TAMPEREEN TEKNILLINEN YLIOPISTO
TAMPERE UNIVERSITY OF TECHNOLOGY

KIMMO KARTASALO

SPECTRAL ANALYSIS OF ORIENTED FEATURES IN MICROGRAPHS

Master of Science Thesis

Examiners:
University lecturer Heikki Huttunen
Professor Pasi Kallio

Examiner and topic approved by the Faculty Council of Computing and Electrical Engineering
on 4th March 2015

ABSTRACT

KIMMO KARTASALO: Spectral analysis of oriented features in micrographs
Tampere University of Technology
Master of Science Thesis, 90 pages, 2 Appendix pages
August 2015
Master's Degree Programme in Electrical Engineering
Major: Signal processing
Examiners: University lecturer Heikki Huttunen and Professor Pasi Kallio

Keywords: spectral analysis, Fourier transform, power spectrum, image analysis, microscopy, orientation, isotropy

Image analysis of micrographs is highly important in modern biomedical research. Oriented features, such as the myofibrils of muscle cells and other fibrillar structures of the cytoskeleton, the mechanical backbone of cells, are examples of targets which can be analyzed based on micrographs. Image based analysis of such structures is often limited to qualitative evaluation by human experts, hampering the throughput, repeatability and reliability of the analyses. Software tools are not readily available for this purpose and the existing methods typically rely at least partly on manual operation.

In this study, CytoSpectre, an automated software tool based on spectral analysis was developed. The software allows analysis of oriented features via quantification of orientation and also size distributions of the structures of interest. In principle, the software is compatible with most 2D images. CytoSpectre utilizes classical spectral estimation methods to estimate the power spectrum of a micrograph and based on the spectrum, computes parameter values describing, among others, the mean orientation, isotropy and size of target structures. By constraining the spectral analysis to a limited range of spatial frequencies, the analysis can be further tuned to focus on targets of particular size at cellular or subcellular scales. The software was implemented in MATLAB and it can be operated via a graphical user interface on basic hardware without any programming expertise.

Using a combination of freely available software and algorithms developed as part of this study, thousands of artificial micrographs with realistic characteristics and varying image quality were generated to quantitatively analyze the performance of CytoSpectre. The software was also benchmarked against FibrilTool, a recently published orientation analysis tool. CytoSpectre was found to be tolerant against noise and blurring and superior to FibrilTool when analyzing realistic targets with degraded image quality. Performance of the method was evaluated by comparisons with manual measurements performed for real images by a panel of human experts. The analysis of real images indicated general good agreement between computational and manual results while also revealing notable expert-to-expert variation. Moreover, the experiment showed that CytoSpectre can handle images obtained of different cell types using different microscopy techniques. Based on these results, the tool is expected to be useful in diverse applications dealing with biological structures whose orientation and size distributions are of interest. While designed for use within the biomedical field, the software could also be useful in non-biological applications such as materials research.

TIIVISTELMÄ

KIMMO KARTASALO: Suuntautuneiden rakenteiden spektrianalyysi mikroskooppikuvista
Tampereen teknillinen yliopisto
Diplomityö, 90 sivua, 2 liitesivua
Elokuu 2015
Sähkötekniikan diplomi-insinöörin tutkinto-ohjelma
Pääaine: Signaalinkäsittely
Tarkastajat: yliopistonlehtori Heikki Huttunen ja professori Pasi Kallio

Avainsanat: spektrianalyysi, Fourier-muunnos, tehospektri, kuva-analyysi, mikroskopia, suuntautuneisuus, isotropia

Mikroskooppikuvien analyysi on keskeinen työkalu modernissa biolääketieteellisessä tutkimuksessa. Suuntautuneet rakenteet, kuten lihassäikeet ja muut solujen mekaanisen tukirangan eli sytoskeletonin säiemäiset rakenteet ovat hyviä esimerkkejä kohteista, joita voidaan analysoida kuvaan perustuvilla mittauksilla. Tällaisten rakenteiden kuvapohjainen analyysi rajoittuu kuitenkin usein asiantuntijoiden tekemiin subjektiivisiin arvioihin, mikä heikentää paitsi tulosten toistettavuutta ja luotettavuutta, myös rajoittaa analysoitavan datan määrää. Tällaisiin tutkimuksiin tarkoitettuja vapaasti saatavilla olevia ohjelmistoja ei juuri ole ja harvat olemassa olevat työkalut eivät ole automatisoituja.

Tässä työssä kehitettiin automatisoitu spektrianalyysipohjainen työkalu, CytoSpectre, jonka avulla on mahdollista tutkia suuntautuneiden rakenteiden suunta- ja kokojakaumia. Työkalu on periaatteessa yhteensopiva useimpien 2D-kuvien kanssa. CytoSpectre hyödyntää klassista spektriestimointia analysoitavan kuvan tehospektrin muodostamiseksi. Tehospektrin perusteella ohjelma analysoi muun muassa kohteiden keskimääräistä suuntaa, isotropiaa ja kokoa. Analyysi voidaan lisäksi kohdistaa halutun kokoluokan kohteisiin rajaamalla spektrianalyysi tietylle taajuuskaistalle, mikä mahdollistaa esimerkiksi solunsisäisten pienempien kohteiden erillisen tarkastelun. Ohjelma toteutettiin MATLAB:illa ja sitä voidaan käyttää graafisen käyttöliittymän kautta ilman erityisiä vaatimuksia laitteiston suorituskyvylle tai käyttäjän erikoisosaamiselle.

CytoSpectren suorituskykyä mitattiin kvantitatiivisesti kattavilla simulaatioilla, jotka perustuivat sekä olemassa olevilla että tässä työssä kehitetyillä menetelmillä luotuihin laadultaan vaihteleviin keinotekoisiin mikroskooppikuviin. Tulosten oikeellisuutta verrattiin osittain vastaavan työkalun, FibrilToolin, tuottamiin tuloksiin. CytoSpectren havaittiin sietävän hyvin kohinaisia ja sumeita kuvia, sekä saavuttavan FibrilToolia parempia tuloksia erityisesti kuvanlaadultaan huonojen kuvien tapauksessa. Työkalun soveltuvuus aidoille kuville varmennettiin vertaamalla tuloksia asiantuntijapaneelin vaihekontrasti- ja fluoresenssimikroskopiakuvista tekemiin subjektiivisiin määrityksiin. Tulokset olivat pääasiassa yhteneviä, mutta yksittäisten asiantuntijoiden välillä havaittiin myös suurta vaihtelua. Koe osoitti myös, että CytoSpectre soveltuu käytettäväksi erilaisista solutyypeistä eri mikroskopiatekniikoilla otettujen kuvien analysoimiseksi. Yhteenvetona näiden tulosten perusteella tässä työssä kehitetyn työkalun voidaan odottaa olevan hyödyksi sovelluksissa, joissa kiinnostuksen kohteena ovat biologisten rakenteiden suunta- ja kokojakaumat. Ohjelma voi myös soveltua biolääketieteellisten tutkimusten lisäksi ei-biologisiin sovelluksiin esimerkiksi materiaalitutkimuksen alalla.

PREFACE

This study was performed as part of the Human Spare Parts Project at the Institute of Biosciences and Medical Technology (BioMediTech), a joint institute of University of Tampere and Tampere University of Technology. The research groups involved were the Micro- and Nanosystems Research Group of Tampere University of Technology and the Heart Group of University of Tampere. The contents of this thesis parallel those of a scientific journal article, prepared during the study, in which I am the first author. At the time of writing, the article was being evaluated for publication.

First, I wish to thank Professors Pasi Kallio, Katriina Aalto-Setälä and Jukka Lekkala for the confidence they have placed on me and the opportunity of working in a very interdisciplinary and open-minded environment. I also wish to thank Heikki Huttunen for his feedback in the role of the examiner of this thesis. I am most grateful to Heart Group members Risto-Pekka Pölönen and Marisa Ojala for the crucial contributions they have made to this project in the form of providing test images, user feedback and expert opinions on cell biology and tissue engineering. I would also like to thank Heimo Ihalainen from the Department of Automation Science and Engineering, Tampere University of Technology and Jyrki Rasku, from the School of Information Sciences, University of Tampere for the fruitful discussions we had especially during the early stages of this study.

I am thankful to Heart Group members Markus Haponen, Mostafa Kiamehr, Janne Koivisto, Eeva Laurila, Suvi Marttila, Mari Pekkanen-Mattila, Kirsi Penttinen, Chandra Krishna Prajapati, Disheet Shah, Henna Venäläinen and Leena Viiri for their time spent on manual expert evaluation of images. I also appreciate all the comments and suggestions made by fellow Micro- and Nanosystems Research Group members, as well as members of the Sensor Technology and Biomeasurements Group. I would like to congratulate everyone in the Human Spare Parts Project for succeeding in building up a community of scientists, which welcomes new ideas while avoiding many of the artificial barriers between different disciplines, research groups, departments or universities. Hopefully ‘the Human Spare Parts spirit’ will continue to spread! Finally, I wish to thank my family and friends, especially my beloved Miina, for their love and support.

Tampere, 24.8.2015

Kimmo Kartasalo

CONTENTS

LIST OF SYMBOLS AND ABBREVIATIONS

BSA	bovine serum albumin
CS	CytoSpectre
DAPI	4’,6-diamidino-2-phenylindole
DFT	Discrete Fourier Transform
DIC	Differential Interference Contrast
END-2	visceral endoderm-like cells
FFT	Fast Fourier Transform
FT	FibrilTool
FWHM	full width at half maximum
GUI	graphical user interface
hiPSC	human induced pluripotent stem cell
hiPSC-CM	human induced pluripotent stem cell derived cardiomyocyte
hiPSC-PSN	human induced pluripotent stem cell derived peripheral sensory neuron
IR	infrared
MCM	Modulation Contrast Microscopy
NDS	normal donkey serum
PBS	phosphate buffered saline
PCR	polymerase chain reaction
PFA	paraformaldehyde
PSD	power spectral density
qPCR	quantitative polymerase chain reaction
ROI	region of interest
SNR	signal-to-noise ratio
UV	ultraviolet
WOSA	Weighted Overlapped Segment Averaging

α_i	sample of angular data
$\bar{\alpha}$	mean direction
Γ	Gamma function
Δ_C	camera pixel size
Δ_I	image pixel size
ε	angular tolerance parameter
θ	orientation
$\bar{\theta}$	initial mean orientation of the detail component
$\Lambda(\lambda)$	wavelength distribution
λ	wavelength
σ_γ	scale parameter of the Gamma distribution
σ_f	standard deviation of the Gaussian detail component model
σ_θ	initial angular standard deviation of the detail component
$\sigma_{K,f}$	standard deviation of the Gaussian mask along the spatial frequency axis
$\sigma_{K,\vartheta}$	standard deviation of the Gaussian mask along the orientation axis
τ	convergence threshold parameter
$\Phi(u,v)$	phase spectrum
$\Psi(\theta)$	orientation distribution

A	significance level parameter
b	circular skewness
$B(f,\theta)$	background power spectrum in polar coordinates
c	correction factor for binned angular data
d	bin spacing for binned angular data
E	prior size parameter
f	spatial frequency
$\bar{f}$	mean of the Gaussian detail component model
f_H	higher spatial frequency cutoff
f_L	lower spatial frequency cutoff
$f_{Nyquist}$	Nyquist frequency
$F(u)$	1D discrete Fourier transform of $g(x)$
$F(u,v)$	2D discrete Fourier transform of $g(x,y)$
$g(x)$	1D input signal
$g(x,y)$	2D input signal
h_γ	shape parameter of the Gamma distribution
$H(u,v)$	frequency response of a low-pass filter
I_0	modified Bessel function of order zero
$K(f,\theta)$	Gaussian mask
k	circular kurtosis
L	segment length parameter
M	image height, signal length
M_W	WOSA segment height
N	image width
N_W	WOSA segment width
N_θ	number of angle bins in the polar power spectrum
N_f	number of spatial frequency bins in the polar power spectrum
N_L	number of low spatial frequency spectral bins to exclude
N_{max}	maximum number of iterations parameter
N_S	number of segments averaged in Bartlett's or Welch's method
N_{uv}	number of spatial frequency bins in the Cartesian power spectrum
n	number of elements in a sample of angular data
O	magnification
$P(u,v)$	raw periodogram
$P_j(u,v)$	raw periodogram of segment j in Bartlett's method
$\hat{P}_j(u,v)$	raw periodogram of windowed segment j in Welch's method
$P_{Daniell}(u,v)$	Daniell's power spectral density estimate
$P_{Bartlett}(u,v)$	Bartlett's power spectral density estimate
$P_{WOSA}(u,v)$	Welch's power spectral density estimate
$P_{WOSA}(f,\theta)$	Welch's power spectral density estimate in polar coordinates
$P_{WOSA}(f)$	Welch's power spectral density estimate averaged over orientations
$P_{WL}(\lambda,\theta)$	power spectrum in wavelength form
$P_M(f,\theta)$	mixed component
$P_D(f,\theta)$	detail component
$p_{VM}(\theta;\mu,\kappa)$	probability density function of the von Mises distribution
R	mean resultant vector length
R_c	corrected mean resultant vector length
r_i	sample of unit vectors
$\bar{r}$	mean resultant vector
S	circular variance

s	angular standard deviation
u	vertical spatial frequency coordinate
v	horizontal spatial frequency coordinate
$VM(\mu,\kappa)$	von Mises distribution with mean μ and spread parameter κ
$w(x,y)$	windowing function
$w_{Hann}(x)$	Hann window
W	resolution parameter
$X(f,\theta)$	probabilistic detail component in polar coordinates
x	vertical spatial coordinate
y	horizontal spatial coordinate
$Y(f,\theta)$	normalized power spectrum in polar coordinates
$Z(f,\theta)$	probabilistic power spectrum in polar coordinates

1. INTRODUCTION

Modern biological and biomedical research relies heavily on different imaging systems (Eliceiri *et al.* 2012; Jang *et al.* 2015). Advancements in optics and other microscope hardware during the past two decades have facilitated the monitoring of different biological phenomena with ever-increasing resolution, specificity and speed. Many laboratories generate hundreds of gigabytes of image data each day and emerging technologies allowing highly automated image capture are set to accelerate this trend even further. This data explosion has led to a growing interest in bioimage informatics, that is, the development of computational methods for biomedical image processing, storing and analysis. The sheer volume of the data often necessitates adopting computational approaches instead of manual ones. Another key benefit of using software for image analysis is that it can increase the accuracy and reproducibility of results. Moreover, computational methods can even enable new types of experiments and biological discoveries that would not be otherwise possible. For these reasons, the importance of the bioimage informatics field is expected to keep growing and the use of computational techniques and software tools is becoming a necessity for the majority of biomedical researchers.

Although there are some general-purpose bioimage analysis tools available, it is often necessary to develop software tailored for a particular research problem (Eliceiri *et al.* 2012). One such research problem is the analysis of oriented features in micrographs, which is the focus of this study. Orientation and the degree of isotropy or anisotropy are important concepts in many biological systems (Marquez 2006; Sander & Barocas 2009; Boudaoud *et al.* 2014). Most notable examples are different fibrillar structures of the cytoskeleton, which forms the mechanical backbone of cells, and the extracellular matrix which surrounds cells in tissue. The properties of these structures are dictated by their microstructure and in turn influence the behavior and function of entire cells, tissues, organs or even organisms. For example, the structural properties of myofibers, which are a special class of fibrillar structures present in muscles, including heart muscle, are closely linked to the functional properties of myocytes (*i.e.*, muscle cells) (Karlon *et al.* 1998). Abnormal myofibrillar structure can prevent cells from functioning normally, which leads to a pathologic condition. For example, aberrant organization of cardiac myofibers in heart muscle cells is a hallmark of familial hypertrophic cardiomyopathy. Analyzing the orientation distribution of myofibers and other cellular components is also important in the context of engineered stem cell derived cardiac cells and tissues (Bray *et al.* 2010; Feinberg *et al.* 2012; Pasqualini *et al.* 2015; Khan *et al.* 2015). Production of such cells and tissue could benefit toxicological testing, drug development and also regenerative treatments, which seek to replace damaged tissue (Sinnecker

et al. 2014; Zanella *et al.* 2014; Khan *et al.* 2015). Other examples of cases where oriented features are important include the study of neurite development (Kim *et al.* 2011) and collagen structures in different tissues and biomaterials (Chaudhuri *et al.* 1987; Petroll *et al.* 1993; Marquez 2006; Pang *et al.* 2009; Sander & Barocas 2009; Rezakhaniha *et al.* 2012; Schriefl *et al.* 2012; Boudaoud *et al.* 2014). The reorientation of cells subjected to different external stimuli has also been studied (Karlon *et al.* 1999; Umeno & Ueno 2003). Despite being frequent in many biological systems, oriented features are still often analyzed based only on subjective assessments of micrographs performed by human experts (Boudaoud *et al.* 2014). Typically, such procedures are too arduous to allow high-throughput analysis. Moreover, subjective measurements hinder the repeatability and accuracy of such assays and the lack of quantitative data prevents or complicates the use of computer models and other more advanced analysis methods.

In order to obtain quantitative data describing oriented structures, both optical techniques and software tools have been developed (Boudaoud *et al.* 2014). The optical techniques use specialized microscopes allowing the imaging to be performed in an orientation-sensitive manner. The software approach has the benefit of wider applicability, as the analyses can be usually performed for images that have been routinely captured using common microscopy techniques instead of special instruments. The proposed computational approaches can be divided to three main categories based on three different concepts: segmentation, edge detection and spectral analysis. The segmentation-based methods attempt to segment the objects of interest from the image, which is often highly challenging due to intensity variations and the presence of other objects. The methods based on edge detection typically employ partial derivatives of pixel intensity to detect the sharp changes in intensity near the edges of oriented structures. Parameters describing the orientation of the structures are then calculated based on the directions of intensity gradients. Methods based on this concept have been used to analyze, for example, the arrangement of myofibers in cardiac tissue (Karlon *et al.* 1998), the response of endothelial cells to fluid flow (Karlon *et al.* 1999), the organization of collagen fibers in arteries (Rezakhaniha *et al.* 2012) and the orientation of fibers in plant cells (Boudaoud *et al.* 2014).

Methods based on spectral analysis are markedly different, as they handle the image in an alternative frequency domain representation rather than in the original spatial domain (Boudaoud *et al.* 2014). The frequency domain representation of the image consists of a spectrum of periodic components called the power spectrum or power spectral density (PSD) of the image. The orientation distribution of the different structures present in the image can be analyzed by studying the PSD. In addition, the PSD also contains information concerning the dimensions of the structures. Methods based on spectral analysis have been applied to, for example, the analysis of collagen and other fibers in different biomaterials and tissues (Chaudhuri *et al.* 1987; Petroll *et al.* 1993; Marquez 2006; Pang *et al.* 2009; Sander & Barocas 2009; Schriefl *et al.* 2012), myocytes subjected to mag-

netic fields (Umeno & Ueno 2003), neuronal processes (Kim *et al.* 2011), diseased corneas of the eye (Lo *et al.* 2012) and myofibers of stem cell derived cardiac myocytes (Pasqualini *et al.* 2015; Khan *et al.* 2015). In addition to these three main types of methods, there have been reports of approaches combining frequency domain preprocessing steps with subsequent analysis in the spatial domain (Bray *et al.* 2010; Feinberg *et al.* 2012).

