean image, from all frames in the sequence, and calculate the absolute difference for each pixel in the image and each frame of the video. This approach did indeed enhance the moving leukocytes over the background, however as the background is not fully static, subtle changes in background illumination over the image sequence created artefacts in the images. Artefacts from endothelial contraction, firmly adhering leukocytes and debris in the flow can cause artefacts in the background subtracted image.

The background removal approach was improved by calculating a running average background image for the duration of the video. For each frame (F) in the image sequence, a background image was calculated as the mean of pixels of (n) frames prior and proceeding the current frame. The value of a background frame (B) is calculated at each (x) and (y) coordinates for all time points (t) in the video as,

$$B(x,y,t) = \frac{1}{2n+1}\left(\sum_{i=1}^{n} F(x,y,t-i) + \sum_{i=0}^{n} F(x,y,t+i)\right). \tag{4.4}$$

The absolute difference between the current frame and the background image for that frame is then then calculated as:

$$\left|F(x,y,t) - B(x,y,t)\right|. \tag{4.5}$$

Taking the absolute difference between the original image and a rolling averaged mean image in Equation 4.5 was advantageous over a simple subtraction to expose illumination changes beyond the bit depth of the image by maintaining negative values. The resulting pixel values were normalized in the 8-bit range 0 to 255 to enhance the background difference.

A running average frame differencing method responds to gradual changes in endothelial morphology throughout the video duration and does not create significant image artefacts. The number of frames before and after the current frame, from which to calculate the average background frame is called n. Typically, an n of 50 was effective, this can be adjusted depending on the flow rate and frame rate of the setup. The algorithm creates a background model from the closest n frames, and for frames in the middle of the sequence uses n frames ahead and n frames behind. For frames at the start of the sequence the running average uses more frames that are after the current, and for frames at the end more frames prior to the current are used.

4.3.3 Object detection

In computer vision and image processing the task of identifying objects of a certain semantic class within a digital image is known as object detection. Face detection and pedestrian detection are especially well researched domains of object detection given their applications in photography, human computer interfaces, and surveillance.

Detection of leukocytes in THP-1 interaction assays can be approached as an object detection problem. The sometimes subtle phase variation of leukocytes makes their detection via thresholding non-trivial. Thresholding and edge-detection approaches are not suitable for detection of leukocytes due largely to the halo and shade-off effects introduced by phase contrast microscopy [113]. A machine learning approach implemented for leukocyte detection is capable of detecting objects quickly and accurately. This method was implemented based on a framework originally developed for real time face detection using a set of Haar-like features [119, 114], and is implemented in OpenCV.

Haar wavelet feature sets are a faster alternative to image intensity features, that have been used in object detection [120]. An adaption of Haar features is Haar-like features, which considers adjacent rectangular regions at a specific location in a detection window, sums up the pixel intensities in each region and calculates the difference between these sums. This difference is then used to categorise subsections of an image (see Figure 4.4).

The detection of leukocytes with a Haar-like features classifier involves three steps:

1. a data preparation step to gather image sets of positive examples (containing the object of interest in this case leukocytes) and negative samples (without the object of interest),

2. a classifier training step that uses the positive and negative samples to train a model for the object, and

3. a detection step to apply the model to detect the object in images.

The data preparation and classifier training only need to be performed once and are then applicable over multiple experiments.

Data preparation

A cascade of boosted classifiers working with Haar-like features was trained with positive and negative example images. A positive example is an image containing a large, centered instance of the object of interest, in this case leukocytes. A negative image is an arbitrary background region with the same dimensions as the positive example (see Figure 4.3). The size of each leukocyte in the video was dependent on its velocity, where leukocytes travelling at higher velocities, travel further during the exposure time of the camera and therefore appear longer. Positive images of leukocytes with a variance in intensity and orientation were selected to cover all stages of the leukocyte adhesion cascade.

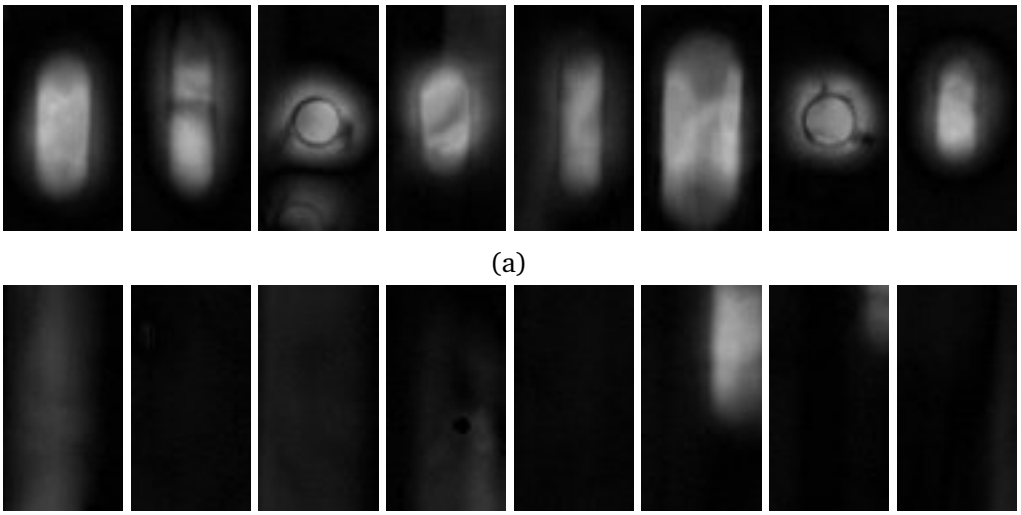

(a)

(b)

Fig. 4.3 A sample of positive and negative images used to train the object detection classifier. Training was performed on images where the background was removed via frame differencing. The halo artefact produces a distinctive ring around leukocytes. Images in (a) are positive samples containing a single leukocyte in each image. Images in (b) do not contain leukocytes and are negative background region images.

The image regions for positive and negative images have the same image dimensions, where positive samples (Figure 4.3a) are views containing a leukocyte. These regions were selected manually from background removed frames, and were selected to include a range of leukocytes in differing background illumination conditions and at a variety of stages in the adhesion cascade. Negative samples (Figure 4.3b) were selected as arbitrary random image regions in background removed video frames.

To acquire a set of negative images an ImageJ macro was written to extract a set of fixed size image regions from a video at random frames and positions, the logic of this process is shown in Listing 4.1. The generated set of images from random frames and regions of the video sequence were then manually pruned to exclude any examples that happen to contain positive leukocyte instances, either wholly or partially.

Listing 4.1 Generate negative samples

```
1 FOR number of negative imgaes
2   CALL select a random video frame
3   CALL create a rectangle: width (w), height (h), at random coordinates (x,y)
4   CALL duplicate selection
5   CALL save as image
6 END FOR
```

From the positive and negative samples a larger set of training sample objects was created using the opencv_createsamples utility. The utility provides functionality for generation of training datatsets, writing and viewing. Samples are created by applying perspective transformations and manipulations to a positive sample, to create more samples for training. Transformations and manipulations include: rotations around the x, y and z axis, the addition of white noise, and performing inversions. A positive sample image is then superimposed on a negative sample background and a binary output file with a *.vec extension is generated.

Listing 4.2 Create samples and merge vector files

```
1 find ./positive_images -iname "*.jpg" > positives.txt
2 find ./negative_images -iname "*.jpg" > negatives.txt
3 perl bin/createsamples.pl positives.txt negatives.txt samples 3000 "
      opencv_createsamples -bgcolor 0 -bgthresh 0 -maxxangle 0.05 -maxyangle 0.05
      maxzangle 0.0 -maxidev 10 -w 45 -h 85"
4 find ./samples -name '*.vec' > samples.txt
5 ./mergevec samples.txt samples.vec
```

In conjunction with the opencv_createsamples utility a method written as a Perl script applies the utility in a loop, (see Listing 4.2). For each positive image a series of transformations are applied with the opencv_createsamples utility. The output *.vec files are merged using the mergevec utility available under the MIT License written by Naotoshi Seo [1].

Classifier training

Object detection relies on Haar-wavelet features to perform a binary classification of regions within an image. Haar-wavelet features provide a compact representation of an object class and are computationally efficient as compared to working with pixel intensities or edges directly [120]. Haar-like features are arrangements of rectangular regions within a window, the difference between sums of these areas provides a weak descriptor that was used to categorise regions in the window. The descriptor indicates a certain characteristic of a particular area of the image. The combination of multiple features was used to indicate the existence or absence of certain image characteristics such as edges or texture changes.

[1] N. Seo, Tutorial: OpenCV haartraining (Rapid Object Detection With A Cascade of Boosted Classifiers Based on Haar-like Features), http://note.sonots.com/SciSoftware/haartraining.html, accessed 2015-01-09

Object detection was performed using a set of Haar-wavelet like features cascade classifiers as proposed by Viola et al. in their paper, 'Rapid Object Detection using a Boosted Cascade of Simple Features' [119]. Their approach employs machine learning techniques to train classifiers from positive and negative images. The generated classifier can then be used to detect objects in images.

This important work (Viola et al, 2001) introduces three major contributions in object detection making rapid real time object detection possible. Firstly, they introduce the use of *integral images*, which are essentially two-dimensional lookup tables. These *integral images* allow more efficient calculations of the sum of rectangular areas in the image, at any position or scale. Secondly, for fast classification a method of classifiers are constructed using a small number of critical features with minimal error rates from the large number available, this is provided by the AdaBoost algorithm [121]. Thirdly, they introduced a method which combines successively more complex classifiers in a cascade structure. The detection speed is significantly increased by focusing computational processing only on promising regions.

Training of a Haar-wavelet like object detector classifier for detection of leukocytes was achieved with the opencv_traincascade utility. The classifier was trained on the generated samples.vec file, as shown in code Listing 4.3.

Listing 4.3 Haar-wavelet cascasde classifier training

```
opencv_traincascade -data classifier -vec samples.vec -bg negatives.txt -numStages 20
    -minHitRate 0.999 -maxFalseAlarmRate 0.5 -numPos 2000 -numNeg 1000 -w 45 -h 85 -
    mode ALL -precalcValBufSize 1024 -precalcIdxBufSize 1024
```

The opencv_traincascade utility command line arguments pertain to: directory locations and settings, parameters for the cascade, boosted classifier parameters and Haar-like features parameters. The choice of parameters for classifier training was heavily guided by the findings of Lienhart et al, 'Empirical Analysis of Detection Cascades of Boosted Classifiers for Rapid Object Detection' [122]. In code Listing 4.3 the following command line arguments describe common locations and settings: *data* is the directory name to save the classifier file, *vec* the file name containing positive samples, *bg* contains a file list of images of background regions to superimpose the positive samples, *w* and *h* are the size of training samples in pixels, *n*umPos and *numNeg* determines the number of positive samples used in training at every classifier stage, *precalcValBufSize* and *precalcIdxBufSize* describe the size of buffer for precalculated feature indices. Making more memory available speeds up the training process.