Despite the large number of possible applications and previous studies dealing with orientation analysis, most of the published methods are currently not available to biomedical researchers in the form of software. However, two general purpose orientation analysis tools have been published as plug-ins for ImageJ (Schneider *et al.* 2012), an open source image processing and analysis tool. They are OrientationJ (Rezakhaniha *et al.* 2012) and the more recent FibrilTool (Boudaoud *et al.* 2014), which are both based on the detection of edges on the basis of intensity gradients. While allowing quantitative analysis of orientation distributions, these methods still have significant limitations. OrientationJ requires the tuning of multiple parameters that lack any physical interpretation and have been poorly documented. FibrilTool is much more user-friendly, but requires significant manual operation in the form of selecting regions of interest from the images and is therefore not suitable for automated, high-throughput analysis. Moreover, we are not aware of any publicly available tools utilizing spectral methods for this purpose. Spectral analysis has the capability of separately analyzing features present within different spatial frequency bands, which could be beneficial in the case of images with heterogeneous content. For example, contributions from larger objects such as cells could be differentiated from those originating from finer details within cells. This kind of separation is not possible in the spatial domain unless some kind of preprocessing steps are introduced. Such preprocessing steps can introduce different types of bias into the results and avoiding them is therefore an advantage (Boudaoud *et al.* 2014). As the PSD of an image also contains information on the distribution of spatial frequencies in addition to orientation distributions, a tool based on spectral analysis could simultaneously analyze the spatial frequency content of the image, a feature of spectral analysis that has been largely overlooked in previous studies in this field. Spatial frequencies can then be converted to wavelengths, which are easier to interpret as measures of object dimensions. In addition, methods based on spectral analysis are known to be relatively insensitive to noise (Boudaoud *et al.* 2014), computationally efficient and easily tunable to suit different applications (Sander & Barocas 2009). For these reasons, the spectral analysis approach was selected in this study.

To allow quantitative, highly automated spectral analysis of oriented features in micrographs, the primary goal of this study was to develop a user-friendly software tool. The aim was to provide biomedical researchers with a tool which could be operated via a graphical user interface (GUI) on basic hardware without prior experience of programming or image processing. A tool fulfilling these requirements, CytoSpectre, was de-

signed and implemented using MATLAB (The MathWorks, Inc., Natick, MA, USA). The software can be used to analyze the complete power spectrum of the image, in line with previous approaches, but it also allows focusing the analysis to a limited region of the spectrum. The former produces results dominated by the gross properties of the image, while the latter yields information representative of finer details of the image such as different intracellular fibrils or other structures. A number of parameter values describing the orientation and spatial frequency content of the images are computed based on the power spectrum. These summary statistics can then be exported to spreadsheets, plain text files or images for further study. The software is freely available in either MATLAB source code form or as a standalone Windows application. The source codes, the standalone installation package, a user's guide and example images for testing are available at www.tut.fi/cytospectre[1]. The second goal of this study was to develop a suitable simulation test bench in order to understand the functionality of the software and quantitatively analyze its performance. The simulator developed as part of this thesis is able to generate artificial images with varying characteristics resembling those of real micrographs. With the aid of this simulator, the performance of the analysis software was quantified through an extensive set of simulations based on thousands of artificial images. A subset of these images was also analyzed using FibrilTool (Boudaoud *et al.* 2014) and the performance of the two tools was compared. To avoid relying solely on simulations, the performance of CytoSpectre was evaluated by comparing computational and manual results for real fluorescence and phase contrast microscopy images. The real images of human induced pluripotent stem cell (hiPSC) derived cardiomyocytes (hiPSC-CM) and peripheral sensory neurons (hiPSC-PSN) were manually analyzed by a panel of human experts routinely working with these cell types at University of Tampere. This experiment allowed the quantification of expert-to-expert variation and the degree of correspondence between computational and manual measurements. Finally, a sensitivity analysis using artificial images was performed to understand the effects of adjustable parameter values on the results produced by the software.

In Chapter 2 of this thesis, basic concepts of microscopy, spectral analysis and orientation analysis are reviewed. The functionality and technical implementation of the software are presented in Chapter 3. In Chapter 4, these are followed with descriptions of the methodology used for creating artificial images and the experimental protocols used during cell culture, differentiation and image acquisition. Results of performance evaluation experiments using artificial and real images are presented and discussed in Chapter 5. Final conclusions on the results of this study are drawn in Chapter 6.

[1] Username: CytoSpectre_reviewer_2015, password: 2DSpectralAnalysis

2. BACKGROUND

In this chapter, we will review basic concepts and theories central to understanding the rest of the thesis. First, some of the optical microscopy techniques most commonly encountered in the field of biosciences are introduced in Section 2.1. Basic theory on spectral analysis with a particular focus on images is then presented in Section 2.2, followed by a review of previously reported orientation analysis techniques, software and applications in Section 2.3.

2.1 Optical microscopy

Microscopy is certainly one of the most widespread methods used in biosciences with applications ranging from obtaining basic information on the growth, death and dynamical behavior of cells to clinical diagnosis of different diseases (Jang *et al.* 2015). Microscopy techniques relying on electromagnetic radiation (in contrast to scanning probe techniques) can be divided to two main categories, namely optical microscopy and electron microscopy. While electron microscopy allows capturing images at a nanometer resolution, it is not well suited to the monitoring of biological samples due to their inherent low contrast and the requirement of very thin samples. Optical microscopy techniques are more widely used in biosciences and in many cases allow biological specimens to be imaged in a non-destructive way. For these reasons, this overview of microscopy techniques is limited to optical microscopy. Optical microscopy comprises a tremendous number of techniques relying on different principles to produce the contrast necessary for resolving targets in the image with each having different advantages and limitations (Jang *et al.* 2015).

Microscopes have been used since the 17th century (Evennett & Hammond 2005) but many of the breakthrough inventions which still influence the design of modern microscopes were made in the late 19th century and during the 20th century (Davidson & Abramowitz 2002). The work of Ernst Abbe and Carl Zeiss led to the production of apochromatic objectives in 1886, which were for the first time based on proper lens design relying on sound optical principles, greatly reducing spherical and chromatic aberrations. Since then, microscopy has developed rapidly with the introduction of new methods for enhancing contrast, labeling specific molecules and electronic acquisition and processing of images. In the overview presented in this chapter, the main focus is on techniques which are nowadays used routinely rather than on the newest techniques of this rapidly evolving field, which typically require high-cost special instruments. In the following sections, phase contrast microscopy, Differential Interference Contrast

(DIC) microscopy, Modulation Contrast Microscopy (MCM), epi-fluorescence microscopy, confocal microscopy and two-photon microscopy will be briefly introduced. A summarized comparison of these techniques, as discussed by Murphy (Murphy 2002), Jang (Jang *et al.* 2015) and colleagues, is presented in Table 1.

Table 1. *Comparison of common optical microscopy techniques.*

Technique	Phase contrast	DIC & MCM	Epi-fluorescence	Confocal	Two-photon
Principle	Phase shifts caused by differences in optical path length in transparent specimens are converted to amplitude differences.	Relative phase shifts caused by gradients of optical path length in transparent specimens are converted to amplitude differences.	Fluorescence emitted by fluorophores throughout the sample is used instead of transmitted or reflected light.	Fluorescence emitted by fluorophores in small sample volumes is captured by using a pinhole in order to block out-of-focus photons.	Fluorophores are excited based on non-linear optical phenomena to allow confocal operation without a pinhole.
Resolution, lateral/axial	200/500 nm	200/500 nm	200/500 nm	200/500 nm	300/900 nm
Contrast source	Optical path length	Optical path length	Fluorescence	Fluorescence	Fluorescence
Benefits	- Suitable for living, unstained samples - High contrast - High resolution	- Suitable for living, unstained samples - High contrast - High resolution	- Specificity - High contrast - High resolution	- Specificity - Excellent contrast - High resolution - 3D sectioning	- Suitable for thick samples - Minimized signal loss - Low phototoxicity
Limitations	- Not suited for thick samples - Non-specific - Halo effect	- Not suited for thick samples - Non-specific	- Not suited for thick samples - Phototoxicity - Weak signals	- Limited sample thickness - Phototoxicity - High cost	- Relatively low resolution - Limited fluorophore selection - IR heating

2.1.1 Phase contrast microscopy

During the first half of the 20th century, phase contrast microscopes were developed by Zeiss based on the research conducted by Frits Zernike in the 1930s to enhance the contrast of transparent specimens (Murphy 2002). Zernike was awarded the Nobel Prize in physics in 1953 for his discoveries. This technique allowed living cells and other transparent biological samples to be imaged without the need for staining and is today used routinely within the biomedical field. Studying such samples was difficult with conven-

tional brightfield microscopy, which relies on the absorption of light travelling through the specimen, leading to amplitude differences in the image which are then perceived as changes in intensity. Biological samples are usually transparent, that is, they do not absorb visible light, leading to weak contrast when viewed with brightfield microscopy. In the past, such samples could only be examined by using colored dyes. However, even though transparent biological samples do not absorb light, they do diffract light, leading to phase shifts in the light rays that pass through them. Phase contrast microscopy takes advantage of this phenomenon by transforming the differences in phase to differences in amplitude. This is achieved by an optical design which causes interference between light diffracted by the specimen and non-diffracted light, producing the contrast necessary for observing the specimen. Since staining or other sample processing is not needed, even living cells can be imaged in a non-invasive manner at high contrast and resolution.

2.1.2 Differential Interference Contrast and Modulation Contrast Microscopy

The downsides of the phase contrast technique are different halos and other aberrations, which can be problematic especially for quantitative image analysis (Murphy 2002). An improved method for contrast enhancement, DIC, was developed in the 1950s by George Nomarski, followed by MCM, introduced by Hoffman and Gross in 1975. These techniques produce images with a distinctive three-dimensional appearance without the halos of phase contrast microscopy. DIC devices rely on dual-beam interference optics using polarized light and beam-splitting components called Wollaston prisms. The optical system of a DIC microscope creates contrast in the image based on local gradients in optical path lengths. Briefly, the specimen is sampled using coherent wave bundles, that is, closely spaced pairs of light rays generated by a beam splitting component. If there is a gradient in the refractive index or thickness in the region traversed by the pair of light rays, this results in a difference of optical path length and leads to a relative phase shift between the rays. The two wavefronts are then recombined by a Wollaston prism, leading to interference and changes in the amplitude of the light. Brightly illuminated and shadowed edges in the resulting image then correspond to optical path gradients, allowing details of the specimen to be examined. The related MCM technique, based on oblique or off-axis illumination, produces images which appear quite similar to those obtained using a DIC microscope while offering some advantages over DIC, for example, the ability to examine samples on birefringent plastic substrates such as cell culture dishes (Murphy 2002). This comes at the expense of slightly reduced resolution compared to DIC. The phase contrast, DIC and MCM techniques are widely used in modern microscopes.

2.1.3 Epi-fluorescence microscopy

A different family of microscopy techniques, fluorescence microscopy, also gained popularity during the latter half of the 19th century (Murphy 2002). Fluorescence microscopy allows the examination of a single molecular species in the sample. In some cases, it is even possible to detect individual molecules. Fluorescence microscopy relies on the detection of fluorescent light emitted by molecules called fluorophores upon irradiation by excitation light. The enabling technologies were the discovery of methods to conjugate proteins to fluorophores in 1941, the incorporation of interference filters and dichroic mirrors into an epi-illuminator in 1967 as well as the introduction of special high numerical aperture objective lenses, sensitive photographic films and electronic imaging devices.

While some metabolites are inherently fluorescent, the typical procedure involves labeling the non-fluorescent molecules of interest with a fluorescent dye. In the much used technique of immunofluorescence microscopy, antibodies labeled with fluorophores are used to tag target proteins. In the case of widefield fluorescence microscopy, also called epi-fluorescence microscopy, illumination of the whole sample is performed such that a band of shorter wavelength excitation light is directed to the specimen. A filter is used to ensure that the excitation light has the correct wavelength for the fluorophore in question. The fluorescent light emitted by the fluorophore, which has a longer wavelength than the excitation light, is detected with the help of another filter tuned to the emission wavelength of the fluorophore. Moreover, it is possible to use multiple fluorophores with non-overlapping emission wavelength bands to simultaneously visualize several molecular species. The ability to examine molecules of interest with high specificity and resolution has made fluorescence microscopy the most frequently used mode of optical microscopy in modern biomedicine. Countless applications of fluorescence microscopy relying on different fluorescent labels have been developed, allowing, for example, real-time measurement of gene expression or protein binding dynamics. However, the usual requirement of fluorescent labeling and the phototoxicity caused by the excitation light often mean that fluorescence microscopy cannot be used to examine living cells.

2.1.4 Confocal microscopy

A further development of the epi-fluorescence microscope is the confocal microscope (Murphy 2002). The underlying principle of confocal imaging was developed and patented already in 1957 by Marvin Minsky of the Harvard University, but the first commercial instruments only appeared decades later in 1987. Imaging thick fluorescent samples such as round cells or tissue sections with high image quality can be difficult for epi-fluorescence systems due to the out-of-focus light emitted from outside the focal plane. This increases background signals and decreases the contrast of the images. A confocal microscope excludes light originating from out-of-focus planes above and be-

low the features of interest, thereby producing very sharp images and optical sectioning in 3D. This is achieved either by a mechanical spinning disk system or a laser scanning setup, combined with a pinhole detector, allowing the capture of light originating only from a small volume of the sample at a time.

The laser scanning variant includes a scan head, which delivers the excitation light obtained from a laser to the sample using a raster scanning setup. The scan head also features filters which are used similarly to the epi-fluorescence microscope. As the laser scans through the sample area, exciting the fluorophores, the emitted fluorescent light is captured through a pinhole aperture and detected with a photomultiplier tube or another electronic sensor. A laser-scanning confocal microscope is inherently an integrated electronic microscope system, as the actual image is not formed by the optical parts of the device alone but together with electronics and computer software. The detector only measures the intensity of light originating from a single spot of the sample at a time, and the actual image is assembled using a computer. Laser-scanning confocal microscopy is thus an example of the importance of computational methods in biological imaging, as even the acquisition of an image would not be possible without software.

In the spinning disk setup, a rotating disk featuring tiny apertures is used to direct the light to different parts of the sample simultaneously. In these systems, an actual confocal image is generated, because the detection is performed at multiple points of the sample in parallel. The time required for image acquisition is reduced compared to a laser-scanning setup, reducing the amount of photobleaching (*i.e.*, the 'saturation' of fluorophores) and phototoxicity to living samples. On the other hand, the spinning disk system has a lower optical sectioning resolution than the laser-scanning systems and can in some cases limit the use of some fluorophores with ultraviolet (UV) excitation wavelengths. Moreover, confocal systems are rather expensive and share the problem of phototoxicity with epi-fluorescence microscopy. Nevertheless, confocal microscopes are nowadays available in many well-equipped imaging laboratories.

2.1.5 Two-photon microscopy

A newer fluorescence microscopy technique, two-photon microscopy, shares many of the advantages of confocal microscopy, but does not require the pinhole setup (Murphy 2002). This is achieved by means of a non-linear optical phenomenon known as multiphoton excitation. In this process, two or more photons are simultaneously absorbed by a fluorophore molecule and it allows excitation using longer wavelengths than normally. This rare process requires intense illumination, which is provided by pulsed infrared (IR) lasers in two-photon microscopes. Similarly to laser-scanning confocal microscopy, these systems employ a raster scanning mechanism for sample illumination and detection. Careful adjustments of different laser parameters allow the excitation of fluorophores only within the focal plane, improving the signal-to-noise ratio of the system. A pinhole setup is therefore unnecessary, as there is no out-of-focus fluorescence

to be excluded, increasing the sensitivity of the device. Other benefits are that the rate of photobleaching and phototoxicity are reduced, as the excitation is limited to small volumes of the sample at a time and the less energetic, longer wavelength excitation light is less damaging to the sample. The possibility of using near-IR wavelengths for excitation also means that relatively thick specimens can be imaged due to the ability of infrared light to permeate deeply into biological material. The penetration depth of a two-photon microscope can thus be even tenfold compared to a confocal microscope. Limitations in laser technology, in turn limiting the use of different fluorophores, and the heating of water induced by the IR light still limit the use of two-photon microscopy, but the technology has already proved useful in many applications.

2.2 Spectral analysis of images

2.2.1 Introduction

Spectral analysis is a field of signal processing that aims to characterize the frequency content of a measured signal (Marple 1987). The mathematical foundation for relating a signal in the temporal or spatial domain with its counterpart in the frequency (or spatial frequency) domain is the Fourier transform. Only estimates of the frequency content, typically expressed as power spectral densities (PSD), can be obtained based on a finite observation of a signal. The power spectral density is often also called the power spectrum or simply spectral density. It describes how the variance of a random process is distributed with frequency (and orientation in the case of images). Spectral estimation refers to this process of constructing an estimate of the PSD of a random signal based on a recorded sample of the signal.

In the context of microscopy, the signal of interest is typically a two-dimensional micrograph, defined on a grid of spatial locations. In this case, the amplitude and phase of the signal's constituents are obtained as a function of orientation and spatial frequency (Gonzalez & Woods 2001). Features that vary slowly along the dimensions of the image, such as decreases in illumination towards the edges of the image or the alternation between dark background pixels and the bright cytoplasm of fluorescently labeled cells, are represented at the low frequency end of the spectrum. On the other hand, features that exhibit varying intensity at shorter length scales, such as the striated pattern of bright and dark regions along the length of a myofibril, are captured at intermediate frequencies in the spectral representation. The high-frequency end of the spectrum often only contains contributions from noise sources and different processes causing variation at very short length scales. Signals from natural sources rarely exhibit perfect sinusoidal periodicity or alignment along a single direction. For this reason, their spectral representations typically feature cloud-like distributions spread over a range of frequencies and orientations. By analyzing the spectrum, one can estimate orientation and spatial frequency distributions of different features present in the image.

An example image and its power spectrum in Cartesian and polar coordinates are shown in Figure 1A, Figure 1B and Figure 1C, respectively. The image (Figure 1A) features an epi-fluorescence micrograph of a hiPSC-CM at 40X magnification with the intracellular fibrils visible in the green channel (see Section 4.2 for details). The characteristic periodic pattern is clearly visible in the magnified view. An estimate of the PSD of the green channel image (Figure 1B) has been obtained using the so called Welch's spectral estimator (discussed in Section 2.2.4) and shifted to place the zero frequency component in the center (see Section 2.2.3 for details). Furthermore, the values of the power spectrum (given in arbitrary units) have been $\log_{10}$-transformed to increase contrast for visual examination of the power spectrum. Otherwise, the low frequency components tend to dominate the visual appearance of the PSD. The power spectrum expressed in Cartesian coordinates shows the low frequency components of the spectrum near the origin. As will be discussed in Section 2.2.3, the orientation of spectral features is usually orthogonal to the orientation of corresponding structures in the image. The overall shape of the cell, also partly captured in the green channel, gives rise to the bright spectral region (arrow 1) stretching out from the origin at an angle that can be seen to be approximately transverse relative to the longitudinal orientation of the cell. This region also includes higher spatial frequency components, which probably originate from the intensity variations in the image caused by the fibrils. The spectrum also features a distinct 'cloud' both above (arrow 2) and below the origin at an approximate spatial frequency of 0.5 μm^{-1}. Based on its orientation and the spatial frequency corresponding to a wavelength of 2 micrometers, this feature can be rather safely assumed to represent the sarcomere subunits of the myofibrils. They are protein structures with an approximate size of 2 µm, spaced regularly along the fibril (Pasqualini *et al.* 2015). In this case, the orientation of the spectral region reflects the intensity variation caused by these structures along the longitudinal orientation of the fibrils. The high-frequency end of the spectrum does not contain any biologically interesting components in the case of this image.

The features discussed above can often be more easily examined by transforming the power spectrum into polar coordinates (Figure 1C), as discussed in Section 2.2.3. In this representation, the same spectral information is given as a function of orientation and spatial frequency. Starting counter-clockwise from the positive x-axis, the upper half circle of Figure 1B is shown in Figure 1C. The same low-frequency shape seen in Figure 1B can now be observed in the 0°-20° orientation range at spatial frequencies ranging from zero to 0.5 μm^{-1} in Figure 1C (arrow 1). Similarly, the spectral region corresponding to the sarcomere subunits can be seen as a cloud present at spatial frequencies slightly over 0.5 μm^{-1} and orientations ranging from approximately 60° to 130° (arrow 2). This qualitative example demonstrates how power spectra can be utilized for analyzing the orientation and wavelength distributions of structures in images. However, obtaining an accurate estimate of the PSD is far from trivial, as will be discussed in the following sections of this chapter.

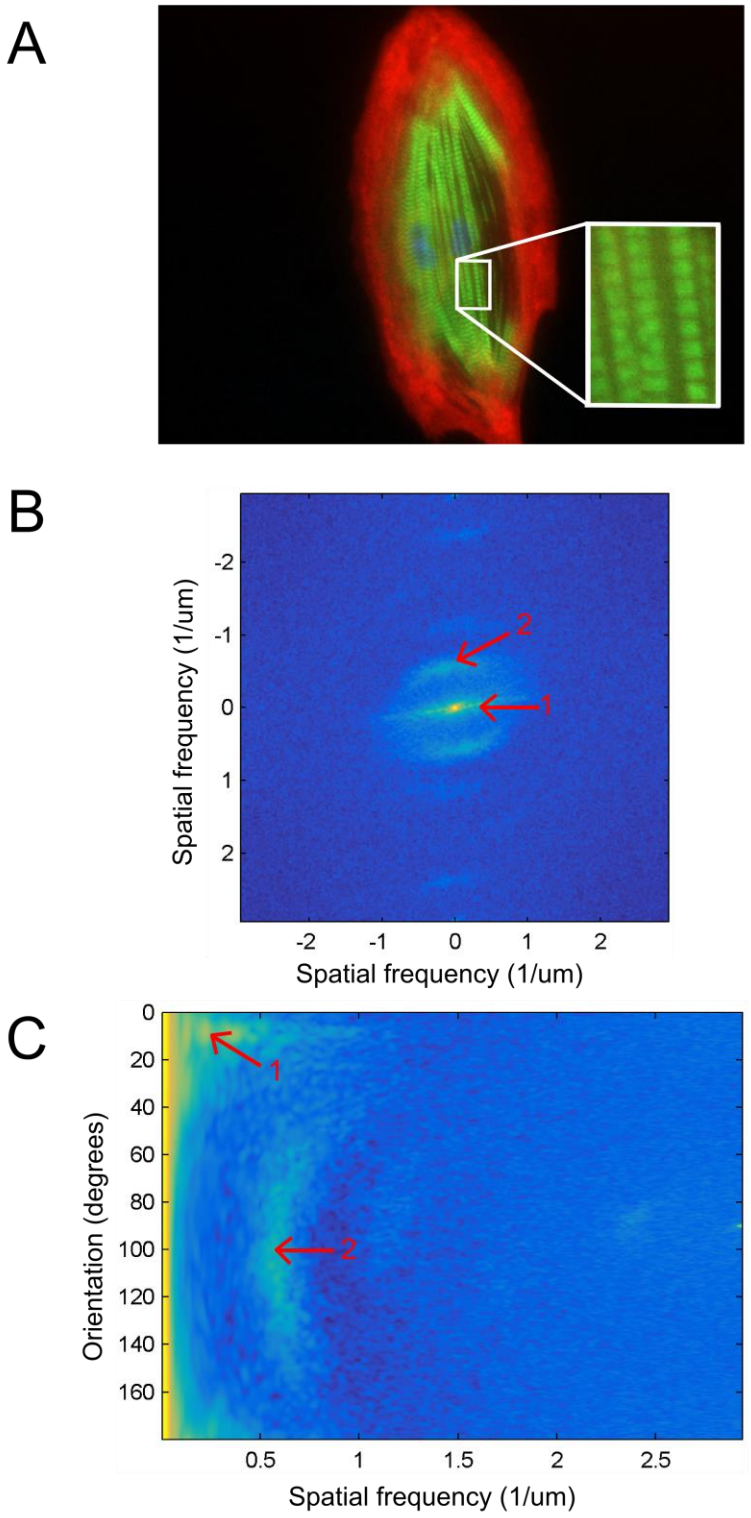

Figure 1. *An example epi-fluorescence micrograph at 40X magnification depicting a hiPSC-CM (A) and its power spectrum expressed in Cartesian (B) and polar coordinates (C). A magnified view of the periodic sarcomere subunit pattern is shown in (A). The power spectral values have been* $\log_{10}$*-transformed and the origin of B has been centered. Only the upper half-circle of B is shown in C. Corresponding spectral regions in Cartesian and polar representations, arising from the myofibrils and the cell shape (arrow 1) and from the sarcomere subunits (arrow 2), are indicated with red arrows.*

Methods for spectral estimation can be divided into two main classes: parametric and non-parametric methods (Marple 1987). Parametric methods rely on the assumption that the underlying stochastic process can be approximated by a parametric model and the spectral estimation task is then to find suitable values for the parameters. In contrast, methods of non-parametric spectral estimation, also referred to as classical spectral estimation, do not assume the process to have any particular structure. In the case of natural images, classical methods are often preferred. Different methods employ different measures to address the many issues and limitations faced during spectral estimation, but they share the common mathematical basis of the discrete Fourier transform (DFT). These issues and the properties of the DFT are briefly introduced in Sections 2.2.2 and 2.2.3, respectively. Methods of classical spectral estimation are then described in Section 2.2.4.

2.2.2 Problems in spectral estimation

Spectral estimation, that is, forming an estimate of the PSD of a random process based on samples (and possible prior knowledge of the process) is not a simple task and constitutes an entire field of research on its own (Marple 1987). Successful spectral estimation requires addressing a number of important problems and limitations, which often means making compromises between conflicting performance goals based on the requirements of the application at hand. Different spectral estimation methods have been proposed to solve these issues. However, despite attempts and eventual progress in the development of the theory of spectral estimation, the reality still is that spectral estimation often relies on empirical considerations and rules of thumb. This is mainly due to the difficulty of devising objective, universal performance measures for spectral estimators. There are, however, some key issues, which all spectral estimation methods somehow have to take into account. These issues are briefly highlighted in this section.

One of the most important performance criteria for spectral estimators is spectral resolution (Marple 1987). In qualitative terms, spectral resolution, often also called frequency resolution, refers to the ability of a spectral estimator to distinguish individual, closely spaced spectral components. Closely spaced in this context means signal components, whose frequency content (or dominant frequency) is very similar but still distinct. The undesirable phenomenon of multiple distinct spectral components merging into one due to insufficient resolution is often called smearing. Different formal definitions for spectral resolution have been proposed, but it is still quite common to convey the relative spectral resolutions of two spectral estimators based on visual comparisons of the spectra they produce. Such comparisons are, however, far from reliable. Spectral resolution is inherently linked to the length of the signal and is approximately the reciprocal of this value according to the commonly used rule of thumb. Thus, a longer recorded signal having a larger number of data samples (assuming fixed sampling frequency) allows better spectral resolution to be obtained.

Another important property of spectral estimators is their ability to detect signal components, that is, to reliably indicate the presence (or absence) of a particular signal component (Marple 1987). Detecting signals in the presence of noise can be problematic or sometimes even impossible if the signal-to-noise ratio (SNR) is very low. In some cases, a spectral estimate which more closely corresponds to the true spectrum may in fact be worse in terms of signal detectability than another, otherwise less truthful estimator. For example, many spectral estimation methods aim to reduce random fluctuations in the spectrum by different smoothing approaches. This can improve the performance of the spectral estimator in the sense that the smoothed estimate more closely resembles the actual spectral shape, which would otherwise be corrupted by random errors due to the random nature of the spectral estimate. However, such smoothing procedures can mask weak signals which can no longer be separated from the background noise. Thus, whether the smoothed estimate is better than the estimate with random fluctuations depends on the relative importance of detecting individual weak signals at certain frequencies and obtaining a truthful estimate of the true spectrum over a larger range of frequencies. This example represents one of the commonly encountered situations, where the performance of a spectral estimator depends on the application and cannot be measured in a universal manner.

The assumption of an infinite input signal, dictated by the properties of the DFT, is of course a practical impossibility (Marple 1987). The inevitable violation of this assumption has some consequences for spectral estimation. First of all, because a spectral estimate formed on the basis of a finite record of data is a random variable, it is not an exact reproduction of the true spectral function, which could only be obtained if an infinite sample of the process of interest were available. This forces spectral analysis to rely on estimates. Because the estimate is a random variable, it must be characterized in terms of its statistical properties. The performance of spectral estimators is typically assessed by their bias and variance. A high bias indicates that the values of the spectral estimate differ strongly from the values of the true spectral function at the corresponding frequencies. On the other hand, a high variance indicates strong random fluctuations in the estimated spectrum. Discriminating true signals from the fluctuations of high-variance spectra can be a difficult task. It turns out that for a given signal length M, any successes in decreasing the variance of a spectral estimate are accompanied by poorer spectral resolution. Conversely, if the length of the data recording is not changed, better spectral resolution can only be obtained at the cost of increased variance. This important ‘uncertainty principle’ is sometimes referred to as the bias-variance dilemma or the stability-time-bandwidth product (Marple 1987; Stoica & Moses 2005). It should be noted, that when referring to the signal length M, the sampling frequency is assumed to be fixed. That is, increasing the sampling frequency to obtain a larger number of samples per an interval of time (or space, in the case of images) does not have an effect on the attainable combination of spectral resolution and stability of the spectral estimate. Only the actual time or space interval of the data sequence has an effect on the time-bandwidth-

product. Typically, the amount of data is fixed, and a compromise between spectral resolution and variance has to be agreed on.

The second consequence of using finite sequences of data instead of the infinite input signal is called leakage (Marple 1987). The finite data sequence can be considered a portion of the theoretical infinite sequence, seen through a rectangular window function. Due to the properties of the DFT, the Fourier transform of a signal implicitly windowed with such a rectangular window will have so called sidelobes, or replications of the shape of the window transform. They will bias the adjacent values of the spectrum due to the 'leaking' of power to a range of frequencies. In addition to increasing bias, leakage may also mask weak signals if they happen to reside on the same band of frequencies. The effects of leakage can be mitigated by windowing the data prior to computing the DFT. Windowing refers to the process of multiplying the data sequence with a so called window function. The window functions, also called tapers, are designed such that they smoothly decrease towards zero at the edges of the data sequence, rather than abruptly cutting the sequence. Decreasing the sidelobes of the window transform can only be achieved by simultaneously broadening the so called mainlobe of the window. This results in the reduction of spectral resolution. Window design thus represents a compromise between spectral resolution and bias caused by leakage. Commonly used windowing functions include, for example, the Bartlett window (triangle), Hann window (squared cosine), Hamming window (raised cosine), Nuttall window (weighted cosines), Gaussian window and Chebyshev window (equiripple).

Another limitation related to the statistical properties of the spectrum is, that it is a function of the second-order statistics of the signal, which are in addition assumed to be stationary, that is, constant over time (Marple 1987). This means that information in higher-order statistics of the process is not included in the spectrum. Moreover, many real signals vary over time and the assumption of stationarity is thus violated. However, this variation is often not significant over short periods of time and the signals can therefore be considered locally stationary. The non-stationarity of the process over longer intervals can then be taken into account by performing spectral estimation using short segments of the whole signal and observing the variation from segment to segment.

2.2.3 The discrete Fourier transform

The discrete Fourier transform, implemented with computationally efficient Fast Fourier Transform (FFT) algorithms, lies at the heart of classical spectral estimation (Gonzalez & Woods 2001). The Fourier transform is also applicable to continuous functions, but since the continuous Fourier transform has no practical relevance for this study, this discussion is limited to the discrete case. Moreover, instead of a lengthy discussion of all of the mathematical properties of the Fourier transform, only those aspects which are of relevance for this work are included in this section. The mathematical foundation for the Fourier transform is the Fourier series, introduced by the French mathematician Jean

Baptiste Joseph Fourier in 1822. His discovery implied that any integrable periodic function, no matter how complex it is, could be expressed as a Fourier series, that is, a sum of sines and/or cosines of different frequencies, each weighted by a different coefficient. Even non-periodic functions can be expressed as the integral of weighted sines and/or cosines of different frequencies, in which case the term Fourier transform is used instead of Fourier series. Moreover, it is possible to reconstruct the original function via an inverse transform without any loss of information, allowing signals to be processed in the Fourier domain, also called the frequency domain, before transforming them back to the temporal or spatial domain. The Fourier transform can be considered as a 'mathematical prism', which separates a function into a number of frequency components, similarly to an optical prism that is able to separate visible light into a spectrum of color components. In the late 1950s, the introduction of digital computers and computational methods collectively referred to as FFT algorithms revolutionized the field of signal processing and allowed the practical application of the Fourier transform in a myriad of applications.

The discrete Fourier transform $F(u)$ of a function $g(x)$, $x = 0, 1, 2, \ldots, M\text{-}1$, is defined by (Gonzalez & Woods 2001):

$$F(u) = \frac{1}{M}\sum_{x=0}^{M-1} g(x)e^{-j2\pi ux/M} \tag{1}$$

for $u = 0, 1, 2, \ldots, M - 1$. The inverse DFT is defined by:

$$g(x) = \sum_{u=0}^{M-1} F(u)e^{j2\pi ux/M} \tag{2}$$

for $x = 0, 1, 2, \ldots, M - 1$. For finite values of $g(x)$, the discrete Fourier transform pair always exists. In the case of image processing, ensuring that all values of the signal are finite is straightforward and we thus make this assumption in the remainder of this chapter. In a spectral analysis context, mean (or some higher-order model of trend) is usually subtracted from the signal prior to the DFT computation (Marple 1987). This process is often called prewhitening. Failure to remove large sample means or significant trends can result in biased or distorted spectral estimates.

In the case of 2D image processing, the signal is a function of two spatial variables and the one-dimensional DFT and inverse DFT must be therefore extended to two dimensions. This is straightforward, and the two-dimensional DFT of a signal array $g(x,y)$ having dimensions M x N is given by (Gonzalez & Woods 2001):

$$F(u,v) = \frac{1}{MN}\sum_{x=0}^{M-1}\sum_{y=0}^{N-1} g(x,y)e^{-j2\pi(ux/M+vy/N)} \tag{3}$$

for $u = 0, 1, 2,\ldots,M-1$ and $v = 0, 1, 2,\ldots, N-1$. The inverse DFT in two dimensions is given by the equation:

$$g(x,y) = \sum_{u=0}^{M-1}\sum_{v=0}^{N-1} F(u,v)e^{j2\pi(ux/M+vy/N)} \tag{4}$$

Due to the separability of both the forward and inverse transforms, the 2D Fourier transform can be performed by first computing the 1D transform along one dimension of the input, followed by the computation of another 1D transform along the other dimension. For this reason, the DFT can be efficiently computed using FFT algorithms also for 2D images. In the case of an image the variables x and y correspond to the coordinates of image pixels on a plane and the function $g(x,y)$ indicates the intensity, or brightness, of the image at each pixel position. The values of M and N correspond to the width and height of the image. The variables u and v are spatial frequency variables indicative of the locations of different frequency components in the Fourier domain. The complex function $F(u,v)$ contains information about the amplitude and phase of each component.

As said, $F(u,v)$ is complex and can be expressed in terms of magnitude (Gonzalez & Woods 2001):

$$|F(u,v)| = \sqrt{\mathrm{Re}(F(u,v))^2 + \mathrm{Im}(F(u,v))^2} \tag{5}$$

where $\mathrm{Re}(F(u,v))$ and $\mathrm{Im}(F(u,v))$ are the real and imaginary parts of $F(u,v)$, respectively. This is called the Fourier spectrum and indicates the amplitude of each frequency component. That is, values of $|F(u,v)|$ correspond to the weighting coefficients assigned to each frequency component in the Fourier transform process. Since $F(u,v)$ is complex, it can also be expressed in terms of phase angle:

$$\Phi(u,v) = \tan^{-1}\frac{\mathrm{Im}(F(u,v))}{\mathrm{Re}(F(u,v))} \tag{6}$$

This is called the phase spectrum and it contains information on the phase angle of each frequency component. In the case of images, the phase is extremely important for the visual appearance of the image (Jahne 2004). The phase indicates the relative location of periodic structures, as opposed to the magnitude, which only carries information on the presence or absence of corresponding structures in the image. Without the phase information, visual perception of images becomes difficult if not impossible.

After computing the two-dimensional DFT, it is customary to perform a shifting operation so as to place the zero frequency component $F(0,0)$ to the middle of the M x N array (Gonzalez & Woods 2001). The operation swaps the top-left quadrant of the array with the bottom-right quadrant, and the top-right quadrant with the bottom-left quadrant. The resulting array still has the same dimensions M x N as the original image with

the zero frequency, or dc component, located at the coordinates $u = M/2+1$ and $v = N/2+1$, assuming that M and N are even numbers. An example power spectrum before and after the shifting operation is shown in Figure 2. In the shifted form, increasing radial distance from the dc component at the origin of the array corresponds to increasing spatial frequency. The component at $F(0,0)$ is termed the dc component due to the fact that its value equals the constant component of the signal, that is, the average intensity of the image obtained as the mean of $g(x,y)$ over all x and y. The elements of $F(u,v)$ corresponding to the Nyquist frequency f_{Nyquist}, that is, the highest frequency, which can be sampled without aliasing artefacts are located at the coordinates $(0, N/2+1)$, $(M, N/2+1)$, $(M/2+1, 0)$ and $(M/2+1, N)$ of the shifted array. These coordinates correspond to the midpoints of the sides of the M x N array. If we make the practical assumption that $M = N$, the region of $F(u,v)$ corresponding to spatial frequencies equal to or below the Nyquist limit is a circle with a diameter of $M = N$. This means that the corners of the rectangular array contain spectral components whose spatial frequency exceeds the Nyquist limit. These elements are usually excluded from analysis due to their complicated interpretation. Moreover, having $M \neq N$ implies that the sampling frequency and correspondingly the Nyquist frequency are different along the horizontal and vertical direction of the image. This further complicates the analysis of the spectrum and it is often easier to ensure that $M = N$.

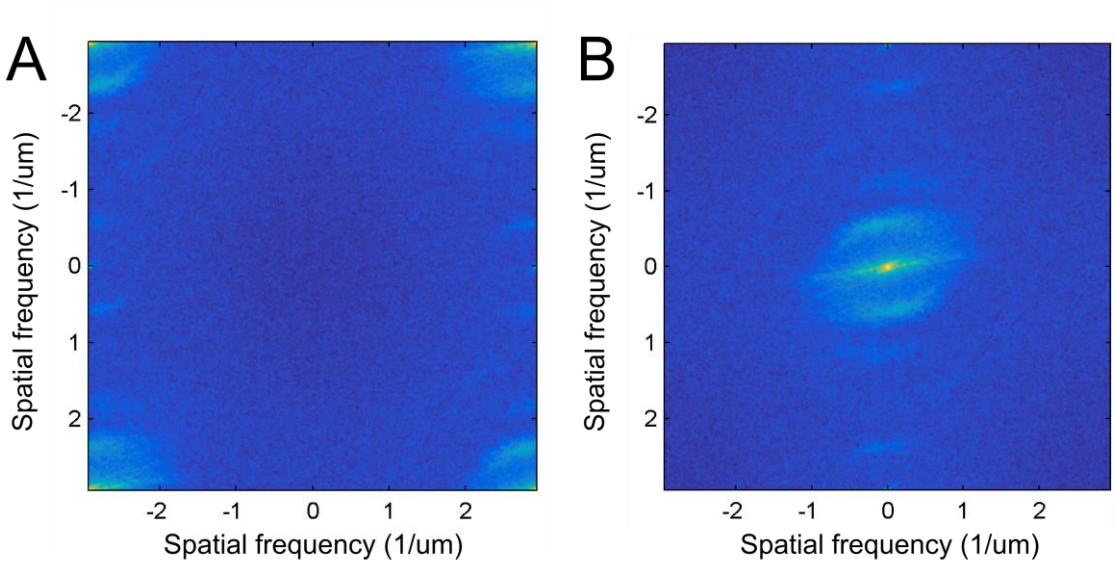

Figure 2. *An example power spectrum before (A) and after (B) the shifting operation.*

In addition to the spatial frequency information, $F(u,v)$ also conveys information about the orientation of image features (Gonzalez & Woods 2001). The use of the Fourier transform for orientation analysis is based on the rotation property of the transform, which states that rotating the image by any angle rotates the Fourier transform of the image by the same angle. The orientation of spectral features relative to the coordinate axes corresponds to the orientation of features in the original image with an added 90

degree offset in most cases. This is due to the fact that in most cases the orientation of object edges in the image is actually orthogonal relative to the orientation of the object itself. The reason for referring to orientation instead of direction is the conjugate symmetry property of the Fourier transform. That is,

$$F(u,v) = F^*(-u,-v) \tag{7}$$

where "*" indicates the complex conjugate operation. This means that

$$|F(u,v)| = |F(-u,-v)| \tag{8}$$

indicating that the magnitude of the complex array obtained after performing the DFT and the shifting operation is symmetric about the origin, that is, a single half-circle already contains all the useful information, while the other half-circle is a redundant mirror image (see Figure 1). This means that the values of the spectrum at any spatial frequency repeat every 180° and are therefore indicative of orientations rather than directions. Another important property of the Fourier transform is that it is periodic with periods *M* and *N* along the respective coordinate axes. The periodicity property together with the conjugate symmetry of the Fourier transform explains why it is perfectly valid to perform the shifting operation described earlier.