Arguments describing the parameters for the cascade in code Listing 4.3 are: *numStages* determines the number of stages in the cascade to be trained. The structure of the cascade classifier is that of a degenerate decision tree, where a positive result from an initial classifier causes the evaluation of a second classification, and a third, up to a maximum number of stages determined by *numStages*. At any point a negative outcome would lead to termination and rejection of the region being classified. Training classifiers with more stages creates a more precise classifier but is more time consuming to train.

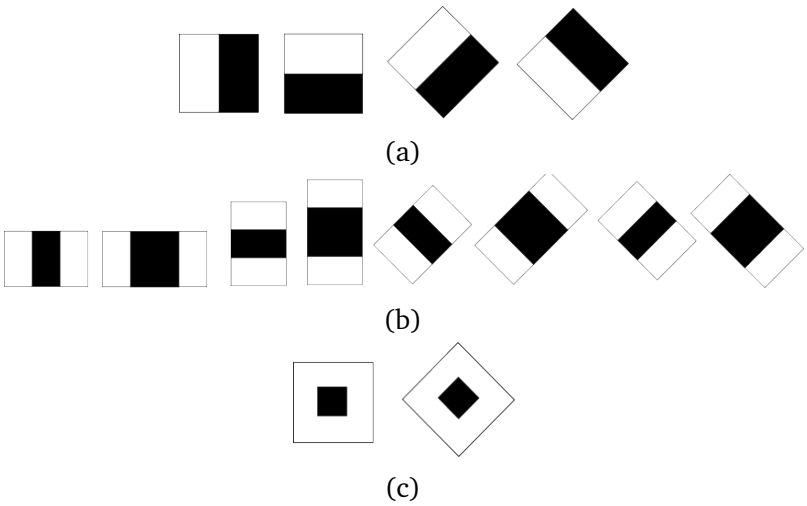

Fig. 4.4 Simple and extended set of Haar-like features. In (a), (b) and (c) features are described as line, edge and ceter surround features, respectively. Black areas have negative and white areas positive weights.

Boosted classifier parameters are: the *minHitRate*, which is the minimal desired hit rate for each stage of the classifier, and the *maxFalseAlarmRate*, which is the maximal desired false alarm rate for each stage of the classifier. The final parameter *mode* determines the Haar-feature set to be used in training, whether it should consist of the basic upright Haar features (Figure 4.4a) or should include extended Haar features, which are a set of features rotated at a 45° angle [114] (Figure 4.4a, Figure 4.4b and Figure 4.4c).

Training of the classifier is a computationally intensive process and depending on the parameter setup, object detection scenario and computer hardware can take several days to complete. Upon completion of training an extensible markup language (XML) classifier file is generated, which can then be used for object detection.

Leukocyte detection

The trained classifier can be used to search for objects over a range of scales in video frames or images. This was achieved with the OpenCV detectMultiScale function, that applies the classifier at multiple scales to search for the object of interest. The function returns an array of detected regions as rectangles with x, y coordinates and width and height.

The detectMultiScale function performs detection with a pixel-by-pixel sliding search window, which traverses the image and searches for objects in each region. For each image, the region within the window will go through the cascade classifier, to check the probability of an object being present. There are two mechanisms for scale-invariant object detection; the first is to resize the search window, the second method is to scale the image. Multiple image rescaling creates an image pyramid, by downsampling the image using neighbouring pixels [123].

For computational efficiency the image rescaling method was used in detection of leukocytes. In the leukocyte flow assay experimental scenario, the size of a leukocyte at the microscope magnification is fairly consistent, so only a small image rescaling is needed. In more broadly applicable tasks of object detection such as face detection, a greater size variance is required.

The cascade classifier is comprised of multiple filter stages, which are gone through at each position of the sliding window. If the search window region does not pass the threshold of a stage, the cascade classifier will reject the region as a leukocyte. If the search window region passes all stages of the classifier successfully it is classified as a candidate leukocyte, which pending further checks maybe classified as a leukocyte. The cascade filter is advantageous in the detection phase because it reduces the computational workload by rejecting regions at early stages in the cascade.

Listing 4.4 Pseudocode for multi-scale object detection

```
FOR number of image scales
  downsample by one scale
  compute integral image for current scale
  FOR each shift step of the sliding detection window do
    FOR each stage in the cascade classifier do
      FOR each filter in the stage do
        filter the detection window
      END FOR
      accumulate filter outputs within this stage
      IF accumulation fails to pass per-stage threshold do
        break the for-loop and reject this window as a leukocyte
```

```
12        END FOR
13      END FOR
14      IF this detection window passes all per-stage thresholds do
15        accept this window as a leukocyte
16      ELSE
17        reject this window as a leukocyte
18      END FOR
19    END FOR
20  END FOR
```

The pseudocode in Listing 4.4 describes the process of multiple size object detection, the implementation of this routine is shown as Listing 4.5, as implemented in OpenCV Python. In Listing 4.5 the parameters are largely self explanatory; a maximum and minimum object size can be selected. Object regions larger or smaller than these values are ignored. The *scaleFactor* parameter specifies by how much the image size is reduced at each image scale. Of particular importance is the *minNeighbors* parameter, this determines the object detection sensitivity. The *minNeighbours* parameter is the minimum number of neighbours required for each candidate region in order that it is retrieved. Objects are usually detected at multiple scales and near other detected objects, creating neighbourhood positive regions. By increasing the *minNeighbours* number, the number of false positives can be reduced, however this can have the consequence of increasing the number of false negatives. Finally the flag CV_HAAR_SCALE_IMAGE, tells the routine to rescale the image rather than the search window size.

Listing 4.5 OpenCV python detectMultiScale function

```
cv2.CascadeClassifier(cascPath).detectMultiScale(frame,scaleFactor=
    detectParticlesScaleFactor,minNeighbors=detectParticlesMinNeighbors,minSize=
    detectParticlesMinSize,maxSize=detectParticlesMaxSize,flags = cv2.cv.
    CV_HAAR_SCALE_IMAGE)
```

The code in Listing 4.5 produces an array of found objects as displayed in Table 4.1. This gives for each particle, the frame number in which it was detected, the x and y coordinates of the upper left corner of the detection region, the width and height of the detection region and the x and y centroid of the detection region. The centroid values were calculated using the rectangle x and y coordinates along with width and height.

Table 4.1 Example results table of objects detected by detectMultiScale

frame_no	x	y	w	h	x_centroid	y_centroid
1	351	42	48	91	375.0	87.5
1	1043	53	50	94	1068.0	100.0
1	867	870	50	95	892.0	917.5
2	24	288	49	92	48.5	334.0
3	1121	18	50	95	1146.0	65.5
⋮	⋮	⋮	⋮	⋮	⋮	⋮

4.3.4 Tracking leukocytes

The distribution of velocities of leukocytes flowing over an endothelial monolayer can be used to study interaction events associated with an inflammatory response. Determination of leukocyte velocities requires detection of leukocyte positions in each video frame (as described in subsection 4.3.3). In addition to this, a method was needed to link leukocytes over multiple frames. Tracking of leukocytes over multiple consecutive video frames was used to collect trajectory data, including the instantaneous velocity and acceleration, which are essential to the study of leukocyte interaction events. The linking of leukocytes into trajectories was achieved by a particle tracking algorithm.

Particle tracking algorithms analyse sequential video frames and record the movement of particles or objects over those frames. This has applications in human-computer interaction, security, surveillance, traffic control and medical imaging. Two commonly used algorithms are meanshift and camshift tracking. These methods use an iterative localisation procedure based on the maximisation of a similarity measure [124].

In a biological context particle tracking tends to focus on scenarios where the motion between frames is effectively Brownian. Two common software packages for particle tracking of particles with a Brownian like motion are *TrackMate* in ImageJ and *TrackPy* in Python. These were found to be suboptimal when applied to tracking leukocytes in an interaction assay where the experimental protocol ensures a linear flow direction.

A specialised tracking protocol was created and here described that exploits the unidirectional movement of leukocytes between frames in the direction of fluid flow. Leukocyte tracking was a video post-processing step, which was applied following detection of leukocyte positions in each video frame. Linking over multiple frames was performed by assigning particle ID values over multiple frames.

Particle linking

Each object or particle detected represents one row in Table 4.1. These particles are linked by iterating through each row in Table 4.1 and searching for corresponding particles in later frames. Particle linking works on the premise that particles move in the direction of flow, and therefore particles in later frames will be within a narrow range of x values and will have y values greater than the original, as illustrated in Figure 4.5. Particles are linked over sequential frames according to the minimal displacement in y value, particle IDs are then matched to form a trajectory.

The logic for particle linking is shown succinctly in Listing 4.6. In the first frame of the video sequence with detected objects present, the objects are assigned a sequential numerical ID depending on their y-coordinate value. From the initial setup with IDs assigned in one frame, numerical IDs can be assigned to all other leukocytes in the data frame.

Listing 4.6 Particle linking

```
ASSIGN ID to objects in frame with first detected object by y value
FOR each detected object
  IF particle_ID == NULL
    GET max particle_ID in particles table
    ASSIGN particle_ID as max(particle_ID) + 1
    GET particles in lookahead frames within x range and larger y value
    IF particles are found
      ASSIGN particle_ID as max(particle_ID) + 1
  ELSE particle_ID != NULL
    GET particles in lookahead frames within x range and larger y value
    IF particles are found
      ASSIGN particle_ID
END FOR
```

To summarise Listing 4.6, for each detected object, if it does not have an ID assigned, it is assigned the next highest particle ID, following this a function looks in proceeding frames for leukocytes to link by assigning the same ID. If the leukocyte has an ID already assigned, then proceeding frames will be searched and leukocytes linked by their ID. In this way after iterating through all the rows, every detected leukocyte in every frame of the video has an ID assigned to it. A parameter specifies how many frames to look ahead when linking particles into trajectories, typically two frames are searched. A lookahead two frames is useful if the detection phase fails in one frame then the trajectory can still be found.

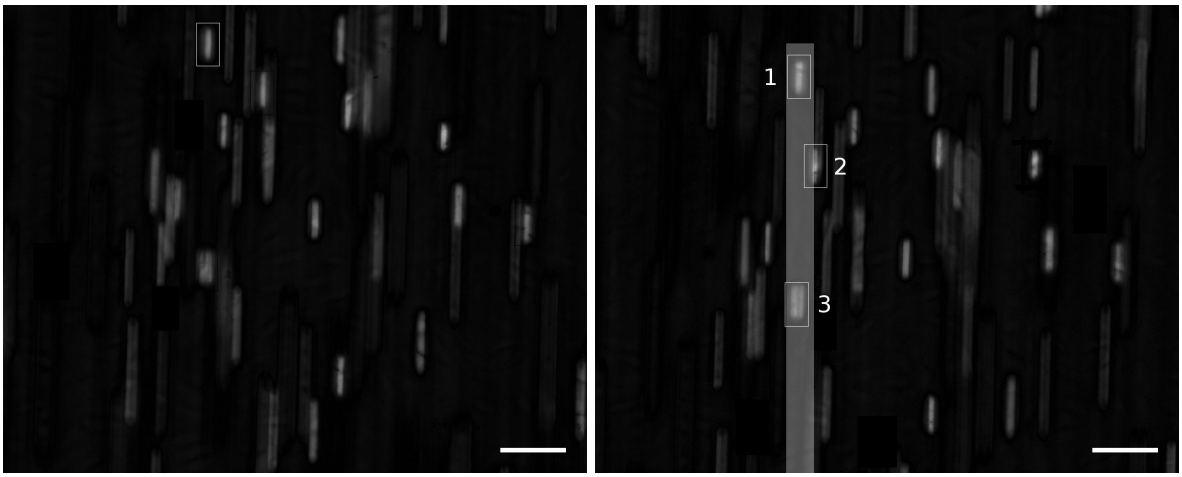

(a) (b)

Fig. 4.5 An example of a detected leukocyte particle being linked over two consecutive frames. Image (a) shows the original detected leukocyte particle to be linked in the first frame. Image (b) is the next consecutive frame, the search area is highlighted as the grey region and the 3 candidate particles for linking are annotated with a box. These candidate particles are ranked in order of their distance from the original particle and linking performed for the closest particle to the orignal particle. Scale bars: 100 μm.

An example of particle linking is shown in Figure 4.5. A leukocyte is detected, the white box in Figure 4.5a and the proceeding frame searched within a range x values. An x coordinate range was typically chosen to correspond to half the width of the search window or $\pm$ 30 pixels. The y coordinate range searched was restricted to a range greater than the y centroid of the original particle. The search area can be seen as the grey shaded area in Figure 4.5b, in which three particles were found. These were ranked in order of their distance from the original particle, the closest particle is then selected to be linked with the particle in the earlier frame. This method does not make any assumptions about particle velocity between frames, and therefore is capable of tracking particles that are accelerating, as they interact with the endothelial monolayer.

Trajectories

Following particle linking trajectories can be aggregated, cleaned, and calculations performed. A trajectory cleaning stipulation is to enforce a minimum trajectory length, removing detected objects that appear in less than n number of frames. This removes very short trajectories that may exist for example as incorrectly detected false positives being linked together. A minimum trajectory length of 10 frames has been used to remove short

trajectories. An additional measure to clean trajectories is to remove trajectories appearing in the first and last frames of the video sequence since they are not complete and may continue beyond the duration of the recording.

An annotated output of each video frame with coloured boxes showing trajectories with particle ID numbers was created as a quality control step and to allow cross referencing with the output results table. This can be seen in Figure 4.6a and Figure 4.6b. Following the linking of particles into trajectories, information about the trajectories was calculated. As the x and y displacement of a particle linked over consecutive frames is known, using Pythagoras's theorem the Cartesian displacement was calculated. From this and knowing the frame rate of the camera the instantaneous velocity and acceleration of the particle was calculated.

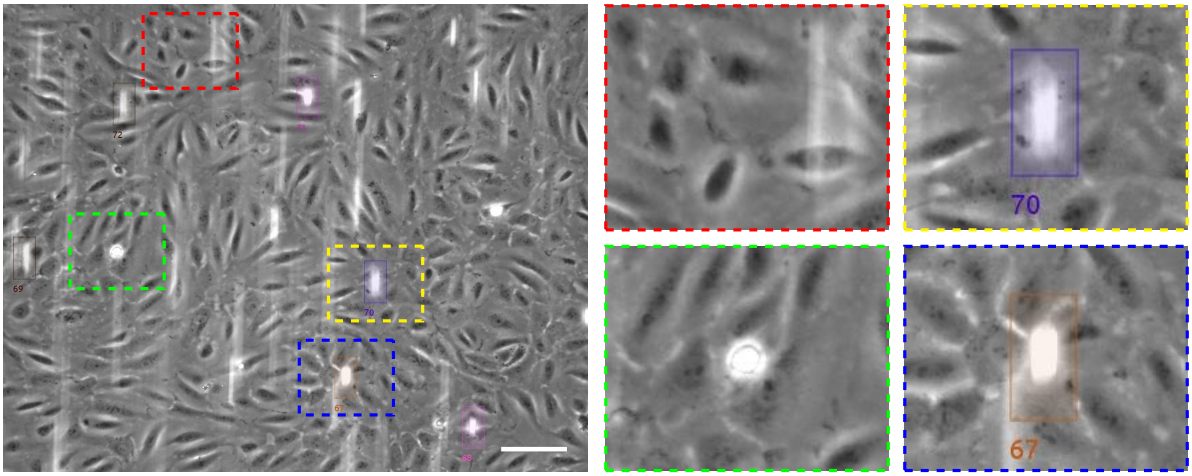

(a) (b)

Fig. 4.6 Image (a) is a video frame obtained from a phase contrast microscope image sequence of leukocytes in an endothelial interactive adhesion assay with the object detection and tracking protocol applied. The frame has been annotated with trajectories displayed in a coloured box with a number. In the magnified regions ($142 \times 114\,\mu m$) in (b) leukocytes rolling over the endothelial monolayer can be seen in the red, yellow and blue boxes, an adhering leukocyte is seen in the green box. Scale bars: $100\,\mu m$.

4.4 Performance evaluation

Two approaches were taken to establish the effectiveness of the leukocyte detection and tracking methodology described in section 4.3. Firstly, an assessment of the leukocyte object classification scheme was performed. This approach was then used to perform parameter optimisation of the object detection. Further to this, an assessment of the accuracy of the tracking protocol was performed with reference to a manual set of tracked leukocytes.

4.4.1 Object detection

The object detection method described in subsection 4.3.3 was validated with reference to a gold standard set of detected leukocytes. The gold standard reference set was generated from 90 frames, selected at random from 12 videos and two experiments. In each frame the x and y coordinate centroids of leukocyte objects were manually recorded, giving a total of 1078 leukocytes.

Selection of objects appropriate for classification is not trivial, since leukocytes travel at a range of velocities, and their profile or streak length is dependent on their velocity. The scope of this detection and tracking algorithm was to evaluate the slower leukocytes, which may be interacting with the endothelial monolayer. Higher velocity trajectories, may only exist in several frames and their detection and tracking would be more error prone. Making a judgement as to what length of object was appropriate for classification is difficult. The classifier was trained with objects in images of size 45 pixels by 85 pixels, this was used as a guide to the maximum object size in the gold standard set.

Leukocytes partially in frame or within one detection window of the frame edge were not labelled since the classifier requires the full object and background for detection. It could therefore be argued that this performance evaluation may overestimate the recall value. However, the detection of leukocyte in the first or final frames of the trajectory are less significant than consistent detection of frames in the middle of a trajectory.

Having generated a gold standard reference coordinate set, the performance of the automated Haar-like features object detection was compared relative to the gold standard. Evaluation of the binary classifier involves classification of true positives, false positives and false negatives, of the automated coordinate set relative to the gold standard coordinate set. This was done for each frame of the 90 frames in the reference set, and an annotated image set produced, an example of which is shown in Figure 4.7a. A range of classification errors can occur depending on the leukocytes present in the local image region being analysed.

Faint longer streaks may be incorrectly identified as leukocytes, this is demonstrated by the blue rectangle in Figure 4.7b. Alternatively, leukocytes with unusual intensity variations and overlapping with other foreground objects can lead to false negatives, as shown by the red rectangle of Figure 4.7b.

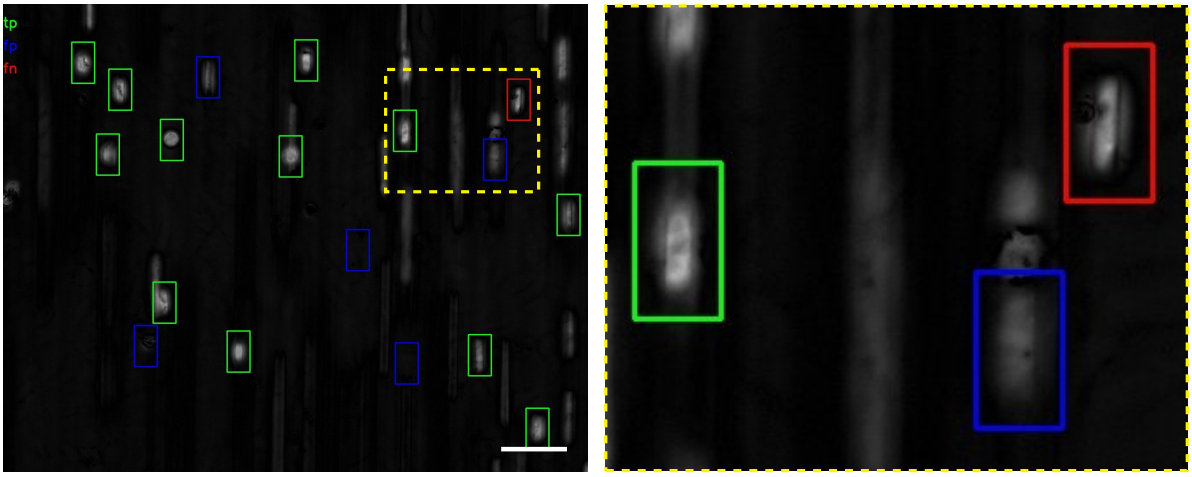

(a) (b)

Fig. 4.7 Image (a) is an example annotated video frame, that demonstrates the evaluation of the automated object detection using Haar-like features. With reference to the hand annotated gold standard objects within green rectangles are true positives, blue rectangles are false positives, and red rectangles false negatives. The magnified region ($232 \times 184\,\mu m$) in (b) shows a region containing a true positive, false positive and false negative. Scale bars: $100\,\mu m$.

Determining true positives, false positives and false negatives was achieved with a script to correlate particles found by the Haar-like features object detection to the gold standard classification for each frame. Code logic is shown in Listing 4.7. In each of the 90 validation frames the labelled gold standard coordinate set was given an initial classification of false negatives. Following this, the automated object detection was run, generating a second set of object coordinates. The generated coordinate set are all initially labelled as false positives. A default state of false positives in the automated set, and false negatives in the gold standard set was given. True positives are found by searching for corresponding coordinates between the automated and gold standard sets. For each coordinate position in the automated leukocyte coordinates, if a corresponding object can be found in the table of gold standard coordinates then it was labelled as a true positive, and the object in the gold standard table removed. Objects are matched if their centroid coordinates are within the same detection window size area. The final stage, having labelled false positives, false negatives and true positives was to concatenate the automatically detected coordinate set

with the gold standard set, the tables are then sorted according to their y values. From the number of true positives, false positives and false negatives the true positive rate or sensitivity and the true negative rate or specificity can be calculated.