To simplify the handling of the spectrum in an orientation analysis context, the original 2D array, expressed in Cartesian coordinates using the variables *u* and *v*, can be alternatively expressed in polar coordinates as a function of orientation θ and spatial frequency *f* by making the following polar transform for each element of the original array:

$$\theta = \tan^{-1}\frac{v}{u} \tag{9}$$

$$f = \sqrt{u^2 + v^2} \tag{10}$$

Moreover, due to the circular symmetry of the 2D spectrum, the density of spectral elements is different at different spatial frequencies, that is, at different distances from the origin. For this reason, the values of the power spectrum in polar form need to be adjusted by a scaling factor that is dependent on spatial frequency. This is performed by multiplying each element of the polar transformed power spectrum with the scaling factor given by:

$$\frac{\pi}{2}\frac{f}{f_{Nyquist}} \tag{11}$$

using the value of *f* corresponding to the spectral element in question. Due to the conjugate symmetry property, there is only need to retain one half-circle of the original array for orientation analysis purposes. Since the array is symmetric about the origin, the circle could be halved at any arbitrary angle. In this study, the upper half-circle is retained.

The interpretation of the angle θ is similarly arbitrary. In this study, 0° corresponds to a horizontal orientation with the other angles increasing counter-clockwise. That is, 45° corresponds to an orientation between North-East and South-West, 90° corresponds to a North-South orientation and 135° corresponds to an orientation along the axis between North-West and South-East. The value of 180° is redundant and is equivalent to 0°. Therefore, according to these conventions, after performing the polar coordinate transform, only the elements having values of θ in the interval [0°, 180°) are retained. In terms of spatial frequency, only elements with $f \leq f_{\text{Nyquist}}$ should be retained. This means that the corners of $F(u,v)$ are discarded in the process and only the upper half-circle region is retained for further analysis. An example of power spectra expressed in Cartesian and polar coordinates was presented Figure 1. Because the polar coordinates do not necessarily correspond to integer Cartesian coordinates, the transform was performed using bilinear interpolation.

The practical computation of the DFT relies on FFT algorithms (Gonzalez & Woods 2001; Stoica & Moses 2005). A direct evaluation of Equation (1) would require on the order of M^2 complex multiplications and additions, which can lead to a very high number of operations for large M. Any procedure capable of computing Equation (1) using less than M^2 flops (1 flop = 1 complex multiplication + 1 complex addition) is classified as an FFT algorithm. The first FFT algorithms were developed in the 1960s, followed by a stream of more and more computationally efficient solutions. The detailed principles of these algorithms are outside the scope of this text, but the central concept of the FFT is to divide the DFT into two or more smaller DFTs, which are computed individually and then linearly combined to obtain the final result. Each smaller DFT can in turn be divided to even smaller subsequences. One of the first algorithms, the radix-2 FFT, requires around ½ $M \log_2 M$ flops to compute an M-point DFT. For comparison, the more recent split-radix FFT takes approximately ⅓ $M \log_2 M$ flops. These solutions require M to be an integer power of two. Since the 2D FFT can be performed via two consecutive passes of the 1D FFT, fulfilling this requirement in the two-dimensional case simply means that both M and N need to be integer powers of two. There are also algorithms capable of handling cases where this condition is not satisfied, but this comes at the expense of computational efficiency.

Implementations taking advantage of multiple FFT algorithms in order to combine computational efficiency and flexibility in terms of signal length are also available. For example, the FFT implementation included in MATLAB utilizes a combination of several algorithms, including the radix-2 FFT, the split radix FFT and a number of other FFT algorithms depending on the value of M (The MathWorks, Inc., 2015). The computation is fastest for integer powers of two and almost as fast for values of M which have only small prime factors. Other lengths can also be handled, but the execution time is several times slower in those cases. Even though it is not strictly necessary with modern hardware and FFT algorithms, it is often straightforward to force M to be an integer

power of two using a procedure known as zero padding (Stoica & Moses 2005). Zero padding involves appending the original signal with a number of zeros to increase *M*, which correspondingly also increases the number of elements in the Fourier domain. However, even though zero padding prior to FFT computation produces spectra defined on a finer grid of frequencies, the spectral resolution is not improved in the process. Instead, the process of zero padding simply corresponds to interpolation in the frequency domain and is in most cases only of aesthetic value. However, zero padding allows the fastest FFT methods to be used even for signals whose length is originally not an integer power of two. The smoother appearance of the spectra may also sometimes aid visual interpretation. Correct use of zero padding is also crucial if one wishes to process images by filtering in the frequency domain and then return them to the spatial domain via the inverse DFT (Gonzalez & Woods 2001).

2.2.4 Classical spectral estimation methods

Classical spectral estimation methods are based either on a direct or an indirect theoretical approach, each one adopting a different definition of the PSD (Marple 1987). These two alternative definitions can, under certain assumptions, be considered equivalent. The indirect approach is based on computing the autocorrelation sequence of an infinite sequence of data and then Fourier transforming the result to obtain the PSD. The direct approach, on the other hand, is based on obtaining the PSD as the squared modulus of the Fourier transform of the infinite data sequence. Methods based on the direct definition of the PSD are collectively termed periodogram methods, while estimators based on first forming the autocorrelation sequence estimate are called correlogram methods. Early dominance of the correlogram-based Blackman-Tukey method has vanished in favor of periodogram approaches due to the efficiency of modern FFT algorithms and hardware. For this reason, the rest of this section will focus on four alternative periodogram methods: the raw periodogram, Daniell's method, Bartlett's method and WOSA. The raw periodogram will yield statistically inconsistent PSD estimates because it does not take into account the random nature of obtaining finite samples from the process under study. The other three methods, on the other hand, employ different averaging approaches to yield smoother and statistically more stable spectral estimates based on a limited amount of data. The challenge of obtaining spectral estimates with the highest possible spectral resolution and statistical reliability based only on a finite amount of data samples means that trade-offs have to be made between different objectives, based on the desired properties of the spectral estimator. All of the methods presented here, apart from the raw periodogram, yield comparable statistical performance and differences in the spectral estimates obtained using different methods are relatively small. For this reason, the method of choice has been often selected based on ease of implementation and computational efficiency. Moreover, due to fundamental differences in 1D and 2D linear systems theory, many approaches originally introduced for the one-dimensional case cannot be simply extended to handle functions of two or more varia-

bles. The four methods described here can, however, be easily extended to handle the 2D case. In addition to the tried and tested nature of these estimators, this is another reason why many of the methods presented here are still widely used despite the fact that they date from the mid-1900s.

The simplest of all periodogram methods is the raw periodogram (Marple 1987). In practice, infinite samples of the random process are not available and the infinite sequence of data is thus replaced by a finite sequence of data samples, which are Fourier transformed by DFT according to Equation (3). The PSD is then simply given by:

$$P(u,v) = |F(u,v)|^2 = \mathrm{Re}(F(u,v))^2 + \mathrm{Im}(F(u,v))^2 \tag{12}$$

It turns out, that this spectral estimator suffers from high variance. In the limit of infinite data, the mean of the raw periodogram estimate will tend to the true value of the PSD, but the variance will then be comparable to the mean in magnitude. This causes the spectral estimate to feature a considerable amount of random errors, rendering the raw periodogram a poor spectral estimator. An example PSD estimated from the green channel of the image in Figure 1A using the raw periodogram method is shown in Figure 3. The high variance is visible as random fluctuations across the spectrum.

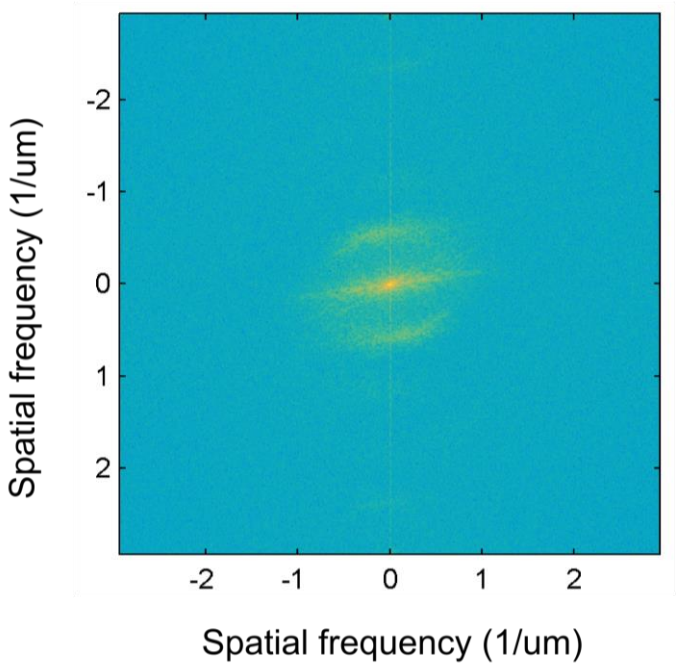

Figure 3. *Raw periodogram estimate of the PSD of the green channel of Figure 1A.*

Daniell's method smooths the raw periodogram by averaging over adjacent frequency bins of the spectral estimate (Daniell 1946; Marple 1987). Originally, Daniell's method was proposed in the form of replacing the values of each spectral element with the arithmetic mean of adjacent elements. The more general form of the Daniell periodogram, also applicable to the 2D case, can be expressed in the form of a convolution as:

$$P_{Daniell}(u,v) = P(u,v) * H(u,v) \quad (13)$$

where $H(u,v)$ is the frequency response of a low pass-filter. Daniell's method thus simply corresponds to low-pass filtering the raw periodogram with a suitable filter. Selecting a suitable cutoff wavelength and other parameters for the filter is an application-dependent problem. Low-pass filtering the spectrum will attenuate some of the spurious random fluctuations of the raw periodogram, but it will also lead to a decrease in spectral resolution due to the smoothing of spectral features. A compromise between loss of resolution and decreased variance has to be thus made when designing the low-pass filter. An example PSD estimated from the green channel of the image in Figure 1A using Daniell's method is shown in Figure 4. An averaging filter with a size of 10 x 10 pixels was used. In comparison to the periodogram estimate, the random fluctuations have been reduced but the averaging process has also affected the spectral resolution, which can be seen as blurring of the bright regions close to the origin. The decreased variance has also revealed a bright vertical line artefact, which is caused by leakage due to the lack of windowing in Daniell's method. The same artefact was faintly visible also in Figure 3.

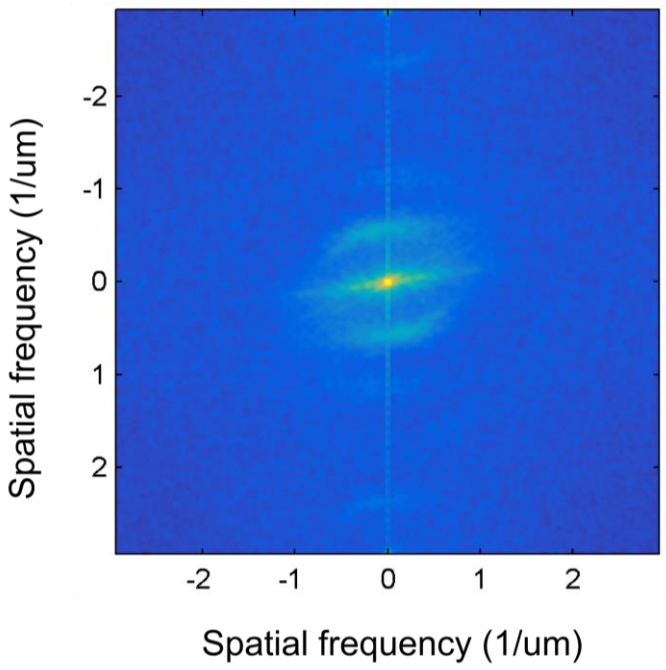

Figure 4. *Daniell's estimate of the PSD of the green channel of Figure 1A. An averaging filter with a size of 10 x 10 pixels was used.*

Bartlett's method averages multiple periodograms which have been formed by segmenting the original data sequence into multiple subsequences, or segments (Bartlett 1948; Stoica & Moses 2005). This is a form of pseudo ensemble averaging, where the finite amount of available data is simply divided into multiple portions to emulate the situation of actually having multiple data recordings of the same process. In the case of images, it is often practical to use square segments, as this removes the complication of

having horizontal and vertical frequency axes of different length in the Fourier domain, as discussed in Section 2.2.3. The image is then divided into N_S subimages, and Bartlett's spectral estimate is given by:

$$P_{Bartlett}(u,v) = \frac{1}{N_S}\sum_{j=1}^{N_S} P_j(u,v) \tag{14}$$

where P_j (u,v) is the raw periodogram of segment j, computed according to Equation (12). Because the size of the segments is inversely proportional to their number N_S and the spectral resolution is approximately inversely proportional to segment size, as discussed in Section 2.2.2, it follows that Bartlett's method leads to an approximate loss of resolution by a factor N_S when compared to the resolution of the raw periodogram. Moreover, it has been shown that Bartlett's method correspondingly decreases the variance of the spectral estimate by the same factor N_S. Therefore, Bartlett's method represents a case of trading off some spectral resolution for an improvement in the statistical stability of the spectral estimate. The balance between these two objectives can be tuned by changing the size of individual segments, which essentially corresponds to changing the total number of segments N_S. An example PSD estimated from the green channel of the image in Figure 1A using Bartlett's method is shown in Figure 5. Square segments with a side length equal to 1/3 of the image's shorter dimensions were used. The random fluctuations are again less severe than in the case of the raw periodogram. Spectral resolution appears better than in the case of Daniell's method, at least with the parameters used here. The leakage artefact, manifested as bright horizontal and vertical lines, is severe in this case. The raw periodogram, Daniell's method and Bartlett's method all suffer from leakage, but in the case of the periodogram the artefacts are partly hidden beneath random noise. In the case of Daniell's method, the effect is partly remedied by the smoothing procedure at the expense of spectral resolution.

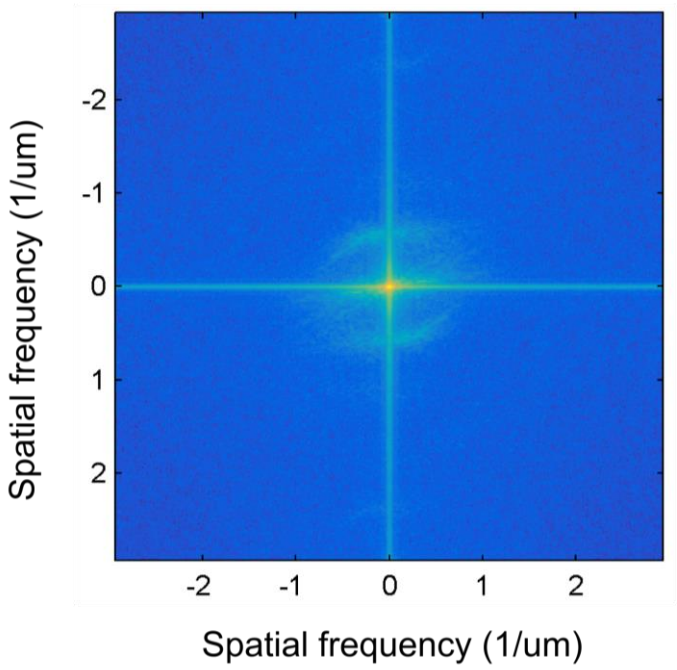

Figure 5. *Bartlett's estimate of the PSD of the green channel of Figure 1A. Square segments with a side length equal to 1/3 of the image's shorter dimensions were used.*

One of the most widely used classical spectral estimation methods is Welch's method, frequently referred to as Weighted Overlapped Segment Averaging or Welch's Overlapped Segment Averaging (WOSA) (Welch 1967; Stoica & Moses 2005). It is an improvement over Bartlett's method. The difference between Bartlett's method and WOSA is that in the case of WOSA, each of the periodograms is modified using a windowing function that decays smoothly towards the edges of the segment. Another difference is that in WOSA, the signal segments are allowed to overlap. Thus, the WOSA procedure is as follows. First, the signal is divided into N_S segments, overlapping by some percentage of the segment size. Unless the overlap percentage is set to zero, N_S will thus be larger than the corresponding value using Bartlett's method, assuming the segments have identical size. Again, it is practical to use square segments for two-dimensional data arrays such as images for reasons discussed above. Next, each segment is multiplied by a windowing function $w(x,y)$ and its raw periodogram is computed according to:

$$\hat{P}_j(u,v) = \left| \frac{1}{M_W N_W} \sum_{x=0}^{M_W-1} \sum_{y=0}^{N_W-1} w(x,y)\, g_j(x,y) e^{-j2\pi(ux/M_W + vy/N_W)} \right|^2 \qquad (15)$$

where $g_j(x,y)$ is the jth subimage segment and M_W and N_W refer to the dimensions of the segments rather than the dimensions of the whole image. This equation is nothing more than a combination of Equation (3) and Equation (12), applied to a windowed data segment. The WOSA spectral estimate is then given by:

$$P_{WOSA}(u,v) = \frac{1}{N_S}\sum_{j=1}^{N_S}\hat{P}_j(u,v) \tag{16}$$

WOSA has a number of advantages compared to Bartlett's method. Using the windowing function prior to computing the DFT removes some of the bias caused by leakage, discussed in Section 2.2.2. The ability of using different window functions also allows some control over the bias/resolution properties of the spectral estimate. Allowing the segments to overlap, on the other hand, serves the purpose of further reducing the variance of the spectral estimator by allowing a larger number of subimages to be averaged. The overlapping also balances the uneven weighting of input samples caused by the windowing function, provided that the overlap percentage of the segments is at least 50 %. In other words, the windowing process, despite mitigating the effects of leakage, causes bias by increasing the relative importance of signal samples close to the center of the window while decreasing the importance of data near the window edges. Overlapping of the segments by at least 50 %, however, ensures that all data samples get both high and low weightings in different segments. For example, in the case of exactly 50 % overlap, a sample located right at the edge of one windowed segment would be at the center of an adjacent data segment. Furthermore, WOSA offers additional flexibility due to the fact that the averaging operator applied on the spectral estimates need not be the simple arithmetic mean but the estimates computed from different segments can instead be unevenly weighted based on, for example, the properties of the corresponding segment of the input signal. This way, the effects of signal artefacts, for example, can be decreased. Analyzing the variance of the WOSA estimator analytically is complicated, but there is empirical evidence that WOSA offers lower variance than Bartlett's method. Due to these advantages, WOSA still is one of the most widely used classical spectral estimation methods. An example PSD estimated from the green channel of the image in Figure 1A using WOSA is shown in Figure 6. Square segments with an overlap percentage of 50 % and a side length equal to 1/3 of the image's shorter dimensions were used. Windowing was performed using a circularly symmetric Hann window. The leakage artefact present in the previous examples is absent in this case due to the Hann window. Spectral resolution appears to be comparable to Bartlett's method, but the variance is visibly lower. The decrease in variance is achieved without sacrifying resolution due to the larger number of segments for averaging in comparison to Bartlett's method, obtained by allowing overlap between neighboring segments.