Listing 4.7 Object detection performance relative to gold standard

```
FOR each frame
  GET gold_standard coordinates
  LABEL all gold_standard_objects as FN
  RUN automated object detection
  LABEL all automated_objects as FP
  FOR each row in automated_objects
    IF particle exists +/- w/2 and +/- h/2 of any gold_standard_objects
      LABEL object as TP
      DELETE row from gold_standard_objects table
  END FOR
  CONCAT gold_standard_objects and automated_objects
  SORT by y-coordinate
END FOR
```

An important parameter to optimise in Haar-like features object detection is the *minNeighbors* parameter, described in subsection 4.3.3. This parameter specifies how many neighbours each candidate rectangle should have to retain it, thereby acting as a discrimination threshold. The choice of a *minNeighbors* parameter balances the recall and the precision rate of the classification system. To find the optimal choice of *minNeighbrs* parameter the validation protocol in Listing 4.7, was run at every integer parameter value between 0 and 300. The mean precision and recall over the 90 image and 1078 leukocyte validation image set was plotted in Figure 4.8b, along with the F1 score. A maximal F1 score was achieved with a *minNeighbors* parameter of 160. However, a value of 140 was was used in video analysis to favour a higher recall than precision, since for trajectory tracking high recall is more important than precision. A false negative is more problematic than a false positive, because a false negative can result in a trajectory being broken. On the other hand, false positives usually do not correlate with any particles in proceeding frames so are expunged from the data set during trajectory linking. Using a *minNeighbors* parameter value of 140 a recall of 0.90 and precision of 0.81 was achieved. This value is displayed as a dashed line in Figure 4.8a, and a red dot in Figure 4.8b.

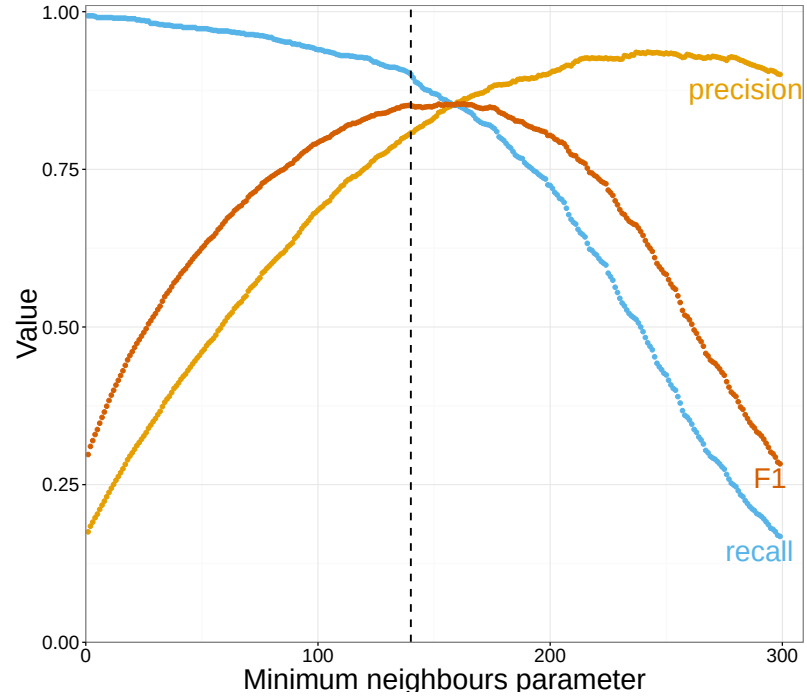

(a)

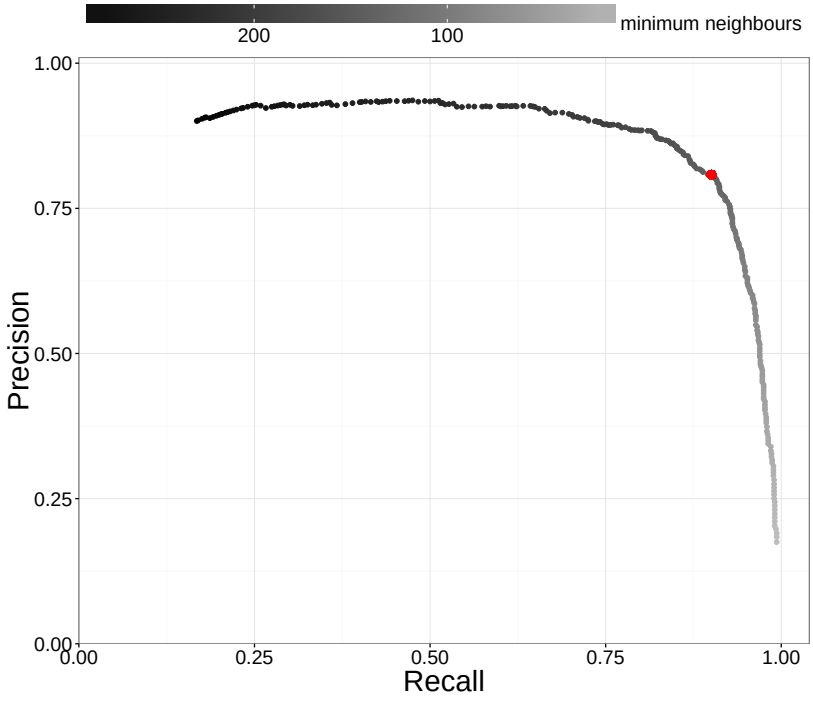

(b)

Fig. 4.8 Plot (a) shows the precision, recall and F1 curves of the automated object classifier over a 90 image validation set, for the *minNeighbours* parameter in a range from 0 to 300. The dashed line is the selected *minNeighbors* parameter of 140. The precision-recall curve in (b) shows the relationship between precision and recall over the parameter range. The red dot is the chosen parameter value of 140.

4.4.2 Tracking

A comprehensive evaluation of leukocyte endothelial interaction assay analysis should include a systematic performance evaluation of the tracking approach. Systems tracking single objects can be evaluated with a straightforward benchmarking approach. However, for multiple object tracking systems as in leukocyte endothelial interaction assays, there is no generally agreed upon evaluation procedure.

Bernardin et al., 2008 [125] present a method to detect the basic types of errors produced by multiple object trackers. They introduce two metrics, multiple object tracking precision (MOTP) and multiple object tracking accuracy (MOTA). These metrics express the tracking procedures ability to estimate object positions and to follow objects consistently over time.

The multiple object tracking precision (MOTP), was calculated as:

$$MOTP = \frac{\sum_{i,t} d_t^i}{\sum_t c_t}, \tag{4.6}$$

where, (d) is the distance between a matched pair of automatically tracked and gold standard particles, and (c) is the count or the number of matched objects. This metric represents the ability of the tracker to precisely find object positions, calculated as the average error in position between the automated classifier and a gold standard set. To calculate this 864 leukocyte positions were recorded in 200 consecutive video frames. For each matched gold standard and automated leukocyte pair the euclidean distance between the gold standard and automated position was calculated. A MOTP value of 3.5 pixels was achieved from this, showing a high degree of precision in estimating leukocyte positions.

A second metric the multiple object tracking accuracy (MOTA), was calculated from:

$$MOTA = 1 - \frac{\sum_t (m_t + fp_t + mme_t)}{\sum_t g_t}, \tag{4.7}$$

where, (m_t), (fp_t) and (mme_t) are the number of misses, of false positives and of mismatches, respectively, for time t. The value g is the sum of the number of trajectories. A miss is an object that the tracking method failed to record, a false positive is an incorrectly tracked particle, that is not actually an object and a mismatch is a particle that is incorrectly reassigned a new ID whilst tracking. These three basic error types are summed to give an error ratio and the resulting tracking accuracy.

To calculate a MOTA value the number of objects, misses, false positives and mismatches were counted in each frame of 200 consecutive video frames. A total of 846 objects were recorded, with 51 misses, 13 false positives and 0 mismatches and a MOTA value of 0.92 was achieved.

A high MOTP and MOTA suggest that the detection and tracking methodology presented in section 4.3 is well suited for the task of tracking leukocytes. The MOTP and MOTA values of 3.5 pixels and 0.92 were achieved from a video of unstimulated endothelial cells. The MOTA value may be reduced in videos with more leukocyte endothelial interaction, leading to more slow particles and therefore greater probability of occlusion of leukocytes.

4.5 Results

Presented here are results from an experiment intended to demonstrate the effectiveness of the automated leukocyte tracking approach, as well as exhibiting some ways in which the leukocyte trajectories data can be visualised. For this experiment three 4 min videos were collected using the methods described in section 4.2. The first video acted as a control, and was acquired capturing frames from a flow chamber with leukocytes and an unstimulated monolayer of HUVECs, the second video was acquired from a HUVEC monolayer stimulated for 5 min with histamine, and a third video was acquired from HUVECs that had undergone a period of 5 min of stimulation with PMA.

In the control video 3475 trajectories were recorded, 3873 trajectories were recorded in the PMA stimulated video, and in the histamine stimulated video 3288 trajecories were recorded. Using Equation 4.3 a critical rolling velocity of $215\,\mu m\,s^{-1}$ was calculated for the experiment. Trajectories containing leukocytes with instantaneous velocities beneath this critical velocity were defined as rolling or interacting significantly. According to this definition 143 leukocytes were rolling in the control video, 195 rolling in the PMA video, and 211 leukocytes were rolling in the video of histamine stimulated HUVECs (shown in Figure 4.9a). This output is similar to what is measured in a manual analysis, except the definition of a rolling or interacting leukocyte has been more rigorously codified, using a calculation of the critical velocity and the instantaneous leukocyte velocity measurements.

This increase in trajectories containing rolling leukocytes can be observed by the difference in the distribution of the minimum velocity in a trajectory from Figure 4.9b. The initial bump in the slowest trajectories for the PMA and histamine groups, represent the slowest interacting leukocytes. In prior manual analysis this data subset was the only quantified output from these experiments. In addition to the minimum trajectory velocities, in Figure 4.9c the distribution of trajectory median velocities for each group is plotted. The treatment displaying the greatest velocity reduction is the histamine stimulated group.

In Figure 4.9d the endothelial response over time can be assessed, this normalised plot shows how the number of slow rolling (beneath the critical velocity) leukocytes changes throughout the duration of the three videos. Histamine has the fastest response leading to rolling earlier in the video sequence, whilst PMA seems to lead to increased rolling at a rate similar to unstimulated endothelial cells.