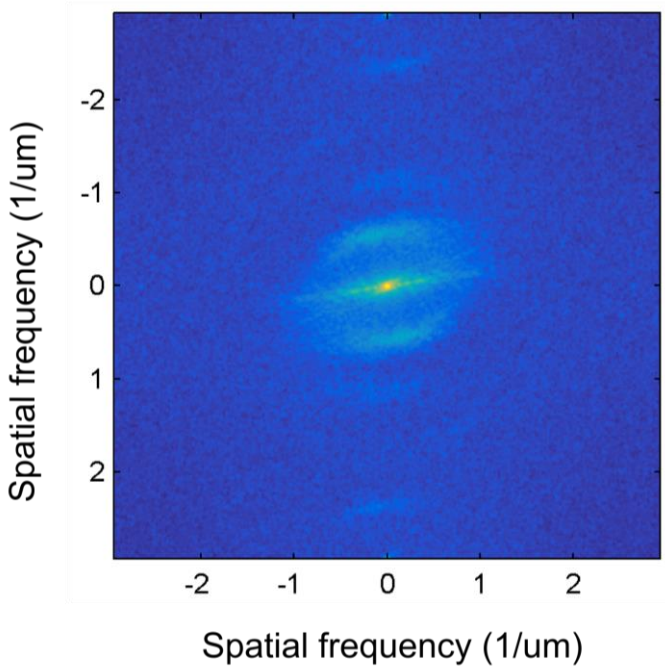

Figure 6. *WOSA estimate of the PSD of the green channel of Figure 1A. Square segments with an overlap percentage of 50 % and a side length equal to 1/3 of the image's shorter dimensions were used. A circularly symmetric Hann window was used to reduce leakage.*

In contrast to Bartlett's method, users of WOSA are not only required to select the segment size (or the number of segments) but also the type of windowing function and the amount of overlap between segments (Stoica & Moses 2005). The common rule of thumb is to use a Hann window with 50 % overlap (Marple 1987; Johnson & Long 1999; Stoica & Moses 2005; Wu *et al.* 2012). The one-dimensional Hann window is given by (Stoica & Moses 2005):

$$w_{Hann}(x) = 0.5 + 0.5\cos\frac{\pi x}{M} \tag{17}$$

for $x = 0, 1, 2 \ldots M\text{-}1$, where x is the sample index and M is the data segment length. A circularly symmetric 2D Hann window can be easily obtained by transforming the coordinates of an M x N array, whose origin is at the center of the array, to polar coordinates. Then one only has to evaluate Equation (17) at each element of the array, substituting x with the radial polar coordinate of the corresponding element. Empirical studies have shown that exact tuning of the overlap percentage is not crucial, as long as the percentage is over 50 % (Marple 1987; Jokinen *et al.* 2000). Overlap percentages below 50 % are not typically used due to the desire to balance the effects of windowing. Higher percentages offer little improvement in terms of decreasing variance because highly overlapping segments are also highly correlated with each other, decreasing the effectiveness of the averaging process. A higher overlap percentage thus offers little improvement in performance while increasing the required computation time due to a larg-

er number of DFTs. While using a Hann window with an overlap percentage of 50 % is acceptable in most cases, users of WOSA are still faced with the bias-variance dilemma when selecting the number of segments to average (or, correspondingly the segment length *M*). While averaging more segments reduces variance, in the practical case of a finite number of input samples it also reduces the spectral resolution as fewer samples are available for spectral estimation per segment, similarly to Bartlett's method. A suitable value for the number of segments thus ultimately depends on the application and the desired spectral resolution. For example, Wu and colleagues (Wu *et al.* 2012) tested several options for their study. They ultimately used segment sizes of 256 x 256 pixels (corresponding to ¼ of the full image dimensions of 1024 x 1024 pixels) and 512 x 512 pixels (corresponding to ½ of the image dimensions) as reasonable compromises.

2.3 Orientation analysis

The earliest reports on the development of computational methods for image-based orientation analysis originate from the early days of digital microscopic imaging (Chaudhuri *et al.* 1987). Almost all of the proposed methods are based either on spectral analysis in the frequency domain or different approaches utilizing pixel intensity gradients in the spatial domain (Boudaoud *et al.* 2014). Previously proposed frequency domain methods are presented in Section 2.3.1, while an overview of methods operating in the spatial domain is given in Section 2.3.2. Some approaches combine spatial and frequency domain analysis steps. Such methods, here referred to as hybrid methods, are discussed in Section 2.3.3. Existing software tools intended for image-based orientation analysis, namely OrientationJ and FibrilTool, are introduced in Section 2.3.4. Finally, some important concepts concerning circular statistics, a subfield of statistics dealing with circular data, are introduced in Section 2.3.5 due to their central importance for handling and analysis of data describing orientations or directions.

2.3.1 Frequency domain methods

In one of the earliest studies, Chaudhuri and colleagues (Chaudhuri *et al.* 1987) analyzed the orientation of collagen fibers in ligaments based on directional filtering in the spatial frequency domain. Their method is based on the linearity and rotation properties of the 2D Fourier transform and involves Fourier transforming the image, performing filtering with sector filters to separate different angle bands from the spectrum, and then inverse Fourier transforming to obtain component images containing only the collagen fibers whose orientation is within the corresponding angle band. The component images are then thresholded using Otsu's method and the areas covered by fibers of each orientation are used to form an orientation distribution. A modified version of the method was adopted by Petroll *et al.* and applied to the study of stress fiber orientation in corneal wounds (Petroll *et al.* 1993). In this study, a single dimensionless parameter termed 'orientation index' was used to quantify the orientation of the stress fibers relative to the

wounds in rabbit corneas. Spectral methods based on the Fourier transform were also successfully applied to the analysis of the morphology and orientation of vascular endothelial cells (Palmer & Bizios 1997) and to the study of smooth muscle cells subjected to magnetic fields (Umeno & Ueno 2003).

More recent research has focused on refining and characterizing the previously developed methods and validating their use in new applications. The use of spectral methods has been extensively discussed in the context of orientation analysis of different fibers (Marquez 2006; Ayres *et al.* 2008; Sander & Barocas 2009; Schriefl *et al.* 2012). Marquez (Marquez 2006) proposed a filtering scheme where instead of summing over the whole PSD, only a certain spatial frequency band is used to estimate the orientation distribution of the target fibers. Sander and Barocas (Sander & Barocas 2009) also suggested using different spatial frequency cutoff values depending on the application in order to increase the accuracy of the method and to decrease the effects of noise on the results. A related approach was adopted by Schriefl and colleagues (Schriefl *et al.* 2012) to obtain orientation information at different length scales. However, in practical applications dealing with biological rather than methodological questions, simpler versions of the spectral approach based on computing a single alignment index from the complete PSD (Pang *et al.* 2009; Khan *et al.* 2015) or on modeling the shape of the PSD with an ellipse (Lo *et al.* 2012) have still retained their popularity. Perhaps surprisingly, most of the published studies still rely on the raw periodogram as the spectral estimator, rather than using WOSA or some other more advanced method of spectral estimation. This approach has probably been adopted for its simplicity and the resulting lower number of tunable parameters.

2.3.2 Spatial domain methods

Approaches operating in the spatial domain, based on different forms of edge detection and intensity gradients, were also reported already in the late 90's by Karlon and colleagues (Karlon *et al.* 1998; Karlon *et al.* 1999). The Hough transform -based technique tested by Karlon and colleagues involves a large number of consequent processing steps. First, edges are highlighted by applying a difference of Gaussians operator on the image. Then, four directional images are formed by convolving the edge-enhanced image with Sobel gradient filters of four different orientations. The directional images are then thresholded and relevant regions are selected based on their shape. Finally, the Hough transform is used to determine the primary orientation of each region. The intensity gradient technique is based on filtering the image with horizontal and vertical Sobel filters to obtain horizontal and vertical intensity gradient values for each pixel. Weighted sums of the gradient orientations within square 20-pixel subimages are then computed to obtain dominant local orientations for different parts of the image. These techniques were applied to the analysis of myofibers in cardiac tissue of mice (Karlon *et al.* 1998) and endothelial cells subjected to fluid flow (Karlon *et al.* 1999).

Refinements on methods operating in the spatial domain have also been reported in recent publications (Rezakhaniha *et al.* 2012; Boudaoud *et al.* 2014). Most importantly, these methods have been implemented in the form of two openly available software plug-ins: OrientationJ (Rezakhaniha *et al.* 2012) and FibrilTool (Boudaoud *et al.* 2014). Both of them are based on structure tensors, which are matrix representations of partial derivatives of pixel intensity and are thus close relatives of the methods proposed earlier by Karlon and colleagues (Karlon *et al.* 1998; Karlon *et al.* 1999). OrientationJ has since been used to analyze cardiac myofibers (Pasqualini *et al.* 2015). OrientationJ and FibrilTool will be discussed in more detail in Section 2.3.4.

2.3.3 Hybrid methods

In addition to orientation analysis methods operating mainly in the spatial domain or the frequency domain, hybrid approaches featuring analysis steps in both domains have been reported. However, these methods are very application specific and have been adopted and fine-tuned by individual research laboratories to analyze particular types of cells and images. Researchers from the Disease Biophysics Group at Harvard University have used a method combining frequency domain preprocessing with spatial domain analysis when dealing with cardiac myocytes (Bray *et al.* 2010; Feinberg *et al.* 2012). The method is based on the periodic striated patterns visible in images captured of the myofibers of cardiomyocytes. Because the spacing of the protein structures which form the striated pattern is well known *a priori*, it is possible to utilize a bandpass filtering operation in the spatial frequency domain to highlight these structures in images obtained via inverse Fourier transforming after filtering. The cutoff frequencies of the bandpass filter are simply tuned so that the passband closely matches the spatial frequency of the striations. After the filtering operation, an algorithm originally developed for the processing of fingerprints is utilized to skeletonize the image. Intensity gradients are then used to compute an orientation distribution.

Images of nerve cells have also been analyzed using a hybrid approach (Kim *et al.* 2011). In contrast to the method described above, this approach features preprocessing in the spatial domain, followed by orientation analysis in the frequency domain. In short, the protocol consists of edge detection based on intensity gradients, followed by thresholding and morphological operations to form a binary image which contains the neuronal processes of interest. The binary image is then Fourier transformed and an orientation distribution is computed based on the whole power spectrum. Similarly to the method used by Bray, Feinberg and colleagues (Bray *et al.* 2010; Feinberg *et al.* 2012) in the context of cardiomyocytes, this method is also optimized for a particular cell type having a distinct morphology.

2.3.4 Existing software

Most studies utilizing image-based orientation analysis have relied on programs developed and probably also operated by image processing experts within the research group. Typically the software implementations used in these studies have not been shared with the research community. However, there are some exceptions in the form of publicly available software tools. We are only aware of two such tools: OrientationJ (Rezakhaniha *et al.* 2012) and the more recent FibrilTool (Boudaoud *et al.* 2014), which are both plug-ins for the general purpose image processing tool ImageJ (Schneider *et al.* 2012). Next, these tools are briefly introduced.

OrientationJ was originally introduced in order to analyze the local orientation of collagen fibers in confocal microscopy images (Rezakhaniha *et al.* 2012). It is based on structure tensors, which are commonly used in image processing (Jahne 2004). For each pixel of an image, the structure tensor is defined as a matrix containing the partial spatial derivatives of the image intensity along the horizontal and vertical directions. Based on the structure tensor, OrientationJ computes the local orientation, energy and coherency for each pixel. The local orientation within the area of interest corresponds to the largest eigenvector of the tensor while the energy and coherency are parameters indicative of the degree of isotropy or anisotropy within the region. The energy parameter is the trace of the tensor matrix while the coherency is the ratio between the difference and sum of the largest and smallest eigenvalue. These parameters can be used to exclude orientation information from areas which are unimportant by setting different thresholds for the energy and coherency values. Color-coded images may be used to help in this process but manual selection of the different parameter values is still a time-consuming task for the user, although the analysis is otherwise performed automatically. Moreover, functionality of the software has been rather poorly documented and there is no user guide available. A screenshot of the OrientationJ user interface is shown in Figure 7. The analysis parameters and the computed orientation distribution are shown in separate windows.

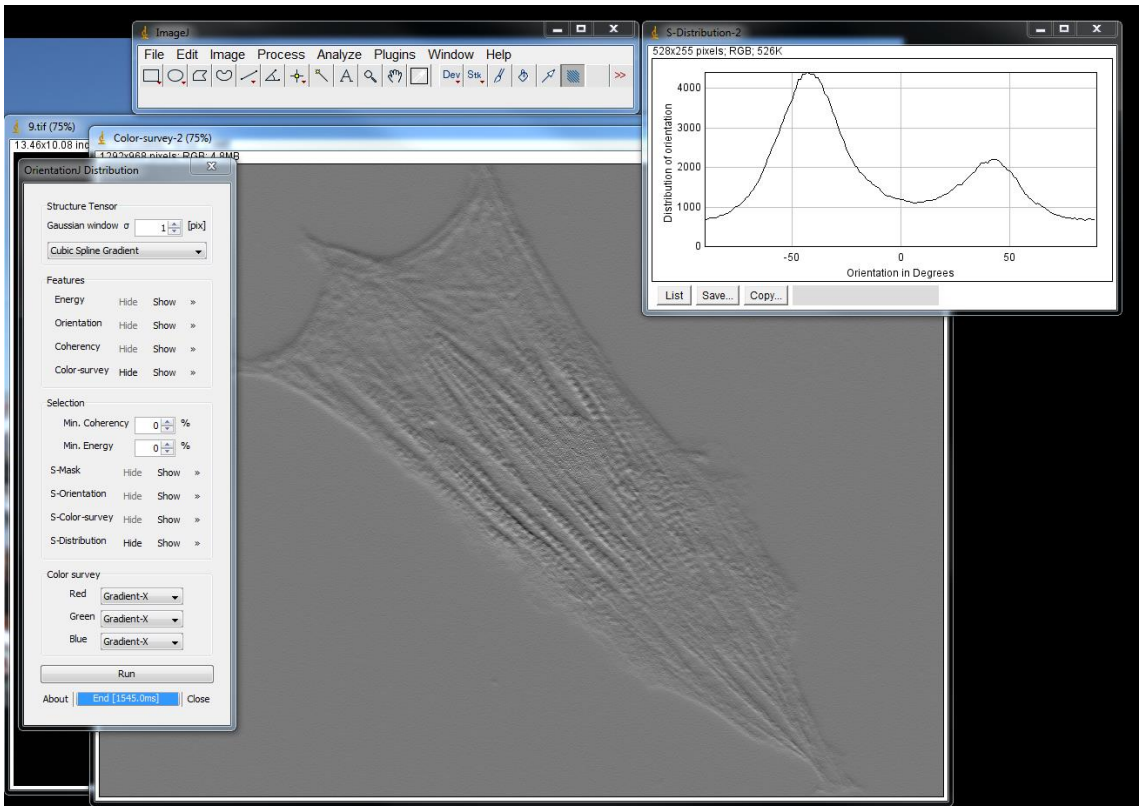

Figure 7. *A screenshot of the OrientationJ user interface. The analysis settings can be customized using the panel on the left-hand side. The estimated orientation distribution is shown in the upper-right corner.*

FibrilTool is functionally very similar to OrientationJ and is based on the concept of nematic tensors (Boudaoud *et al.* 2014). The nematic tensor is a matrix whose components are computed based on the first spatial partial derivatives of pixel intensity. The dominant orientation within a region of interest (ROI) can be calculated by averaging the nematic tensors of each pixel within the ROI and by computing the eigenvector corresponding to the largest eigenvalue of the averaged tensor. A value describing the degree of isotropy or anisotropy within the region can also be computed based on the eigenvalues of the averaged nematic tensor. In contrast to OrientationJ, FibrilTool features only a single tunable parameter and is extremely simple to use. In addition to a graphical representation of the dominant orientation and the degree of isotropy within each ROI, FibrilTool creates a plain text file containing the estimated values. Despite its simple user interface and operation, FibrilTool has the significant disadvantage of requiring manual selection of ROIs. This means that the user has to manually select each desired target from each image to perform the analysis. Since selecting and processing a single ROI typically takes approximately 20 seconds (Boudaoud *et al.* 2014), analyzing just a single image can take minutes of manual labor, essentially ruling out any high-throughput experiments. Moreover, FibrilTool has no means of excluding any of the features in the image, apart from the manual selection of ROIs. This means that the results are affected by noise and any interfering structures if they are present within the ROI. Nevertheless, FibrilTool can currently be considered the state-of-the-art freely

available orientation analysis tool. A screenshot of the FibrilTool user interface is shown in Figure 8. The manually selected region of interest is shown as a yellow box and the estimated mean orientation is indicated with a blue line.

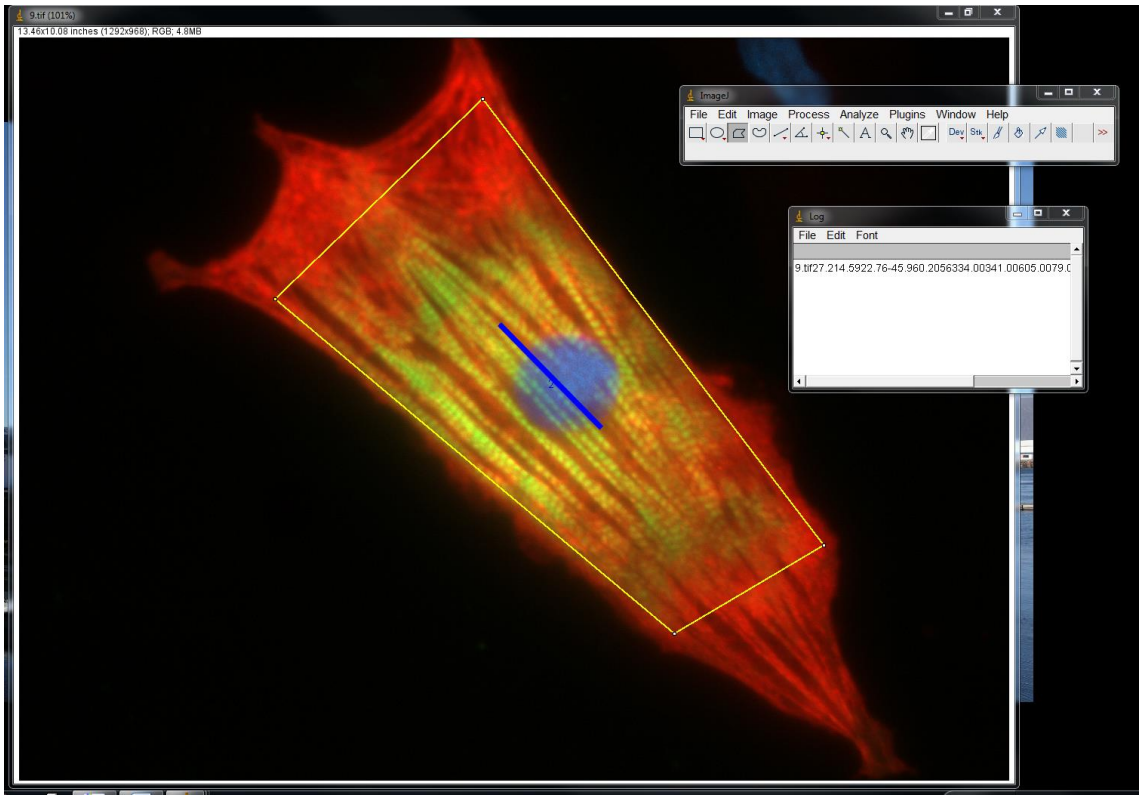

Figure 8. *A screenshot of the FibrilTool user interface. The image has been post-processed to make the blue line, indicative of the mean orientation, more visible. The region of interest selected by the user is shown as a yellow box.*

2.3.5 Circular statistics

Orientation analysis depends heavily on circular statistics, a subfield of statistics dealing with circular data, that is, data defined on an angular scale (Berens 2009; Mardia & Jupp 2009). An angular scale does not have a designated zero or any meaningful designations of low or high values. Orientations and directions are clearly angular variables. For example, it does not make any sense to ask if North is larger or smaller than South. Values on an angular scale also repeat themselves with a certain periodicity. As an example, if we consider the previous case given in numeric format as degrees running clockwise from North, then it is evident that 360° (North) or 720° (North) have the same meaning as 0° (North). In this sense, the angular scale is finite but endless. It should be noted that selecting North as the direction 0° is totally arbitrary and a matter of agreement. Axial data, such as orientations, are closely related to circular data like directions. The difference between the two is that in the case of axial data, opposite directions on a circle are considered equivalent. That is, an orientation of 0° (given ac-

cording to the same convention as used in the example above) is equivalent to an orientation of 180°. This orientation could also be called North-South. For example, a road built between a city in the North and another city in the South does not have a direction but it does have a North-South orientation whereas a car travelling on the road does have a direction. In addition to directions and orientations, many other types of data are defined on an angular scale. For example, the time of day is an angular variable with a periodicity of 24 hours.