An informative way to visualise the tracking data from the leukocyte interaction assays was to render interactive three dimensional plots, as is shown in Figure 4.10. These visualisations give useful insight into the experimental findings and can be used to investigate the

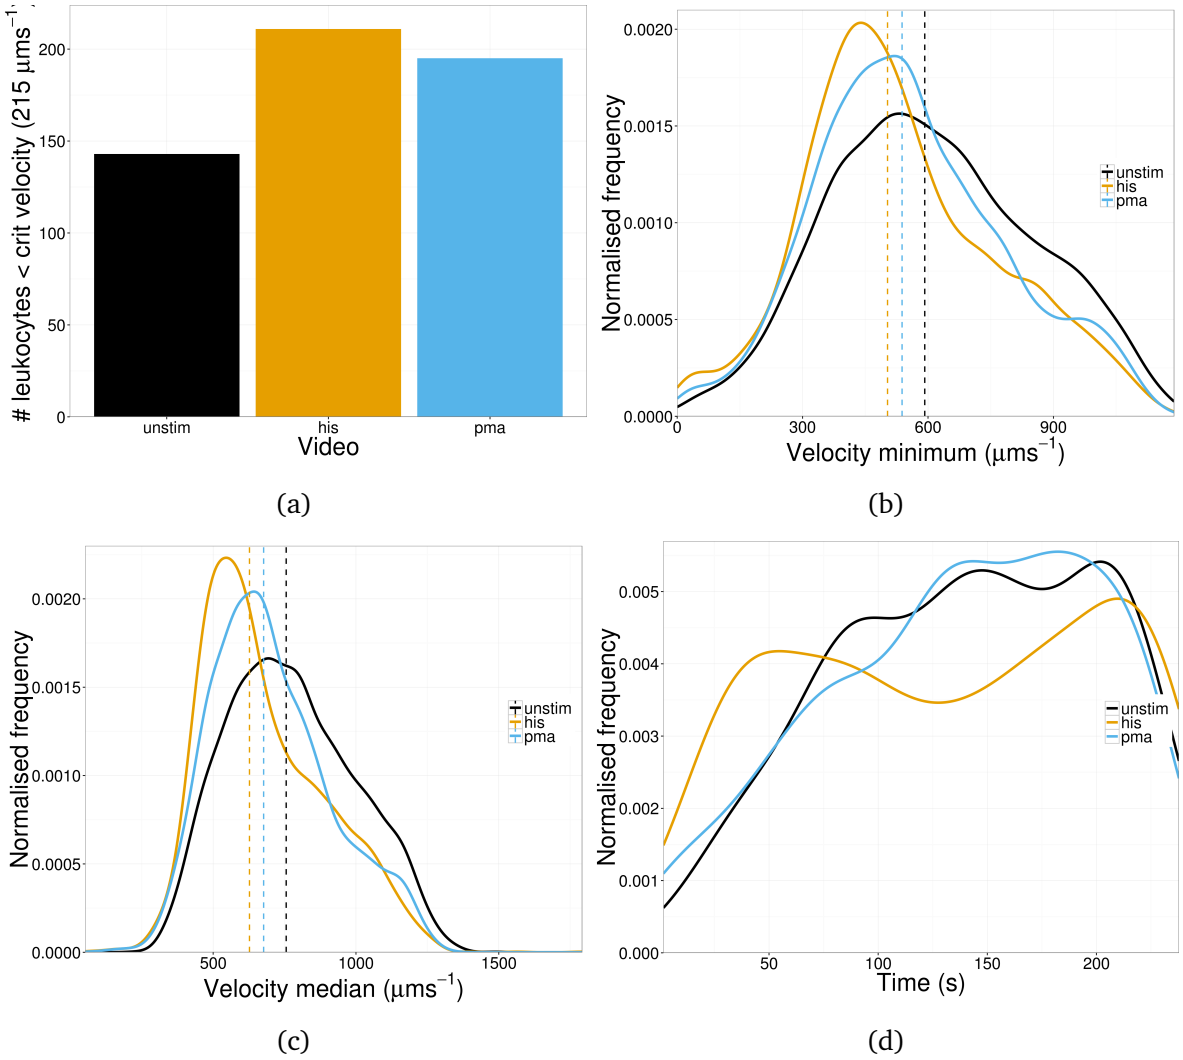

Fig. 4.9 Plot (a) is a bar graph of the number of leukocytes beneath the calculated critical rolling velocity ($215\,\mu m\,s^{-1}$) for the three video sequences. In plot (b) the distribution of minimum leukocyte trajectories is shown for the three treatment groups, where the dashed line is the median for each distribution. Similarly plot (c) is a normalised density plot of the distribution of median velocities for trajectories from the three treatment groups, with the dashed lines representing the median for each distribution. Plot (d) is a normalised density plot of the number of leukocytes beneath the critical rolling velocity ($215\,\mu m\,s^{-1}$) plotted over time.

accuracy of the detection and tracking method. The *z-axis* plots the frame number or time of the video starting at frame 0 up to frame 2000. In Figure 4.10a we can see that fewer leukocytes were tracked in total and that the majority were tracked after the first ∼500 frames. In Figure 4.10b more rolling leukocytes were tracked, and perhaps endothelial cells to the right of the frame interacted more with leukocytes throughout the duration of the recording. A variant of Figure 4.10 that was also informative used a colour gradient to display the instantaneous velocity of each of the leukocytes across the trajectories.

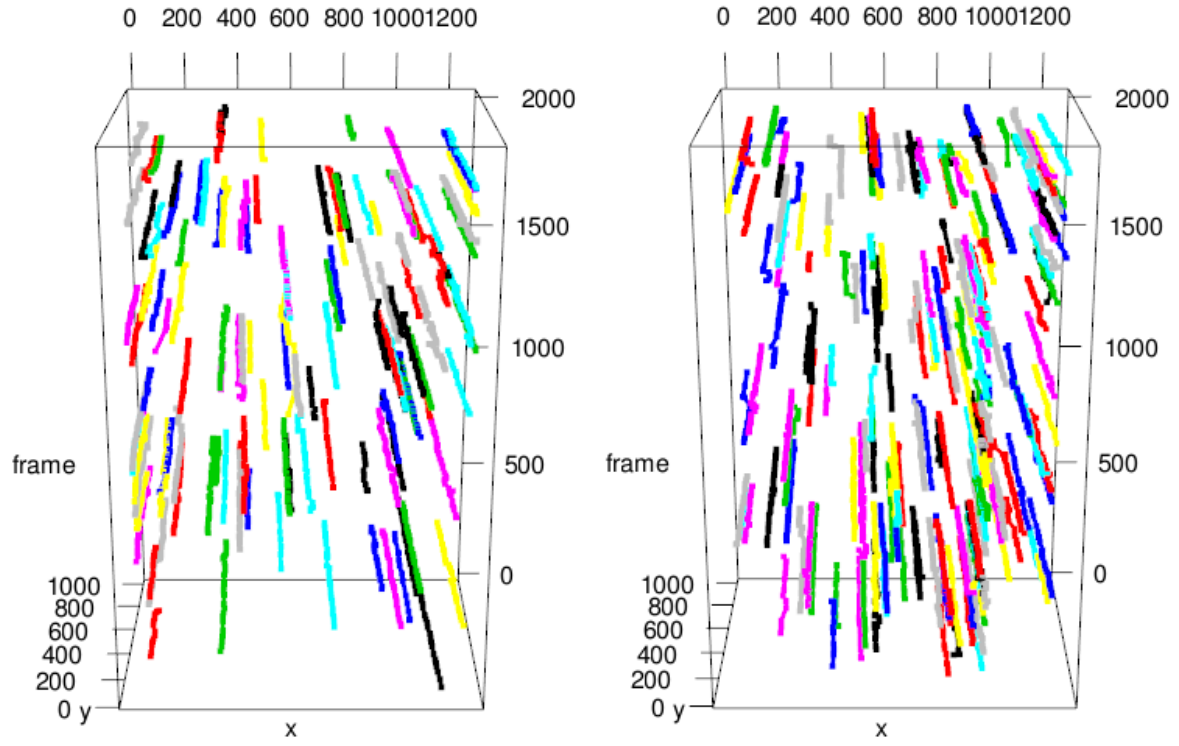

Fig. 4.10 Plots showing the tracking of rolling leukocytes throughout the duration of the acquired videos. The *x-axis* and *y-axis* are Cartesian coordinates in the frame, and the *z-axis* is the frame number, which is analogous to time. Each trajectory is mapped in a random colour to better distinguish its path. Plot (a) shows the rolling leukocytes tracked from the control video and plot (b) shows the rolling leukocytes tracked from the histamine stimulated HUVEC video.

4.6 Summary

This chapter presented a method to detect and track leukocytes in interactive adhesion assays using Haar-like feature object detection and a custom tracking algorithm. In section 4.1 an introduction is given, describing the physiological relevance of shear stress and how this is calculated in an *in vitro* flow assay, along with a description of methods of quantitative analysis, and finally an explanation and formula to calculate the maximum velocity of rolling leukocytes. In section 4.2 the experimental method used to perform these *in vitro* flow assays is described.

In section 4.3 the processing of acquired videos is described from pre-processing to enhance contrast and remove background, to feature detection, linking and leukocyte tracking.

In section 4.4 a comprehensive evaluation of the leukocyte tracking method is described. The Haar-like feature object detection method is evaluated by comparing to a hand-labelled gold standard and the accuracy measured. Parameter optimisation is performed with reference to this gold standard dataset. Leukocyte tracking was evaluated using two metrics described by Bernardin et al., 2008 [125]. The MOTP assessed the spatial precision of leukocyte detection, and an MOTP value of 3.5 pixels was found. The accuracy of the tracking protocol was evaluated using the MOTA metric and a value of 0.92 found.

In section 4.5 the results of a simple experiment provided further validation of the leukocyte detection and tracking methodology. Results from an experiment were presented comparing leukocyte interaction with unstimulated endothelial cells to interaction with endothelial cells that had undergone a period of stimulation either with histamine or PMA. The results corroborate a manual analysis and fit with the biological hypothesis. In this section some examples of data visualisation are presented, although there are potentially many more ways to represent data from this rich quantitative output.

The methodology presented in this chapter represents a precise and accurate way to hugely increase the quantitative yield from interactive adhesion flow assays with leukocytes, and can therefore be applied to give higher level biological insight. To take this work further it would be useful to acquire an experimental control to distinguish between interacting leukocytes and non-interacting leukocytes in the slowest laminar flow stream. An attempt at this was performed in an experiment without an endothelial monolayer in a parallel flow chamber. However, removing the endothelial layer altered the height of the profile as well as the frictional coefficient of the surface so this could not be reliably used as a control. An alternative approach could be to record a video using non-interacting beads instead

of leukocytes and observe the distribution of bead velocities compared to leukocytes. The slowest population of beads could then be used as a guide of the critical velocity or the slowest non-interacting leukocyte velocity.

Chapter 5

Conclusions and future work

Scientific endeavour increasingly requires analytical methods combining scientific inquiry, statistical knowledge, and computer programming. Biological insight in cell biology is increasingly reliant on computerised techniques along with advances in microscopy and data acquisition. This body of work has sought to utilise techniques in image processing, data mining, statistics, and machine learning to further advance understanding of cardiovascular processes.

Three methods have been introduced in this thesis to facilitate or enhance the output from established methods in cardiovascular experimentation. The three automated methods described in chapter 2, chapter 3, and chapter 4 extract quantitative data from microscopy images. Each of these methods has been tested in a biological context, and the performance validated. Two papers have thus far been published using the methods described in this thesis. The methods described in chapter 2 contributed towards Ferraro et al., 2016: 'Weibel-Palade body size modulates the adhesive activity of its von Willebrand Factor cargo in cultured endothelial cells' [126]. The methods described in chapter 4 contributed towards Westmoreland et al., 2016, 'Super-resolution microscopy as a potential approach to platelet granule disorder diagnosis' [37].

This chapter contains some concluding remarks on each of the methods introduced in this thesis, with additional statements about their significance and their continued usage. Finally areas for potential improvement and avenues for further work are mentioned.

High-throughput morphometric analyses of endothelial organelles

In chapter 2 methods were presented to analyse microscopy images acquired in high-throughput of immunostained endothelial cells and their organelles. Segmentation of

endothelial cells and assignment of organelles to these cells enables both whole population level and cellular level analysis. Where the study of endothelial cells and WPBs under different physiological and pathophysiological conditions can be used to gain an understanding of underlying haemostatic processes. Additionally, vWF exocytic pores have been extensively studied with this technique, providing a method that is complementary and potentially more sensitive than the commonly used ELISA method.

The method described here has contributed to several articles in preparation for publication and is frequently used to analyse data sets. Its routine usage should continue to be beneficial to members of the lab. Its usage could in the future be expanded further to include segmentation methods for more endothelial organelles. In addition the semi-automated quality control step could be improved to be fully automated, thus reducing the workload.

Image analysis for platelet dense-granule deficiency diagnostics

An automated image analysis pipeline discussed in chapter 3 provides a quantitative, unbiased, rapidly acquired data set, that forms the basis for a platelet dense-granule disorder diagnostic tool. This proof-of-principle study described in chapter 3 has huge potential for further study and development. The effectiveness of this method has been demonstrated in Westmoreland et al., 2016 [37].

The described method allows platelets to be fixed and sent to a central facility for diagnosis and re-analysed and imaged at a later date if required. This would relieve pressure on specialised facilities that are dependent on functional analyses that require measurements to be taken within hours of the blood being collected, and also relieve pressure on the patient who is not required to travel to a specialised facility.

The method could also be extended with the addition of fluorescent markers to label additional platelet structures. These markers could for example include α-granules, vWF, P-Selectin, and serotonin. The segmentation and feature extraction of multiply labelled platelets would provide an increasingly rich feature set to improve the diagnostic ability. If more samples were acquired and imaged from healthy volunteers and patients with a multitude of platelet disorders, large data sets could be built up. This would provide fertile ground for more sophisticated and specific diagnostic tools. For example supervised machine learning methods could be used to model subtle disease phenotypes, and then used to classify platelets according to the phenotype represented. This could have the additional benefit of being able to identify milder HPS phenotypes, as there is no method

currently available that allows alpha and dense-granule to be measured in high numbers from the same sample.

Automated detection and tracking of leukocytes *in vitro*

The technique presented in chapter 4 for improving the quantitative yield from leukocyte interactive flow assays has not yet been extensively used, but it represents a promising analytical tool. It has been shown to be highly accurate at detecting and tracking leukocytes over multiple frames. However, it does have some limitations. The object classification has difficulty when particles are very close or overlapping. This problem can be alleviated by reducing the flux of leukocytes and increasing the video length.

An improvement that could be very interesting biologically and allow for much deeper analysis would be to segment the background HUVEC monolayer. A segmentation protocol could be applied to do this or it could even be done manually, as there are usually less than 100 cells to segment in an image. By doing this the role of individual endothelial cells in the leukocyte adhesion cased could be observed and quantified.

Final remarks

As modern cell biology becomes increasingly a quantitative science it is hoped that the role of bioinformatics gains more recognition as being integral to the process of discovery. Moreover the skillset within the scientific community will need to adapt to better understand and assess the quantification and data analytics methods and their impact on experimental findings. This relatively recent, niche, and interdisciplinary field of bioinformatics will benefit greatly from continued advances in the fields of artificial intelligence and machine learning. These techniques should provide exciting new avenues to further improve image segmentation and recognition tasks.

The methods presented in this thesis were developed in close collaboration with experimentalists at the MRC Laboratory for Molecular Cell Biology in London. Code was concisely written and documented so as to be extensible and easily understandable. It is hoped that these methods will continue to be of use within the laboratory, in this way they can continue to provide unique insight and understanding of endothelial processes. The code base provides a foundation which could be used for the design and architecture of future computational methods. The ideas and explanations presented in this thesis will serve as a useful reference to guide the use and potential development of these tools.

References

[1] B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter, *Molecular Biology of the Cell*. New York: Garland Science, 4 ed., 2002.

[2] S. G. Eskin, C. L. Ives, L. V. McIntire, and L. T. Navarro, "Response of cultured endothelial cells to steady flow," *Microvascular Research*, vol. 28, no. 1, pp. 87–94, 1983.

[3] B. E. Sumpio, J. Timothy Riley, and A. Dardik, "Cells in focus: endothelial cell," *International Journal of Biochemistry and Cell Biology*, vol. 34, no. 12, pp. 1508–1512, 2002.

[4] E. R. Weibel and G. E. Palade, "New Cytoplasmic Components in Arterial Endothelia.," *The Journal of Cell Biology*, vol. 23, pp. 101–112, 1964.

[5] F. Ferraro, J. Kriston-Vizi, D. J. Metcalf, B. Martin-Martin, J. Freeman, J. J. Burden, D. Westmoreland, C. E. Dyer, A. E. Knight, R. Ketteler, and D. F. Cutler, "A Two-Tier Golgi-Based Control of Organelle Size Underpins the Functional Plasticity of Endothelial Cells," *Developmental Cell*, vol. 29, no. 3, pp. 292–304, 2014.

[6] D. D. Wagner, J. B. Olmsted, and V. J. Marder, "Immunolocalization of von Willebrand Protein in WeibeI-Palade Bodies of Human Endothelial Cells," *The Journal of Cell Biology*, vol. 95, no. 1, pp. 355–360, 1982.

[7] R. Bonfanti, B. C. Furie, B. Furie, and D. D. Wagner, "PADGEM (GMP140) is a component of Weibel-Palade bodies of human endothelial cells.," *Blood*, vol. 73, no. 5, pp. 1109–1112, 1989.

[8] R. P. McEver, J. H. Beckstead, K. L. Moore, L. Marshall-Carlson, and D. F. Bainton, "GMP-140, a platelet alpha-granule membrane protein, is also synthesized by vascular endothelial cells and is localized in Weibel-Palade bodies," *Journal of Clinical Investigation*, vol. 84, no. 1, pp. 92–99, 1989.

[9] E. R. Weibel, "Fifty years of Weibel-Palade bodies: The discovery and early history of an enigmatic organelle of endothelial cells," *Journal of Thrombosis and Haemostasis*, vol. 10, no. 6, pp. 979–984, 2012.

[10] H. Rasche, "Haemostasis and thrombosis: an overview," *European Heart Journal Supplements*, vol. 3, pp. Q3–Q7, dec 2001.

[11] Z. M. Ruggeri, "The role of von Willebrand factor in thrombus formation," *Thrombosis Research*, vol. 120, pp. S5–S9, jan 2007.

[12] H. O. Ogedegbe, "An Overview of Hemostasis," *Toxicologic Pathology*, vol. 21, no. 2, pp. 170–179, 1993.

[13] J. E. Hall and A. C. Guyton, *Guyton and Hall Textbook of Medical Physiology*, vol. 53. 12 ed., 2011.

[14] G. Pocock, C. A. Richards, and D. A. Richards, *Human Physiology*. Oxford University Press, 4 ed., 2013.

[15] K. De Ceunynck, S. F. De Meyer, and K. Vanhoorelbeke, "Unwinding the von Willebrand factor strings puzzle," *Blood*, vol. 121, no. 2, pp. 270–277, 2013.

[16] J. E. Sadler, "Biochemistry and Genetics of von Willebrand Factor," *Annual Review of Biochemistry*, vol. 67, pp. 395–424, 1998.

[17] M. Serpe, D. Umulis, A. Ralston, J. Chen, D. J. Olson, H. Othmer, M. B. O. Connor, and S. S. Blair, "Re-establishment of VWF-dependent Weibel-Palade bodies in VWD endothelial cells," *Blood*, vol. 14, no. 6, pp. 612–626, 2008.

[18] P. J. Ewenstein BM, Warhol MJ, Handin RI, "Composition of the von Willebrand factor storage organelle (Weibel- Palade body) isolated from cultured human umbilical vein endothelial cells," *The Journal of Cell Biology*, vol. 104, no. 5, pp. 1423–1433, 1987.

[19] T. N. Mayadas, R. C. Johnson, H. Rayburn, R. O. Hynes, and D. D. Wagner, "Leukocyte Rolling and Extravasation are Severely Compromised in P Selectin-Deficient Mice," *Cell*, vol. 74, no. 3, pp. 541–554, 1993.

[20] K. Ley, C. Laudanna, M. I. Cybulsky, and S. Nourshargh, "Getting to the site of inflammation: the leukocyte adhesion cascade updated.," *Nature reviews. Immunology*, vol. 7, no. 9, pp. 678–689, 2007.

[21] M. L. Watson, S. F. Kingsmore, G. I. Johnston, M. H. Siegelman, M. M. Le Beau, R. S. Lemons, N. S. Bora, T. A. Howard, I. L. Weissman, and R. P. McEver, "Genomic organization of the selectin family of leukocyte adhesion molecules on human and mouse chromosome 1.," *The Journal of experimental medicine*, vol. 172, no. 1, pp. 263–272, 1990.

[22] R. P. McEver, "Selectins: Lectins that initiate cell adhesion under flow," *Current Opinion in Cell Biology*, vol. 14, no. 5, pp. 581–586, 2002.