Even though angular data are widespread, the field of circular statistics only started to attract interest during the 1950's (Berens 2009). This is surprising, as using statistical techniques developed for linear data often lead to seriously erroneous results if applied to angular data. Even a simple operation such as computing the mean of two values defined on an angular scale cannot be performed in the same way as it would be done for linear data. For example, let us consider two directions given according to the convention used in the previous examples: 45° (North-East) and 315° (North-West). If one simply computes the arithmetic mean of these two directions, neglecting their angular nature, the result would be 180° (South). However, it is clear that the true mean direction is North, or 0°, and the simple arithmetic mean is not only erroneous, but it is actually completely opposite to the correct value. The correct way to perform this calculation would be to transform the directions to unit vectors and average them to obtain the mean resultant vector. As this simple example illustrates, angular variables have to be handled in a way that is profoundly different from linear variables, and for this reason circular analogues have been developed for many operations and techniques originally developed for linear data. As many of these concepts are not generally well known, the ones relevant to this study are next briefly introduced (Berens 2009).

Transforming axial data to circular data: a sample of axial data can be converted to a unimodal, circular sample by a procedure called angle doubling. That is, a sample of axial values α_i is converted according to the following mapping:

$$\alpha_i \rightarrow 2\alpha_i \ (mod\ 2\pi) \tag{18}$$

This is necessary for some computations such as calculation of the mean direction of a sample of orientations. After the computations, the result can be transformed back to the original scale.

Mean: the mean of a sample of circular data having *n* elements is obtained by first transforming the *n* directions into unit vectors in the two-dimensional plane by:

$$r_i = \begin{pmatrix} \cos \alpha_i \\ \sin \alpha_i \end{pmatrix} \tag{19}$$

Next, the *n* unit vectors r_i are averaged by:

$$\bar{r} = \frac{1}{n}\sum_{i} r_i \tag{20}$$

to obtain the mean resultant vector $\bar{r}$. The mean direction $\bar{\alpha}$ is then obtained by transforming $\bar{r}$ with the four quadrant inverse tangent function (*i.e.*, the inverse tangent function which takes into account the quadrant to yield angles from 0 to 360 degrees). In the case of axial data, the mean direction is usually divided by two to obtain the mean orientation in the original axial scale.

Resultant vector length: the length of the mean resultant vector is used to calculate descriptive statistics measuring circular spread. A longer mean resultant vector length corresponds to a sample that is concentrated more strongly around the mean direction. The mean resultant vector length *R* is computed by:

$$R = \|\bar{r}\| \tag{21}$$

In some situations the sample of circular data may have been obtained in binned form. That is, the distribution of circular values is in the form of a histogram with a given bin size. This is exactly the case when using spectral analysis for studying orientation distributions, because the distribution of power at different orientations is quantified in discrete bins of finite size. Use of binning introduces bias in the estimate of *R*, which can be taken into account by calculating a correction factor *c* based on the bin spacing *d*. The correction factor is computed by:

$$c = \frac{d}{2\sin(d/2)} \tag{22}$$

and the corrected estimate of mean resultant vector length is $R_c = cR$.

Circular variance: circular variance indicates the amount of angular spread. It is computed by:

$$S = 1 - R \tag{23}$$

Circular variance is bounded in the interval [0, 1], having a value close to 1 for data that is distributed evenly around the circle and a value close to 0 for data that is concentrated towards the mean direction.

Angular standard deviation: angular standard deviation is also a measure of angular spread, analogous to the linear standard deviation. It is computed by:

$$s = \sqrt{2(1 - R)} \tag{24}$$

and is bounded in the interval [0, √2].

Circular skewness: circular skewness is a measure of the symmetry of the data around the mean direction. It is computed by:

$$b = \frac{1}{n}\sum_{i=1}^{n} \sin 2(\alpha_i - \bar{\alpha}) \tag{25}$$

Values close to 0 are indicative of a population that is symmetric around the mean direction.

Circular kurtosis: circular kurtosis is a measure of the 'peakedness' of the data. It is computed by:

$$k = \frac{1}{n}\sum_{i=1}^{n} \cos 2(\alpha_i - \bar{\alpha}) \tag{26}$$

A large positive value close to 1 indicates a distribution that is strongly peaked rather than smooth.

The von Mises distribution: the most commonly encountered circular distribution is the von Mises distribution *VM* (μ,κ), which is a circular analogue of the normal distribution. The parameters μ and κ are related to the mean direction and the amount of angular spread of the distribution, respectively. The probability density function of the von Mises distribution is:

$$p_{VM}(\theta;\mu,\kappa) = \frac{1}{2\pi I_0(\kappa)} e^{\kappa \cos(\theta-\mu)} \tag{27}$$

where I_0 is the modified Bessel function of order zero.

As an example, let us consider two sets of axial orientation values: set 1 ({60°, 70°, 80°, 85°}) and set 2 ({15°, 50°, 110°, 160°}). After the axial to circular transformation (*i.e.*, angle doubling), the values are: set 1 ({120°, 140°, 160°, 170°}) and set 2 ({30°, 100°, 220°, 320°}). The mean orientation of set 1 (after division by two) then is 73.802°. For set 2, the mean orientation is 8.023°. The mean resultant vector lengths are 0.945 and 0.180 for sets 1 and 2, respectively, corresponding to circular variances of 0.055 and 0.820. This indicates that the orientations of set 2 are much less concentrated around the mean orientation than those of set 1. In other words, set 1 represents a distribution exhibiting a higher degree of anisotropy than set 2. The same observation can be made on the basis of angular standard deviation values of 0.111 and 1.640 for sets 1 and 2, respectively. The circular kurtosis of set 1 is 0.788, indicative of a peaked distribution when compared to set 2, which has a circular kurtosis of 0.050. The circular skewness of set 1 is close to zero with a value of 0.010, indicating that the orientations are distributed symmetrically around the mean orientation of 73.802°. This seems reasonable,

when considering the original values of 60°, 70°, 80° and 85°. For set 2, the circular skewness is slightly higher at 0.123, indicating a slightly skewed distribution of orientations.

3. SOFTWARE IMPLEMENTATION

The functionality and detailed implementation of the software developed in this study are presented in this chapter. The contents of the chapter follow the workflow of the software. First, a general description of the implementation is given in Section 3.1, followed by explanations of the functionality for reading images, configuring analysis settings and importing segmentation data in Sections 3.2, 3.3 and 3.4, respectively. The detailed methodology used for spectral estimation and separation of the spectrum into different components is then described in Sections 3.5 and 3.6, respectively. Finally, the analysis and storage of results are described in Sections 3.7 and 3.8.

3.1 General description of the software

The software, named CytoSpectre, was developed using MATLAB R2014a and it can be used either under MATLAB or as a standalone 64-bit Windows 7 application. If run under MATLAB, CytoSpectre can be operated either via the command line or via a graphical user interface. The main analysis steps performed for each image are illustrated in Figure 9 and described in the following sections. Briefly, the images are first read into the software together with optional mask images containing segmentation results. The segmentation masks can be used to perform the spectral analysis individually for each cell, if the image contains multiple cells. In that case, basic descriptors of cellular morphology are also calculated. Spectral estimation is then performed using the WOSA estimator (see Section 2.2.4). CytoSpectre separates the spectrum into two components, which are defined in more detail in Section 3.6. The mixed component corresponds to the complete power spectrum limited only by lower and higher cutoff frequencies or corresponding wavelengths, which may be specified by the user. The mixed component thus represents the complete image and is analyzed similarly to the previously published frequency domain approaches discussed in Section 2.3.1. The other component, referred to as the detail component, represents a distinct spectral region limited to a particular range of orientations or spatial frequencies, reflecting contributions from some particular structures of interest in the image. An example of such structures would be the sarcomere subunits of myofibrils, seen in Figure 1. The orientation and wavelength distributions of each component are then separately analyzed, usually followed by exporting the results for further study. A screenshot of the CytoSpectre GUI during a typical analysis session is shown in Figure 10. The GUI features controls for importing and examining images, customizing analysis settings, running the analysis and for exporting results. Analysis results are also presented in graphical form as plots.

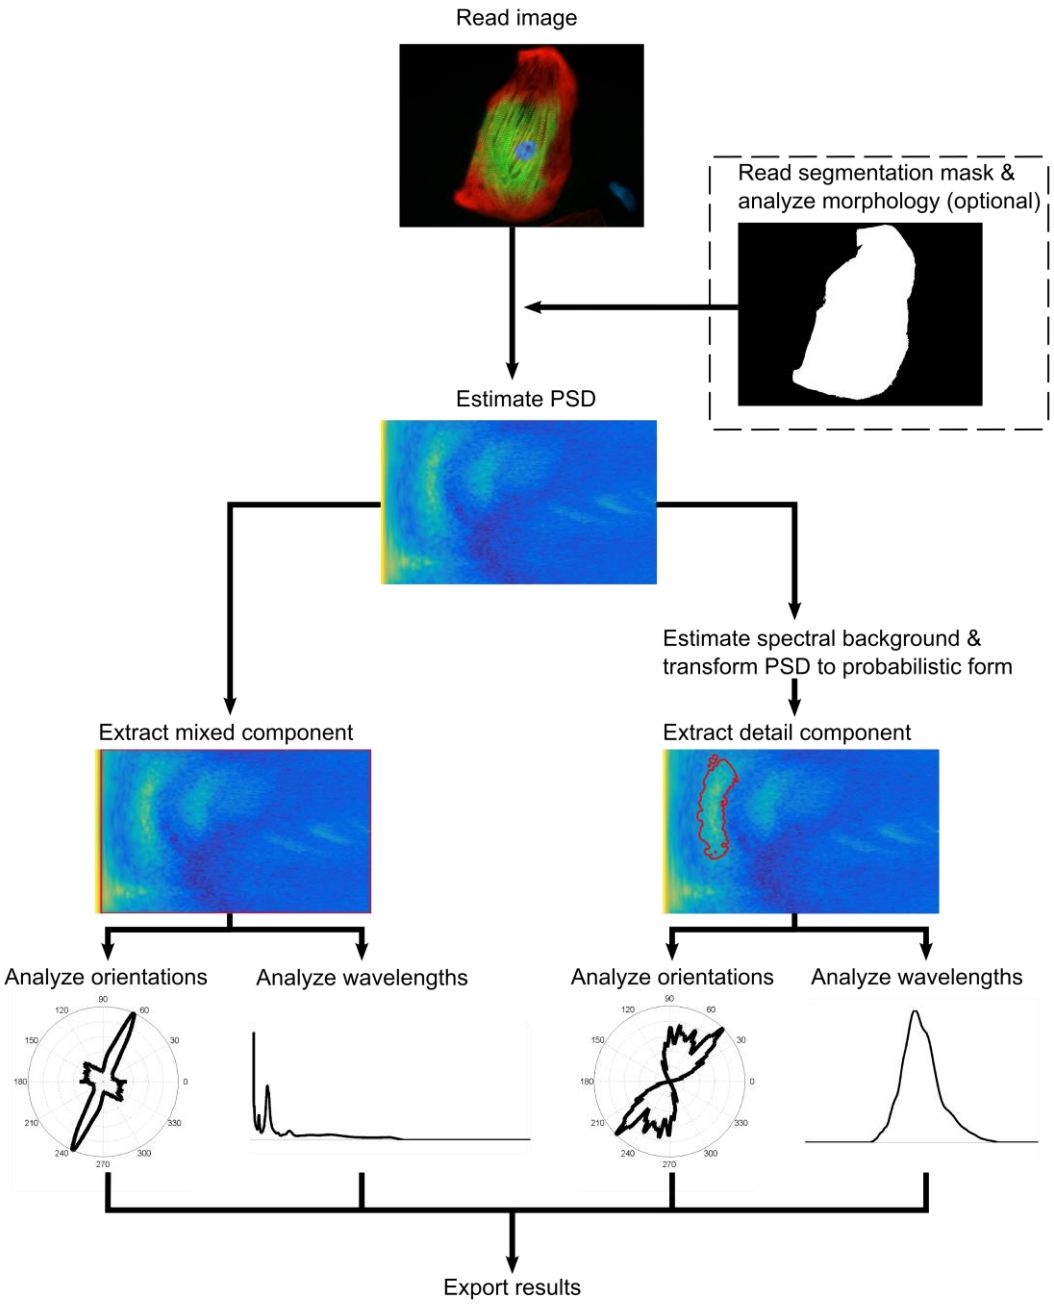

Figure 9. *Analysis workflow. After reading the image, an optional step may be performed to import segmentation results obtained using external software, allowing the rest of the analysis to be performed cell-by-cell. Spectral estimation is then performed and the PSD is separated into two components (red outlines). The mixed component corresponds to the conventional type of analysis where the entire PSD (excluding regions outside cutoff limits) is considered. The detail component, extracted using an estimate of spectral background and an iterative procedure, represents a distinct spectral region corresponding to some features of interest in the image. The orientation and wavelength distributions of each component are then analyzed and the results are exported for further analysis.*

3.2 Reading and handling images

CytoSpectre utilizes MATLAB's *imread* function to read images and is compatible with grayscale and RGB images in .tif, .jpeg, .gif., .bmp and .png formats. In principle, any 2D images are supported, although the software is intended for biomedical micrographs. CytoSpectre neither performs nor requires any particular preprocessing steps for the images, decreasing the risk of bias introduced by such modifications (Boudaoud *et al.* 2014). CytoSpectre uses temporary files written on the hard disk to store settings and results for images during an analysis session. By using the temporary files, only a single image has to be loaded into memory at a time, allowing users to examine large sets of images without particular system requirements in terms of available memory. The temporary files are saved in .MAT format and automatically deleted when no longer needed. Due to the potentially large amount of disk space required for the temporary files, the user may select a preferred location on the disk for storing them. By default, the temporary files are saved to the same directory as the corresponding images.

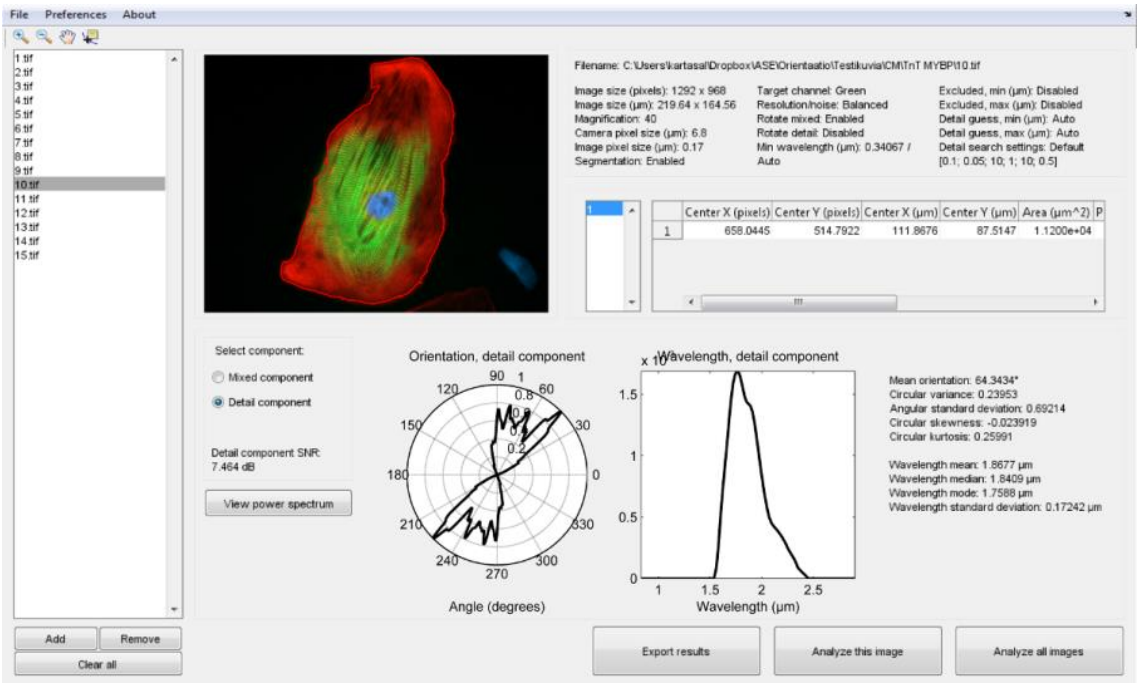

Figure 10. *The GUI includes controls for importing and selecting images (far left), examining the current image (top-left), reviewing analysis settings, image information and morphological data (top-right), selecting spectral components (middle-left), examining orientation and wavelength distributions (center), viewing summary statistics (middle-right) and for running the analysis and saving results (bottom-right).*

3.3 Configuring analysis settings

A number of settings may be configured by the user depending on personal preferences and application dependent requirements. General settings include the preferred location

for saving temporary files and the preferred length unit, the default being µm. The most important actual analysis settings include:

- **Magnification O:** Magnification of the images can be set as fixed or detected automatically from the filename of each image. Filename based detection requires the name to contain a substring which includes the magnification value (for example, 40X).
- **Camera pixel size Δ_C:** The physical size of the camera pixels.
- **Import cell segmentation:** If cell segmentation (*i.e.*, partition of the image into cells and background) is available, it can be imported into CytoSpectre, allowing cell-by-cell analysis. See Section 3.4 for details.
- **Target channel:** The color channel to analyze. This setting has no effect for grayscale images.
- **Rotate orientation by 90 degrees:** As discussed in Section 2.2.3, the orientation of spectral features is usually orthogonal to the orientation of the structures in the image and the orientations calculated from the spectrum thus have to be rotated by 90 degrees. In some cases, intensity variation occurs parallel to the orientation of the structures of interest. In such cases, the rotation correction should be disabled. The correction can be applied separately to the mixed and detail components and is enabled by default.
- **Minimum wavelength:** By default, the minimum wavelength analyzed is limited only by the resolution of the imaging system. A manual minimum wavelength limit can be set to exclude structures whose characteristic dimensions are smaller than this limit. This can be useful, for example, in order to reduce the effects of interfering low wavelength patterns such as grooves on a cell culture substrate or artefacts produced by the camera sensor. See Section 3.6.3 for details.
- **Maximum wavelength:** By default, the maximum wavelength analyzed is set automatically to only remove long wavelength artefacts caused by the spectral estimation process. A manual maximum wavelength limit can be set to exclude structures whose characteristic dimensions are larger than this limit. This can be useful, for example, if the user is only interested in finer details of the image. See Section 3.6.3 for details.
- **Exclude wavelength band:** By default, the whole spectrum between the minimum and maximum wavelengths is analyzed. A given band of wavelengths can also be excluded from the analysis by enabling this setting. This can be useful, for example, if there are interfering patterns or artefacts whose characteristic dimensions are known to be within a certain range of values.
- **Detail component wavelength band:** By default, the detail component is extracted automatically based on signal-to-noise ratios. Semi-automatic detection can be used if the characteristic dimensions of the structures of interest are known or a reasonable guess is available. In this case, the detail component is

automatically extracted based on the given range of expected wavelengths. See Section 3.6 for details.

Based on the values input by the user, the physical size of the image pixels is obtained as:

$$\Delta_I = \frac{\Delta_C}{0} \tag{28}$$

and the Nyquist frequency is then given by:

$$f_{Nyquist} = \frac{1}{2\Delta_I} \tag{29}$$

in accordance with the sampling theorem (Gonzalez & Woods 2001).