[23] E. Larsen, A. Celi, G. E. Gilbert, B. C. Furie, J. K. Erban, R. Bonfanti, D. D. Wagner, and B. Furie, "PADGEM protein: a receptor that mediates the interaction of activated platelets with neutrophils and monocytes," *Cell*, vol. 59, no. 2, pp. 305–312, 1989.

[24] T. A. Springer, "Adhesion receptors of the immune system," *Nature*, vol. 346, no. 6283, pp. 425–434, 1990.

[25] D. C. Anderson and T. A. Springer, "Leukocyte adhesion deficiency: an inherited defect in the Mac-1, LFA-1, and p150,95 glycoproteins," *Annual review of Medicine*, vol. 38, no. 1, pp. 175–194, 1987.

[26] A. Atherton and G. V. R. Born, "Quantitative investigations of the adhesiveness of circulating polymorphonuclear leucocytes to blood vessel walls," *The Journal of Physiology*, vol. 222, no. 2, pp. 447–474, 1972.

[27] D. L. Spector and R. D. Goldman, *Basic Methods in Microscopy : a laboratory manual.* New York: Cold Spring Harbor Laboratory Press, 2006.

[28] K. Thorn, "A quick guide to light microscopy in cell biology," *Molecular Biology of the Cell*, vol. 27, no. 2, pp. 219–222, 2016.

[29] D. B. Murphy and M. W. Davidson, *Fundamentals of Light Microscopy and Electronic Imaging*. New York: John Wiley & Sons, 2 ed., 2001.

[30] V. Ljosa and A. E. Carpenter, "Introduction to the quantitative analysis of two-dimensional fluorescence microscopy images for cell-based screening," *PLoS Computational Biology*, vol. 5, no. 12, pp. 1–10, 2009.

[31] E. Abbe, "Beitrage zur Theorie des Mikroskops und der mikroskopischen Wahrnehmung.," *Archiv für mikroskopische Anatomie*, vol. 9, no. 1, pp. 413–418, 1873.

[32] M. Minsky, "Memoir on inventing the confocal scanning microscope," *Scanning*, vol. 10, no. 4, pp. 128–138, 1988.

[33] C. J. R. Sheppard and A. Choudhury, "Image Formation in the Scanning Microscope," *Optica Acta: International Journal of Optics*, vol. 24, no. 10, pp. 1051–1073, 1977.

[34] C. G. Galbraith and J. A. Galbraith, "Super-resolution microscopy at a glance.," *Journal of cell science*, vol. 124, no. Pt 10, pp. 1607–1611, 2011.

[35] L. Schermelleh, R. Heintzmann, and H. Leonhardt, "A guide to super-resolution fluorescence microscopy," *Journal of Cell Biology*, vol. 190, no. 2, pp. 165–175, 2010.

[36] M. G. Gustafsson, "Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy.," *Journal of Microscopy*, vol. 198, no. 2, pp. 82–87, 2000.

[37] D. Westmoreland, M. Shaw, W. Grimes, D. J. Metcalf, J. J. Burden, K. Gomez, A. E. Knight, and D. F. Cutler, "Super-resolution microscopy as a potential approach to platelet granule disorder diagnosis," *Journal of Thrombosis and Haemostasis*, vol. 14, pp. 839–849, 2016.

[38] B. Huang, M. Bates, and X. Zhuang, "Super resolution fluorescence microscopy," *Annual Review of Biochemistry*, vol. 78, pp. 993–1016, 2010.

[39] T. Otaki, "Artifact Halo Reduction in Phase Contrast Microscopy Using Apodization," *Optical Review*, vol. 7, no. 2, pp. 119–122, 2000.

[40] A. R. Parker, "On the origin of optics," *Optics & Laser Technology*, vol. 43, pp. 323–329, mar 2011.

[41] T. Jones, A. Carpenter, D. Sabatini, and P. Golland, "Methods for high-content, high-throughput image-based cell screening," *Proceedings of the First MICCAI Workshop on Microscopic Image Analysis with Applications in Biology*, pp. 65–72, 2006.

[42] D. Prodanov and K. Verstreken, *Automated Segmentation and Morphometry of Cell and Tissue Structures. Selected Algorithms in ImageJ*. 2012.

[43] A. H. K. Roeder, A. Cunha, M. C. Burl, and E. M. Meyerowitz, "A computational image analysis glossary for biologists," *Development*, vol. 139, pp. 3071–3080, sep 2012.

[44] J. C. Russ, N. Carolina, and C. John, *The Image Processing Handbook*. CRC Press, 6 ed., 2006.

[45] J. B. Pawley, *Handbook Of Biological Confocal Microscopy*. Boston: Springer US, 2006.

[46] J. C. Waters, "Accuracy and precision in quantitative fluorescence microscopy," *Journal of Cell Biology*, vol. 185, no. 7, pp. 1135–1148, 2009.

[47] J. Pawley, "The 39 steps: A cautionary tale of quantitative 3-D fluorescence microscopy," *BioTechniques*, vol. 28, no. 5, pp. 884–888, 2000.

[48] M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis, and Machine Vision*, vol. 3. Thomson-Engineering, 2 ed., 2007.

[49] D. J. Ketcham, R. W. Lowe, and J. W. Weber, "Image Enhancement Techniques for Cockpit Displays," *National Technical Information Service*, no. AD-A014 928, 1974.

[50] C. Pal, A. Chakrabarti, and R. Ghosh, "A Brief Survey of Recent Edge-Preserving Smoothing Algorithms on Digital Images," *Procedia Computer Science*, vol. 00, pp. 1–40, mar 2015.

[51] N. Otsu, "A Threshold Selection Method from Gray-Level Histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.

[52] J. Kittler and J. Illingworth, "Minimum error thresholding," *Pattern Recognition*, vol. 19, no. 1, pp. 41–47, 1986.

[53] W. H. Tsai, "Moment-preserving thresolding: A new approach," *Computer Vision, Graphics, and Image Processing*, vol. 29, no. 3, pp. 377–393, 1985.

[54] J. Roerdink and A. Meijster, "The Watershed Transform: Definitions, Algorithms and Parallelization Strategies," *Fundamenta Informaticae*, vol. 41, no. 1-2, pp. 187–228, 2000.

[55] L. Vincent and P. Soille, "Watersheds in digital spaces: An efficient algorithm based on immersion simulations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 6, pp. 583–598, 1991.

[56] J. Goutsias, L. Vincent, and D. S. Bloomberg, "Mathematical Morphology and Its Applications to Image and Signal Processing," *Image Rochester NY*, vol. 18, no. 8, p. 456, 2000.

[57] M. Benes and B. Zitova, "Performance evaluation of image segmentation algorithms on microscopic image data," *Journal of Microscopy*, vol. 257, no. 1, pp. 65–85, 2015.

[58] H. Zhang, J. E. Fritts, and S. A. Goldman, "Image segmentation evaluation: A survey of unsupervised methods," *Computer Vision and Image Understanding*, vol. 110, no. 2, pp. 260–280, 2008.

[59] J. R. Cardoso, L. M. Pereira, M. D. Iversen, and A. L. Ramos, "What is gold standard and what is ground truth?," *Dental Press Journal of Orthodontics*, vol. 19, no. 5, pp. 27–30, 2014.

[60] R. Unnikrishnan, C. Pantofaru, and M. Hebert, "A measure for objective evaluation of image segmentation algorithms," in *CVPR '05 Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Workshops*, vol. 3, pp. 34–42, 2005.

[61] R. Parikh, A. Mathai, S. Parikh, G. C. Sekhar, and R. Thomas, "Understanding and using sensitivity, specificity and predictive values," *Indian Journal of Opthalmology*, vol. 56, no. 1, pp. 45–50, 2008.

[62] D. Vestweber, "VE-cadherin: The major endothelial adhesion molecule controlling cellular junctions and blood vessel formation," *Arteriosclerosis, Thrombosis, and Vascular Biology*, vol. 28, no. 2, pp. 223–232, 2008.

[63] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu, "scikit-image: image processing in Python," *PeerJ*, vol. 2, p. e453, 2014.

[64] F. Pedregosa, O. Grisel, R. Weiss, A. Passos, and M. Brucher, "Scikit-learn : Machine Learning in Python," vol. 12, pp. 2825–2830, 2011.

[65] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

[66] C.-C. Chang and C.-J. Lin, "LIBSVM : A Library for Support Vector Machines," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 2, pp. 1–39, 2011.

[67] C. W. Hsu, C. Chin-Chung, and C. J. Lin, "A Practical Guide to Support Vector Classification," *BJU international*, vol. 101, no. 1, pp. 1396–400, 2008.

[68] P. Bajcsy, A. C. J. Chalfoun, M. Halter, D. Juba, M. Kociolek, M. Majurski, A. Peskin, C. Simon, M. Simon, A. Vandecreme, and M. Brady, "Survey Statistics of Automated Segmentations Applied to Optical Imaging of Mammalian Cells," *BMC Bioinformatics*, no. Under review, pp. 1–28, 2015.

[69] W. Yu, H. K. Lee, S. Hariharan, W. Bu, and S. Ahmed, "Evolving generalized voronoi diagrams for accurate cellular image segmentation," *Cytometry Part A*, vol. 77, pp. 379–386, 2010.

[70] A. Chambolle, "An Algorithm for Total Variation Minimization and Applications," *Journal of Mathematical Imaging and Vision*, vol. 20, no. 1-2, pp. 89–97, 2004.

[71] N. L. Stevenson, B. Martin-Martin, J. Freeman, J. Kriston-Vizi, R. Ketteler, and D. F. Cutler, "G protein-coupled receptor kinase 2 moderates recruitment of THP-1 cells to the endothelium by limiting histamine-invoked Weibel-Palade body exocytosis.," *Journal of Thrombosis and Haemostasis*, vol. 12, no. 2, pp. 261–272, 2014.

[72] S. R. Sternberg, "Biomedical Image Processing," *Computer*, vol. 16, no. 1, pp. 22–34, 1983.

[73] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, and T. scikit-image contributors, "scikit-image: image processing in Python," *PeerJ*, vol. 2, pp. 1–18, 2014.

[74] ImageJ, "ImageJ User Guide," *IJ 1.46r*, p. 187, 2003.

[75] W. McKinney, "pandas: a Foundational Python Library for Data Analysis and Statistics," *Python for High Performance and Scientific Computing*, pp. 1–9, 2011.

[76] P. J. Clark and F. C. Evans, "Distance to Nearest Neighbor as a Measure of Spatial Relationships in Populations," *Ecology*, vol. 35, no. 4, pp. 445–453, 1954.

[77] L. R. . Dice, "Measures of the Amount of Ecologic Association Between Species," *Ecology*, vol. 26, no. 3, pp. 297–302, 1945.

[78] P. Jaccard, "The distribution of the flora in the alphine zone," *The New Phytologist*, vol. XI, no. 2, pp. 37–50, 1912.

[79] K. M. Valentijn, L. F. van Driel, M. J. Mourik, G.-J. Hendriks, T. J. Arends, A. J. Koster, and J. a. Valentijn, "Multigranular exocytosis of Weibel Palade bodies in vascular endothelial cells.," *Blood*, vol. 116, no. 10, pp. 1807–1816, 2010.

[80] R Core Team, "R : A Language and Environment for Statistical Computing," vol. 1, 2014.

[81] R. O. Duda, P. E. Hart, and D. G. Stork, "Pattern Classification," 2000.

[82] E. Dimitriadou, K. Hornik, F. Leisch, D. Meyer, and A. Weingessel, "e1071: Misc Functions of the Department of Statistics (e1071), TU Wien," 2010.

[83] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and Regression Trees*, vol. 19. 1984.

[84] L. Rokach and O. Maimon, "Top-Down Induction of Decision Trees Classifiers—A Survey," *IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews)*, vol. 35, no. 4, pp. 476–487, 2005.

[85] T. Therneau, B. Atkinson, B. Ripley, and M. B. Ripley, "rpart: Recursive Partitioning and Regression Trees.," *R Package version 4.1-10*, vol. https://CR, 2015.

[86] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.

[87] H. Van Nispen Tot Pannerden, F. De Haas, W. Geerts, G. Posthuma, S. Van Dijk, and H. F. G. Heijnen, "The platelet interior revisited: electron tomography reveals tubular alpha-granule subtypes," *Blood*, vol. 116, no. 7, pp. 1147–1156, 2010.

[88] S. L. Seward and W. A. Gahl, "Hermansky-Pudlak Syndrome: Health Care Throughout Life," *Pediatrics*, vol. 132, no. 1, pp. 153–160, 2013.

[89] F. Rendu and B. Brohard-Bohn, "The platelet release reaction: granules' constituents, secretion and functions.," *Platelets*, vol. 12, no. 5, pp. 261–273, 2001.

[90] F. A. Ruiz, C. R. Lea, E. Oldfield, and R. Docampo, "Human platelet dense granules contain polyphosphate and are similar to acidocalcisomes of bacteria and unicellular eukaryotes," *The Journal of Biological Chemistry*, vol. 279, no. 43, pp. 44250–44257, 2004.

[91] G. L. Reed, M. L. Fitzgerald, and J. Polgár, "Molecular mechanisms of platelet exocytosis: insights into the "secrete" life of thrombocytes.," *Blood*, vol. 96, no. 10, pp. 3334–3342, 2000.

[92] A. T. Nurden and P. Nurden, "Congenital platelet disorders and understanding of platelet function," *British Journal of Haematology*, vol. 165, no. 2, pp. 165–178, 2014.

[93] M. L. Wei, "Hermansky-Pudlak syndrome: A disease of protein trafficking and organelle function," *Pigment Cell Research*, vol. 19, no. 1, pp. 19–42, 2006.

[94] R. E. Rumbaut and P. Thiagarajan, *Platelet-Vessel Wall Interactions in Hemostasis and Thrombosis*, vol. 2. 2010.

[95] M. S. Pols and J. Klumperman, "Trafficking and function of the tetraspanin CD63," *Experimental Cell Research*, vol. 315, no. 9, pp. 1584–1592, 2009.

[96] M. G. L. Gustafsson, "Nonlinear structured-illumination microscopy: wide-field fluorescence imaging with theoretically unlimited resolution," in *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, pp. 13081–13086, 2005.

[97] M. G. L. Gustafsson, L. Shao, P. M. Carlton, C. J. R. Wang, I. N. Golubovskaya, W. Z. Cande, D. A. Agard, and J. W. Sedat, "Three-dimensional resolution doubling in wide-field fluorescence microscopy by structured illumination.," *Biophysical journal*, vol. 94, no. 12, pp. 4957–4970, 2008.

[98] K. O. Holleran and M. Shaw, "Optimized approaches for optical sectioning and resolution enhancement in 2D structured illumination microscopy," *Biomedical Optics Express*, vol. 5, no. 8, pp. 113–121, 2014.

[99] C. A. Schneider, W. S. Rasband, and K. W. Eliceiri, "NIH Image to ImageJ: 25 years of image analysis," *Nature Methods*, vol. 9, no. 7, pp. 671–675, 2012.

[100] M. J. Kraus, H. Neeb, and E. F. Strasser, "Fractal and Euclidean descriptors of platelet shape.," *Platelets*, vol. 25, no. 7, pp. 1–11, 2013.

[101] E. C. Dell'Angelica, V. Shotelersuk, R. C. Aguilar, W. A. Gahl, and J. S. Bonifacino, "Altered trafficking of lysosomal proteins in Hermansky-Pudlak syndrome due to mutations in the beta 3A subunit of the AP-3 adaptor," *Molecular Cell*, vol. 3, no. 1, pp. 11–21, 1999.

[102] M. E. Ambühl, C. Brepsant, J. J. Meister, A. B. Verkhovsky, and I. F. Sbalzarini, "High-resolution cell outline segmentation and tracking from phase-contrast microscopy images," *Journal of Microscopy*, vol. 245, no. 2, pp. 161–170, 2012.

[103] V. Jossen, R. Pörtner, S. C. Kaiser, M. Kraume, D. Eibl, and R. Eibl, "Mass Production of Mesenchymal Stem Cells — Impact of Bioreactor Design and Flow Conditions on Proliferation and Differentiation," *Cells and Biomaterials in Regenerative Medicine*, pp. 119–174, 2014.

[104] M. Sperandio, J. Pickard, S. Unnikrishnan, S. T. Acton, and K. Ley, "Analysis of Leukocyte Rolling In Vivo and In Vitro," *Methods in Enzymology*, vol. 416, no. 1990, pp. 346–371, 2006.

[105] N. Ray, S. T. Acton, and K. Ley, "Tracking leukocytes in vivo with shape and size constrained active contours," *IEEE Transactions on Medical Imaging*, vol. 21, no. 10, pp. 1222–1235, 2002.

[106] S. T. Acton, K. Wethmar, and K. Ley, "Automatic tracking of rolling leukocytes in vivo.," *Microvascular research*, vol. 63, pp. 139–148, 2002.

[107] O. Debeir, P. Van Ham, R. Kiss, and C. Decaestecker, "Tracking of Migrating Cells Under Phase-Contrast Video Microscopy With Combined Mean-Shift Processes," *IEEE Transaction on medical imaging*, vol. 24, no. 6, pp. 697–711, 2005.

[108] R. G. Bacabac, T. H. Smit, S. C. Cowin, J. J. W. A. Van Loon, F. T. M. Nieuwstadt, R. Heethaar, and J. Klein-Nulend, "Dynamic shear stress in parallel-plate flow chambers," *Journal of Biomechanics*, vol. 38, no. 1, pp. 159–167, 2005.

[109] B. R. Munson, T. H. Okiishi, W. W. Huebsch, and A. P. Rothmayer, *Fundamentals of fluid mechanics*. 7 ed., 2013.

[110] G. W. Schmid-Schonbein, Y. Y. Shin, and S. Chien, "Morphometry of Human Leukoocytes," *Blood*, vol. 56, no. 5, pp. 1–4, 1998.

[111] K. Ley and P. Gaehtgens, "Endothelial, not hemodynamic, differences are responsible for preferential leukocyte rolling in rat mesenteric venules.," *Circulation research*, vol. 69, no. 4, pp. 1034–41, 1991.

[112] A. J. Hand, T. Sun, D. C. Barber, D. R. Hose, and S. MacNeil, "Automated tracking of migrating cells in phase-contrast video microscopy sequences using image registration," *Journal of Microscopy*, vol. 234, no. 1, pp. 62–79, 2009.

[113] Z. Yin, M. Kanade, and M. Chen, "Understanding the Phase Contrast Optics to Restore Artifact- free Mic roscopy Images for Segmentation," *Medical Image Analysis*, vol. 16, no. 5, pp. 997–1003, 2012.

[114] R. Lienhart and J. Maydt, "An extended set of Haar-like features for rapid object detection," in *Proceedings. International Conference on Image Processing*, vol. 1, pp. 1–3, 2002.

[115] N. M. Sasi and V. K. Jayasree, "Contrast Limited Adaptive Histogram Equalization for Qualitative Enhancement of Myocardial Perfusion Images," *Engineering*, vol. 5, no. October, pp. 326–331, 2013.

[116] R. Hummel, "Image enhancement by histogram transformation," *Computer Graphics and Image Processing*, vol. 6, no. 2, pp. 184–195, 1977.

[117] S. C. Cheung and C. Kamath, "Robust techniques for background subtraction in urban traffic video," in *Proceedings of Video Communications and Image Processing, SPIE Electronic Imaging*, pp. 881–892, 2004.

[118] M. Wells, C. M., Parsons, *Cell Migration Developmental Methods and Protocols*. Springer, 2nd ed., 2005.

[119] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," in *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, vol. 1, 2001.

[120] C. P. Papageorgiou, M. Oren, and T. Poggio, "A general framework for object detection," in *Computer Vision, 1998. Sixth International Conference on*, pp. 555–562, 1998.

[121] J. Friedman, T. Hastie, and R. Tibshirani, "Additive Logistic Regression: A Statistical View of Boosting," *The Annals of Statistics*, vol. 28, no. 2, pp. 337–407, 2000.

[122] R. Lienhart, A. Kuranov, and V. Pisarevsky, "Empirical analysis of detection cascades of boosted classifiers for rapid object detection," *Proceedings of the 25th DAGM Pattern Recognition Symposium*, pp. 297–304, 2003.

[123] Y.-L. Qiao and T. Guo, "Haar Wavelet Transform Based Fast K-Nearest Neighbor Search Algorithm," *2010 First International Conference on Pervasive Computing, Signal Processing and Applications*, pp. 895–898, 2010.

[124] D. Comaniciu, V. Ramesh, and P. Meer, "Real-time tracking of non-rigid objects using mean shift," *IEEE Conference on Computer Vision and Pattern Recognition*, vol. 2, no. 7, pp. 142–149, 2000.

[125] K. Bernardin and R. Stiefelhagen, "Evaluating multiple object tracking performance: The CLEAR MOT metrics," *Eurasip Journal on Image and Video Processing*, vol. 2008, 2008.

[126] F. Ferraro, S. Mafalda Lopes da, W. Grimes, H. K. Lee, R. Ketteler, J. Kriston-Vizi, and D. F. Cutler, "Weibel-Palade body size modulates the adhesive activity of its von Willebrand Factor cargo in cultured endothelial cells," *Scientific Reports*, vol. 6, 2016.