A number of other parameter values, explained in the remaining sections of this chapter, may also be adjusted, but based on a sensitivity analysis using simulated images the results produced by the software are not overly sensitive to the exact values of these settings (see Section 5.6). Therefore, the default values of these parameters are usually appropriate. These settings include the following parameters:

- **Spectral resolution W**
- **Segment length L**
- **Significance level A**
- **Angular tolerance ε**
- **Convergence limit τ**
- **Maximum number of iterations N_{max}**
- **Prior width E**

Once suitable settings have been found, CytoSpectre allows them to be stored in a user profile file in .MAT format, which can then be loaded during later analysis sessions. This feature also makes it easy to store favorite settings suitable for different types of images.

3.4 Importing cell segmentation data

In addition to analyzing complete images, the spectral analysis can be performed on a cell-by-cell basis, if the image contains multiple cells. Cell-by-cell analysis can be accomplished by providing the software with cell segmentation results in the form of binary mask images, along with the actual images. The mask images are required to have zero intensity for background pixels and positive values for pixels belonging to cells. If the user has enabled cell segmentation, the mask images are automatically read during analysis, provided that certain naming rules are followed for the mask images. In addi-

tion to being able to perform the rest of the analysis separately for each cell, basic descriptors of cellular morphology, such as cell area, perimeter and eccentricity are also computed. Cellular segmentation is a common problem in biomedical image analysis, and a large number of image processing and analysis tools, for example ImageJ (Schneider *et al.* 2012) and CellProfiler (Kamentsky *et al.* 2011), are available for performing the actual segmentation step. In some cases, the segmentation masks may even be formed manually. Since specialized tools fine-tuned for a particular application are often required for successful segmentation, none of the available tools were included as an integral part of CytoSpectre. Instead, users can freely perform the segmentation step using tools they have found suitable for their application and import the final results into CytoSpectre as mask images. However, the segmentation step is optional and in many cases unnecessary, especially if only a single cell is present in the image.

If cell segmentation is enabled, a number of parameters summarizing the morphological features of each cell are calculated using the *regionprops* function of MATLAB. These parameters are not used by the software in any subsequent analysis steps, but such morphological statistics may provide useful additional information for the user. If segmentation is disabled, this step of the analysis is skipped. The parameters include:

- **Center X:** The x-coordinate of the cell's centroid.
- **Center Y:** The y-coordinate of the cell's centroid.
- **Area:** The physical area covered by the cell.
- **Perimeter:** The physical perimeter of the cell, that is, the distance along the outline of the cell.
- **Major axis:** The length of the major axis of an ellipse that has the same normalized second central moments as the cell region, that is, the 'length' of the cell along its longitudinal axis.
- **Minor axis:** The length of the minor axis of an ellipse that has the same normalized second central moments as the cell region, that is, the 'width' of the cell along its transverse axis.
- **Eccentricity:** The ratio of the distance between the foci of an ellipse that has the same normalized second central moments as the cell region and the major axis length of the ellipse. Eccentricity values lie in the interval [0, 1] with a value of zero corresponding to a circle and a value of one corresponding to a line segment. In other words, cells that have an elongated shape have a higher eccentricity than cells whose shape is round.
- **Solidity:** The ratio of cell area and the area of the cell's convex hull. Solidity values lie in the interval (0, 1] with smoothly shaped cells having a high value and cells with protrusions having a low value.

Centroid coordinates, area, perimeter, major axis length and minor axis length are converted from pixels to physical units in accordance with the image pixel size Δ_I.

3.5 Spectral estimation

Spectral estimation is performed with Welch's method (Welch 1967; Marple 1987), or WOSA, described in Section 2.2.4. If cell segmentation is enabled, background pixels are set to zero prior to spectral estimation to eliminate contributions from all other image features except the cell being analyzed. Otherwise, the whole image is analyzed. Mean is first subtracted from the image and the image is then divided into N_S square segments with 50 % overlap between adjacent segments. The size of the segments M_W, controlling the tradeoff between spectral resolution and stability, can be selected by the user by adjusting the spectral resolution parameter W. Square segments are used (*i.e.*, $M_W = N_W$) and the length of the segment side is given by:

$$M_W = N_W = W\ min\{N, M\} \tag{30}$$

where N and M are the dimensions of the image. The segment size is thus selected as a fraction of the shorter image dimension, as dictated by the value of W. The options include maximum resolution (W = 1), high resolution (W = 1/2), balanced (W = 1/3), low noise (W = 1/4) and minimum noise (W = 1/5). Each segment is multiplied by a circularly symmetric Hann window (see Equations (15) and (17)) to reduce spectral leakage and the DFT of the segments is then computed with the FFT algorithm (implemented by MATLAB's *fft2* function). Before computing the DFT, the segment is zero-padded either to a size equaling the next power of two larger than the original segment size or to a minimum size of 1024 x 1024, whichever is larger. The primary reason for this operation is to obtain optimal FFT performance while the secondary reason is to slightly interpolate the PSD to aid in subsequent visual examination (see Section 2.2.3 for details). The final spectral estimate is formed as the arithmetic mean of the power spectral densities of all segments according to Equations (15) and (16). The conventional shifting operation (see Section 2.2.3) is then performed to place the lowest frequencies at the center of the spectrum.

The power spectrum is transformed into polar coordinates according to Equations (9), (10) and (11) for easier processing during later analysis steps. The transform is performed using bilinear interpolation. The number of spectral elements in the interpolated polar power spectrum is selected so as to retain or exceed the resolution of the original Cartesian form of the spectrum along both dimensions at all spatial frequencies below the Nyquist frequency $f_{Nyquist}$. The numbers of bins required along the two dimensions are calculated based on the number of spatial frequency bins N_{uv} in the Cartesian representation. Neither the zero spatial frequency bin nor frequencies above $f_{Nyquist}$ are included in the power spectrum after the polar coordinate transform. Moreover, only angles from 0 to 180 degrees (*i.e.*, the upper half circle in the Cartesian power spectrum) are retained, since the lower and upper half circles of the Cartesian power spectrum are symmetric (see Section 2.2.3 for details). For the angular dimension, the required number of angle bins in the polar power spectrum is calculated based on the circumference

of a circle whose radius corresponds to the number of spatial frequency bins between zero frequency and $f_{Nyquist}$ as:

$$N_\theta = \left\lceil \pi \frac{N_{uv}}{2} \right\rceil \tag{31}$$

For the spatial frequency dimension, the required number of frequency bins in the polar power spectrum is:

$$N_f = \left\lfloor \frac{N_{uv}}{2} \right\rfloor \tag{32}$$

The PSD estimate expressed in polar coordinates, $P_{WOSA}(f, \theta)$, given as a function of spatial frequency f and orientation θ, is used for the remaining analysis steps.

3.6 Separation of spectral components

CytoSpectre separates two components from the PSD, which are hereafter referred to as the mixed component and the detail component. The mixed component represents all features of the image present in the range of spatial frequencies dictated by upper and lower cutoff frequencies. By default, these frequency cutoffs are employed merely to exclude extremely low spatial frequencies, which typically only contain uninteresting, poorly resolved features (Chaudhuri *et al.* 1987; Petroll *et al.* 1993) or windowing artefacts (Wu *et al.* 2012). However, the cutoff values may also be adjusted by specifying corresponding minimum and maximum wavelength cutoffs (see Section 3.3). This may be desirable, for example, to exclude high-frequency noise or to focus the analysis on a particular wavelength range of interest. The detail component, on the other hand, represents some particular group of structures with a limited range of spatial frequencies and/or orientations, such as intracellular fibrils of given size. If present in the image, such structures are reflected in the power spectrum as more or less contiguous regions of high intensity relative to background. The aim of the detail component extraction procedure is to separate these regions from the rest of the power spectrum, allowing the user to obtain results that more specifically represent the structures of interest. CytoSpectre uses an iterative method to separate the spectral region corresponding to the detail component. The procedure consists of first forming an estimate of a background PSD, which is used to obtain a probabilistic representation of the PSD. An initial estimate of the detail component spectral region is then formed using the probabilistic PSD and a 2D Gaussian 'mask' in a manner inspired by a Bayesian classifier. Finally, the obtained region is iteratively adapted to the actual shape of the spectral area of interest. Moreover, if an estimate of the size of the targets of interest is available *a priori*, this information can be provided as a wavelength guess to guide the algorithm (see Section 3.3).

3.6.1 Estimation of spectral background

An estimate of spectral background is formed based on the polar power spectrum estimate to allow detection of spectral signals based on their intensity relative to the background. The background spectrum estimate represents image features, such as noise, which are not of interest when analyzing the structures present in the image. Spectral signals representing meaningful details and structures of interest in the image appear as peaks and local fluctuations on top of the smoother background spectrum. It has been widely reported, that the rotationally averaged power spectra of natural images tend to fall with increasing spatial frequency according to a power law (Field 1987; Tolhurst *et al.* 1992; Ruderman 1997; Field & Brady 1997; Hsiao & Millane 2005). In log-log space this corresponds to a decreasing approximately linear trend. That is, the low spatial frequency region of the power spectrum tends to dominate the high spatial frequency region. An estimate of the general shape of the spectral background is needed in order to differentiate signals from the background.

The background spectrum is estimated based on the polar power spectrum using averaging over all orientations and a smoothing procedure. The advantage of not using a parametric model for the background spectrum is greater flexibility in terms of applications and imaging systems. The downside is that since the background spectrum estimate is based on an estimate of the power spectrum of the actual image being analyzed, it is biased by spectral peaks representing signals of interest (Perron *et al.* 2008). The bias is reduced by smoothing to remove most of the spectral peaks and to retain only the general shape of the spectral background. First, the polar power spectrum is averaged over all orientations to obtain a one-dimensional spectrum:

$$P_{WOSA}(f) = \frac{1}{N_\theta} \sum_{i=1}^{N_\theta} P_{WOSA}(f, \theta_i) \tag{33}$$

where θ_i is the ith orientation bin. The one-dimensional spectrum is then transformed to log-log scale and linear interpolation, retaining the original number of data points, is performed to obtain equidistant data points along the spatial frequency axis. Smoothing via robust local regression using weighted linear least squares and a first degree polynomial model (robust lowess) is then performed. The smoothing procedure is implemented using MATLAB's *smooth* function. The method decreases the effect of outliers (*i.e.*, elements probably belonging to spectral peaks originating from genuine signals) by assigning zero weight to data outside six mean absolute deviations. By default, the span value associated with the smoothing method is 10 % of N_f, the length of the one-dimensional spectrum (*i.e.*, by default L = 0.10). The value of the span parameter L can also be adjusted by the user, but the analysis is not sensitive to the exact value (see Section 5.6). Smoothing is performed in log-log space to effectively vary the span depending on spatial frequency, allowing high frequency signal peaks to be suppressed while

retaining the general spectral shape, especially the typically narrow high intensity peak near zero frequency, at low frequencies. This effect is obtained because the log-log transform followed by linear interpolation effectively corresponds to upsampling at low frequencies and downsampling at high frequencies. Finally, the smoothed one-dimensional spectrum is converted back to linear scale and tiled to form a two-dimensional polar background spectrum $B(f, \theta)$ whose size is identical to the original polar power spectrum. The background spectrum formed this way is thus isotropic and does not introduce any orientational bias during later analysis steps.

Examples of a PSD estimate and the corresponding spectral background estimate are shown in Figure 11. The background estimate was computed for the power spectrum depicted earlier in Figure 1. The isotropic nature of the background estimate is clearly visible, as each row is simply an identical copy of the one-dimensional smoothed background estimate. The smoothing procedure has retained the general spectral shape, showing the typical decrease in the spectral values with increasing spatial frequency. At the same time, the distinct spectral components visible in the original PSD estimate have been successfully eliminated during the background estimation process.

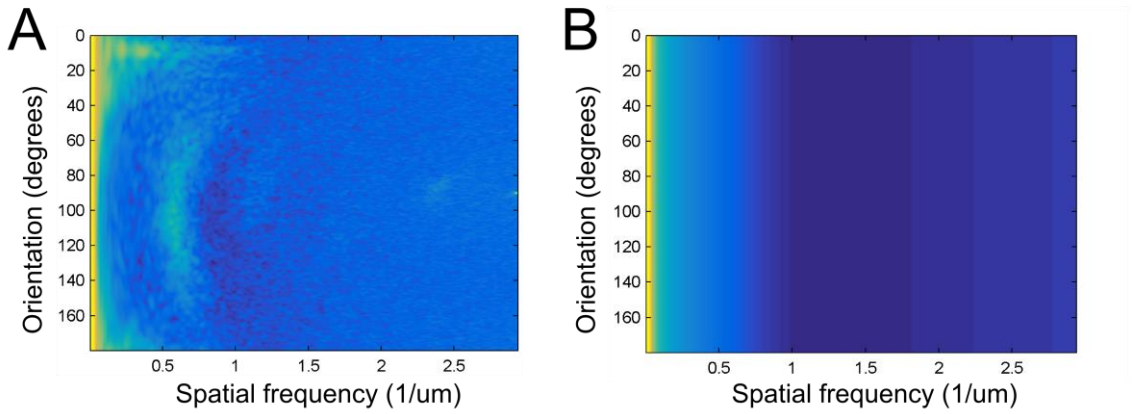

Figure 11. *The PSD estimate from Figure 1 (A) and the corresponding spectral background estimate (B), expressed in polar coordinates. The spectral values have been log_{10}-transformed for higher contrast.*

3.6.2 Expressing the power spectrum in probabilistic form

Based on the estimated background spectrum, the values of the polar power spectrum are converted to probabilities. After the conversion, each element $P_{WOSA}(f, \theta)$ of the spectrum is assigned a value representing the probability that the element belongs to a signal. In this context, 'a signal' should be understood loosely as some feature of the image represented by power spectral values larger than the corresponding values of the estimated spectral background. That is, elements of the power spectrum having a high magnitude relative to the corresponding element of the background spectrum should be assigned a high probability value, while spectral elements with values close to the back-

ground should be assigned low probabilities. To convert the values of the polar power spectrum to probabilities, a normalized WOSA spectrum is first utilized, defined as (Ferraioli *et al.* 2011):

$$Y(f,\theta) = \frac{P_{WOSA}(f,\theta)}{B(f,\theta)} \tag{34}$$

where $B(f,\theta)$ is the background spectrum estimate. In the assumption of vanishing correlations caused by the windowing and overlapping process, the normalized WOSA spectrum is gamma distributed with $\sigma_\gamma = 1/N_S$ and $h_\gamma = N_S$ (Ferraioli *et al.* 2011; Johnson & Long 1999). The probability $Z(f,\theta)$ that the element $P_{WOSA}(f,\theta)$ represents a signal is then obtained as the value of the gamma cumulative density function evaluated at the corresponding value of $Y(f,\theta)$:

$$Z(f,\theta) = \frac{1}{\sigma_\gamma^{h_\gamma}\Gamma(h_\gamma)} \int_0^{Y(f,\theta)} t^{h_\gamma - 1} e^{\frac{-t}{\sigma_\gamma}} dt \tag{35}$$

where $\Gamma(\cdot)$ is the Gamma function.

3.6.3 Selecting spatial frequency cutoffs

Some of the lowest spatial frequencies are usually excluded from analysis because the definition of angle is very coarse near the zero spatial frequency, that is, close to the origin of the Cartesian power spectrum (Chaudhuri *et al.* 1987) and the information present at the lowest frequencies is typically not of interest (Petroll *et al.* 1993). Even if the number of angle bins in the polar power spectrum does not depend on spatial frequency, the values of the spectral elements close to zero spatial frequency have nevertheless been interpolated from a very small number of original data points in Cartesian space. Another reason for excluding some of the lowest spatial frequency bins are artefacts caused by the windowing process (Wu *et al.* 2012). The exact cutoff frequency has typically been selected based on experimentation with different values (Chaudhuri *et al.* 1987; Petroll *et al.* 1993; Marquez 2006; Sander & Barocas 2009; Wu *et al.* 2012). In CytoSpectre, the lowest spatial frequencies are by default excluded based on a fixed minimum angular resolution requirement of five degrees. Based on this criterion and the Cartesian representation of the power spectrum, a suitable lower spatial frequency cutoff f_L can be selected by excluding a circular element of the power spectrum, centered on the zero frequency bin. The requirement for the number of bins along the circumference of this circle to be such that the angular resolution criterion is fulfilled can be expressed as:

$$2\pi N_L = \frac{360°}{5°} \tag{36}$$

where N_L is the number of bins along the radius of the excluded circular region. Thus, the corresponding number of spatial frequency bins to exclude, starting from the origin or zero frequency, can be obtained as:

$$N_L = \left\lceil \frac{360°}{5°} \frac{1}{2\pi} \right\rceil = \left\lceil \frac{36}{\pi} \right\rceil = 12 \tag{37}$$

The actual spatial frequency value corresponding to the 12th bin then depends on the physical dimensions of the image. Provided that random noise is not an issue, specifying a value for the upper cutoff frequency f_H is somewhat arbitrary. By default, the highest spatial frequency included in the analysis in CytoSpectre is limited only by the Nyquist frequency $f_{Nyquist}$. Both of the cutoff frequencies can also be freely modified by the user by selecting corresponding cutoff wavelengths of choice (see Section 3.3). This may be desirable, for example, to exclude high-frequency noise or to focus the analysis on a particular wavelength range of interest.

3.6.4 Extraction of the mixed component

The mixed component represents all features of the image present in the range of spatial frequencies dictated by the upper and lower cutoff frequencies. To extract the parts of the spectrum associated with the mixed component, all spectral elements outside the cutoff frequencies are set to zero. The mixed component can thus be expressed as:

$$P_M(f,\theta) = \begin{cases} P_{WOSA}(f,\theta), & f_L \leq f \leq f_H \\ 0, & \text{otherwise} \end{cases} \tag{38}$$

3.6.5 Estimation of initial values for detail component extraction

CytoSpectre uses an iterative method to separate the spectral region corresponding to the detail component. Before the iterative procedure, a number of initial parameter values are estimated. First, the dominant orientation of the low frequency range of the power spectrum is estimated using a sliding window approach. Mean orientation is first computed for each spatial frequency bin of the polar power spectrum using the corresponding function of the CircStat toolbox (Berens 2009), excluding bins below and above the lower and higher cutoff frequencies, respectively. Angle doubling is performed to account for the axial nature of the data and the mean orientation angle is afterwards divided by two (Berens 2009). The mean orientation values are then transformed to absolute differences relative to the mean orientation of the first bin to avoid problems near zero or 180 degrees.

A sliding window, whose length is controlled by the parameter L and is equal to 10 % of N_f, the total number of spatial frequency bins by default (*i.e.*, L = 0.10), is then ap-

plied to detect the spatial frequency at which a transition in the mean direction takes place. The window is incrementally moved by one bin at a time, starting from the low frequency end, towards higher frequencies, and a Mann-Whitney test is performed at each location to test if the mean orientations within the bins overlapping the window differ significantly from all previous bins. The Mann-Whitney test is a non-parametric statistical test, which can be used to test the null hypothesis that two samples of data come from the same population. By default, a significance level A of 0.05 is used to detect the transition. The window length parameter L and the significance level A can also be adjusted by the user, but the analysis is not sensitive to the exact values of these parameters (see Section 5.6). Mean orientation $\bar{\theta}$ within the range of spatial frequencies defined by the lower cutoff frequency and the detected transition frequency is then calculated as described above. If no transition is detected, the whole spatial frequency range between the cutoff frequencies is used for calculating the mean orientation. The obtained value represents the main orientation at low frequencies, which is mostly dominated by gross properties of the image, such as the overall orientation of a complete cell imaged at high magnification.

The spatial frequency range of the detail component is estimated by first obtaining a slice of the polar power spectrum along the expected mean orientation of the detail component, followed by a peak fitting procedure. By default, the slice contains the elements of the polar power spectrum whose orientations are within ten degrees from the dominant orientation estimated during the previous step. This is the case for typical structures, which exhibit intensity variations along orientations that are orthogonal to the longitudinal axis of the structures. Because orientations in the power spectrum reflect the orientation of intensity variation in the image, spectral signals originating from such structures are offset by 90 degrees relative to the orientation of the structures themselves. By default, this offset is corrected during the analysis (see Sections 2.2.3. and 3.3 for details). For structures exhibiting intensity variations along their longitudinal axis, such as myofibrils and their characteristic striated intensity patterns, the correction should not be made, as the orientations seen in the power spectrum directly correspond to the orientation of the structures of interest. The selection between these two options can be accomplished by adjusting the analysis settings accordingly. In the latter case, the spectral slice is obtained along orientations within ten degrees from the orientation orthogonal to the dominant low frequency orientation. The default tolerance of ten degrees can also be adjusted by the user by modifying the value of the angular tolerance parameter ε, but the exact value is not critical for the analysis (see Section 5.6). Averaging over the selected orientation bins is then performed to obtain a one-dimensional vector of spectral values.

Next, the values of corresponding elements of the background spectrum are subtracted from the one-dimensional spectral slice. A Gaussian peak is then fitted to this background subtracted spectrum using a non-linear least squares fitting routine. If the semi-

automatic detail component detection setting is enabled and the user supplies a guess for the wavelength range of the detail component, spatial frequencies corresponding to the user-specified limits are used to constrain the peak fitting procedure (see Section 3.3). In this case, the mean of the Gaussian peak is constrained between the user-specified limits and the standard deviation of the Gaussian is constrained between ¼ and ½ of the spatial frequency range defined by the limits. Peak height is constrained between zero and the actual maximum value of the background subtracted spectrum between the user-specified limits. If the fully automatic setting is used and no guesses are available, the highest peak between lower and upper cutoff frequencies is searched. The peak is then treated as if it were a perfect Gaussian to find locations at distances corresponding to two standard deviations from the peak location. In other words, nearest locations at both sides of the peak with values less than e^2 times the peak height are searched. The spatial frequencies corresponding to these locations are then used to constrain the peak fitting procedure as in the case of user-specified limits. In both cases, the initial values of all three parameters are set halfway between the lower and upper limits. The mean $\bar{f}$ and standard deviation σ_f produced by the fitting procedure, reflecting the estimated spatial frequency range of the detail component, are then stored.

After the expected mean orientation $\bar{\theta}$ and spatial frequency range of the detail component have been estimated, an initial estimate of the spread of the orientation distribution σ_θ is computed. This is accomplished by calculating the angular standard deviation of the polar power spectrum at spatial frequencies within two standard deviations from the mean of the Gaussian peak. The values obtained for $\bar{f}$ and σ_f during the previous peak fitting step are used here. The angular standard deviation is computed with the corresponding function from the CircStat toolbox (Berens 2009).

3.6.6 Extraction of the detail component

The detail component extraction procedure is inspired by a Bayesian classifier. The previously formed probabilistic form of the polar power spectrum $Z(f,\theta)$ is multiplied by a two-dimensional truncated Gaussian mask $K(f,\theta)$ and the resulting values $X(f,\theta)$ are interpreted as probabilities. Each value now represents the probability that the corresponding element of the power spectrum belongs to the detail component spectral region $P_D(f,\theta)$. This should be understood as a purely methodological concept, since the values do not necessarily represent *bona fide* probabilities. If expressed within the Bayesian classifier framework and terminology, the role of the probabilistic power spectrum corresponds to the data or observations while the Gaussian mask is used similarly to a prior. In keeping with the Bayesian classifier interpretation, the values resulting from the multiplication step are treated as if they were posterior probabilities. Spectral elements with values exceeding 0.5 are therefore accepted as part of the detail component spectral region, as their 'probability' of belonging to the detail component is higher than the

'probability' of the complement case of not belonging to the detail component (*i.e.*, probability of belonging to the background).

The values estimated during the previous analysis step are used to initialize the two-dimensional Gaussian mask $K_0(f,\theta)$. The standard deviation of the Gaussian mask along the spatial frequency axis is set to:

$$\sigma_{K,f} = \frac{4\,\sigma_f}{2\sqrt{2\ln 2}} \tag{39}$$

This somewhat arbitrary decision is motivated by the fact that the full width at half maximum (FWHM) of a Gaussian can be expressed as:

$$FWHM = 2\sqrt{2\ln 2}\,\sigma \tag{40}$$

and since the maximum value of the Gaussian in this case is one, the expected probability of belonging to the detail component exceeds or equals the probability of belonging to the background at all spatial frequencies within the FWHM limits. To select a suitable range of spatial frequencies within which this could be expected, the FWHM is set to:

$$FWHM = 4\,\sigma_f \tag{41}$$

which leads to Equation (39). This range of spatial frequencies, constrained within two standard deviations from the mean of the fitted Gaussian model, is selected simply based on the practical consideration that it includes most (approximately 95%) of the values of the Gaussian peak model while excluding the extreme parts of the tails. The standard deviation of the Gaussian mask along the orientation axis is initially set to:

$$\sigma_{K,\vartheta} = \sigma_\theta\, E \tag{42}$$

where prior size E is an adjustable parameter with a default value of 0.5. The reason for incorporating the parameter E is to allow iterative inflation of the detail component spectral region starting from an initial smaller region, in a manner resembling segmentation algorithms based on region-growing strategies. This avoids the inclusion of non-contiguous spectral regions which are likely to originate from multiple different sources in the image. Initially, the Gaussian mask is centered at the location $(\bar{f}, \bar{\theta})$, as these values represent the expected mean orientation and mean spatial frequency of the detail component. Moreover, the Gaussian function is truncated to assign zero prior probability to spectral regions outside the specified spatial frequency cutoffs. The truncated Gaussian mask with the initial parameter values can thus be expressed as:

$$K_0(f,\theta) = \begin{cases} e^{-\left(\frac{(f-\bar{f})^2}{2\sigma_{K,f}^2} + \frac{(\theta-\bar{\theta})^2}{2\sigma_{K,\theta}^2}\right)}, & f_L \le f \le f_H \\ 0, & \text{otherwise} \end{cases} \quad (43)$$

After the initial mask has been formed, the initial spectral ROI $P_{D,0}$ representing the detail component is obtained by first computing the 'posterior' probabilities of each spectral element as:

$$X_0(f,\theta) = Z(f,\theta)\, K_0(f,\theta) \quad (44)$$

and then accepting all elements of the power spectrum for which $X_0 > 0.5$:

$$P_{D,0}(f,\theta) = \begin{cases} P_{WOSA}(f,\theta), & X_0(f,\theta) > 0.5 \\ 0, & \text{otherwise} \end{cases} \quad (45)$$

After $P_{D,0}$ has been obtained, the process proceeds in iterative manner. By default, the maximum number of iterations N_{max} is 10 and the convergence threshold τ is one degree, but these values can also be adjusted by the user. The iterations $i = 1 \ldots N_{max}$ proceed in the following steps:

1. Calculate mean orientation $\bar{\theta}_i$ and angular standard deviation $\sigma_{\theta,i}$ of the current spectral ROI $P_{D,i}$. Use $P_{D,0}$ during the first iteration.
2. Calculate the change in mean orientation $\Delta\bar{\theta}$ and change in angular standard deviation $\Delta\sigma_\theta$ relative to the previous values. If $\Delta\bar{\theta} < \tau$ and $\Delta\sigma_\theta < \tau$ or $i \ge N_{max}$, stop and return $P_{D,i}$.
3. Get updated mask K_i according to Equation (43) by substituting $\bar{\theta}$ with $\bar{\theta}_i$ and $\sigma_{K,\vartheta}$ with $\sigma_{\theta,i}$.
4. Get updated 'posterior' probabilities X_i according to Equation (44) by substituting K_0 with K_i.
5. Get updated spectral ROI $P_{D,i+1}$ according to Equation (45) by substituting X_0 with X_i.
6. Go to step 1.

3.7 Analysis of spectral components

After the mixed and detail spectral component have been extracted from the polar power spectrum, their spectral elements are averaged over all spatial frequencies and all orientations to obtain one-dimensional distributions of orientations and spatial frequencies, respectively. For the orientation distribution, summary statistics are computed using the functions provided by the CircStat toolbox (Berens 2009). The spatial frequency distribution is first transformed into a wavelength distribution, which in turn reflects the size distribution of target structures, and descriptive statistics are then computed. The wavelengths are given in micrometers by default, but the preferred unit can be selected by the user (see Section 3.3).

3.7.1 Orientation analysis

For each of the two spectral components, the distribution of orientations, normalized to assume the form of a probability density, is obtained as:

$$\Psi(\theta) = \frac{\sum_{i=1}^{N_f} P_c(f_i, \theta)}{\sum_{j=1}^{N_\theta} \sum_{i=1}^{N_f} P_c(f_i, \theta_j)} \tag{46}$$

where f_i and θ_j stand for the *i*th spatial frequency bin and *j*th orientation bin, respectively. The dummy variable $P_C(f,\theta)$ is substituted either with the mixed component $P_M(f,\theta)$ or the detail component $P_D(f,\theta)$ spectral ROI. Angle doubling is performed to account for the axial nature of the data and the mean orientation angle is afterwards divided by two (Berens 2009). The following summary statistics (see Section 2.3.5) are calculated using the corresponding functions of the CircStat toolbox (Berens 2009):

- Mean orientation
- Circular variance
- Angular standard deviation
- Circular skewness
- Circular kurtosis

3.7.2 Wavelength analysis

For each spectral component, the distribution of wavelengths is obtained by first converting the polar power spectrum into wavelength form. As the power spectrum is a density distribution function, expressing it in terms of wavelength λ rather than spatial frequency f is not simply a matter of making the substitution $f = 1/\lambda$ (Soffer & Lynch 1999). In addition to this substitution, the values in each spatial frequency bin need to be multiplied with a correction factor of $1/\lambda^2$ (or equivalently f^2) during the transformation. This correction is simply based on the requirement that total energy be conserved. The power spectrum as a function of wavelength can thus be expressed as:

$$P_{WL}(\lambda, \theta) = \frac{1}{\lambda^2} P_c\left(\frac{1}{\lambda}, \theta\right) \tag{47}$$

where $P_C(f,\theta)$ is substituted either with the mixed component $P_M(f,\theta)$ or the detail component $P_D(f,\theta)$ spectral ROI. The wavelength distribution, normalized to assume the form of a probability density, is then obtained as:

$$\Lambda(\lambda) = \frac{\sum_{j=1}^{N_\theta} P_{WL}(\lambda, \theta_j)}{\sum_{i=1}^{N_f} \sum_{j=1}^{N_\theta} P_{WL}(\lambda_i, \theta_j)} \tag{48}$$

where λ_i and θ_j stand for the *i*th wavelength bin and *j*th orientation bin, respectively. Linear interpolation using a sufficient number of data points to preserve spectral resolution even at the shortest wavelengths is then performed to obtain equidistant spacing of data points. The following summary statistics are then calculated:

- Mean
- Median
- Mode
- Standard deviation

3.8 Exporting results

After finishing an analysis run, the estimated orientation and wavelength distributions can be examined as plots via the CytoSpectre GUI together with summary statistics. If segmentation is enabled, allowing analysis of cellular morphology, the corresponding statistics are also shown in tabular form for each cell. The power spectrum can also be examined visually in order to spot interesting spectral features, which may then be extracted by tuning the wavelength cutoffs and/or the expected detail component wavelength range. CytoSpectre allows the user to export the analysis results in a variety of formats. Summary statistics can be easily exported to spreadsheets or plain text files for further study. The analysis settings used to obtain the results are also stored in these files. In addition, the user can export plots of the estimated orientation and wavelength distributions as images for visualization purposes. It is also possible to export the actual orientation and wavelength distribution values, instead of summary statistics, in plain text files, allowing customized analysis pipelines to be developed utilizing external software. Exporting of results is implemented using MATLAB's in-built file I/O utilities.

4. PERFORMANCE EVALUATION METHODS

In this chapter, supporting methodology, used for analyzing the performance of the software, is described. In Section 4.1, the procedures for generating artificial phase contrast and fluorescence microscopy images for simulation experiments are first explained in detail. Unlike real images, the true characteristics of such simulated images are known exactly. It is thus possible to quantify the amount of error associated with the results produced by the software. The use of artificial images also allows running simulations, where image quality is degraded and the capability of the software to handle such cases is then quantified. Real images were also collected to analyze the performance of the software using authentic data. The experimental methods used to obtain real micrographs of cells are presented in Section 4.2.

4.1 Simulation experiments

4.1.1 Generation of artificial phase contrast images

Although methods capable of generating realistic micrographs of cells are available (Lehmussola *et al.* 2007; Rajaram *et al.* 2012), the first set of artificial images was based only on simple Gaussian targets in order to generate a set of very basic images. Such images are useful for obtaining a general understanding of the software performance. For this dataset, artificial images with bright targets on a dark background, resembling clusters of cells in a phase contrast image, were generated using a custom MATLAB script. An image size of 1024 x 1024 pixels, 10X magnification and an image pixel size of 0.68 μm x 0.68 μm were used as the basis of the images. The targets were randomly located 2D Gaussians with varying sizes and orientations. First, several parameter values were obtained for each image by random sampling from uniform distributions. The maximum number of targets per image was sampled from the interval [10, 100]. The mean width of the targets, that is, the full width at half maximum (FWHM) of the Gaussians, was sampled from the interval (10, 50) μm. The standard deviation of the target width was sampled from the interval (0, 20) μm. The aspect ratio of the targets, that is, the ratio of the target length to the target width, was sampled from the interval (1, 5). The mean direction and concentration parameters of the von Mises distribution were sampled from the intervals (0, π) and (0, 100), respectively.

For each image, targets were located one by one, first sampling the center of each target from a uniform distribution and the direction of the target from a von Mises distribution with the given mean direction and concentration parameters using the corresponding

function from the CircStat toolbox (Berens 2009). Directions were converted to orientations by subtracting π from all direction values larger than or equal to π. The width of each target was then sampled from a normal distribution having the previously obtained mean and standard deviation parameters. Widths less than 5 μm were rejected and the sampling was repeated to ensure that each target had an acceptable number of pixels. Target length was then calculated on the basis of the target width and aspect ratio values. Points belonging to the target were defined as the set of pixels located within an elliptical region at most two standard deviations from the center of the Gaussian. If any of these pixels were located in a so called forbidden region (explained below), the target was rejected and the random placement was repeated. Otherwise, the target was accepted. Initially, the forbidden region comprised all points located within ten pixels from the edges of the image. Whenever a target was accepted, all of the pixels belonging to the target were added to the forbidden region, preventing placement of new targets on existing ones. The target placement process was continued until the maximum number of targets had been placed or more than 1000 consecutive iterations were rejected. The images were output in TIF format and the true orientation and size distributions were saved for performance calculations. An example image is shown in Figure 12. The PSD estimate of the example image in polar coordinates is shown in Figure 13. The estimate was obtained and processed similarly to the PSD estimate of the real epi-fluorescence micrograph in Figure 1. In total, 1000 images were generated.

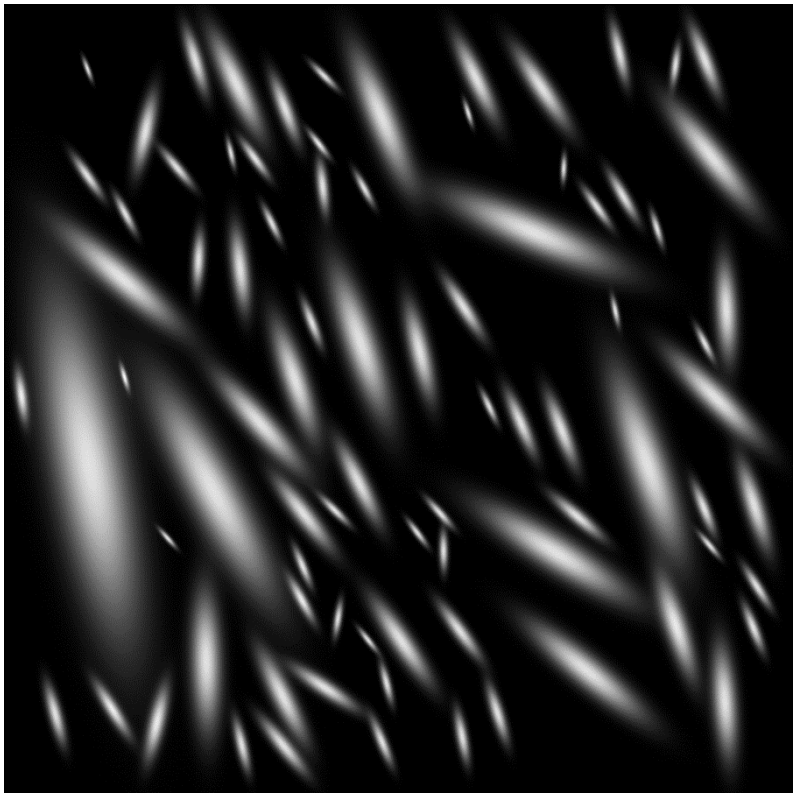

Figure 12. *An example of an artificial phase contrast image.*

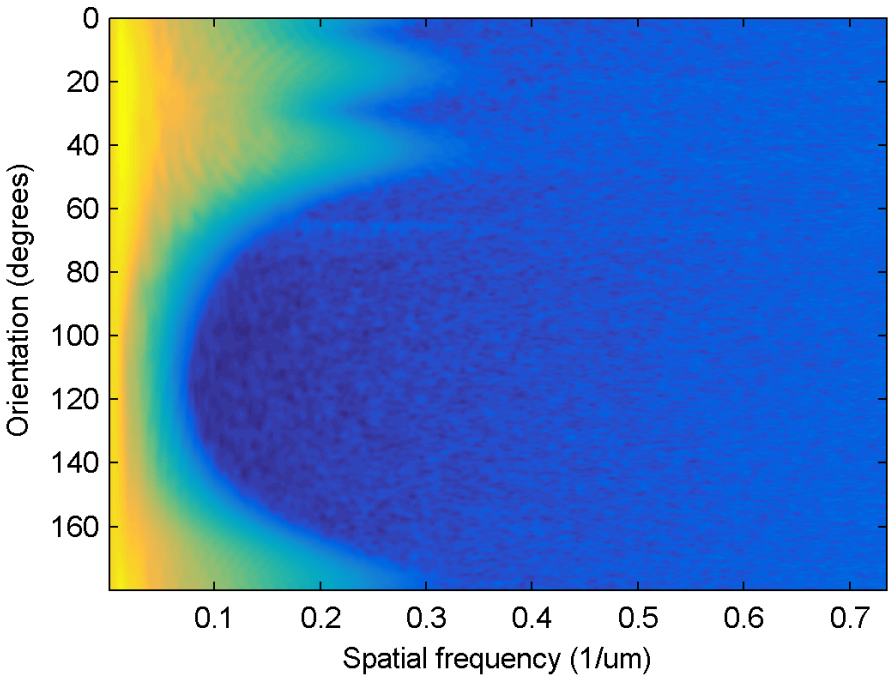

Figure 13. *The PSD estimate of the image in Figure 12 obtained using the WOSA procedure, log10-transformed and expressed in polar coordinates. Bilinear interpolation was used for the polar transform. Only the upper half-circle of the original PSD estimate was retained for the polar representation.*

4.1.2 Generation of artificial fluorescence microscopy images

Another set of artificial images, this time featuring sophisticated intracellular structures as targets, was generated to study the performance of the software in a more realistic setting. Synthetic images of cells with intracellular fibrils were generated using the MATLAB implementation (version 1.0) of SimuCell (Rajaram *et al.* 2012) and custom MATLAB scripts. An image size of 1024 x 1024 pixels, 40X magnification and an image pixel size of 170 nm x 170 nm were used as the basis of the images. For each image, SimuCell was first used to create a single cell with the cytoplasm stained red and the nucleus blue. Random nuclear placement was constrained within two pixels from the midpoint of the image. If any parts of the cell were located less than ten pixels away from the image edge, the cell generation procedure was repeated. For the nuclei, the default nuclear model was used with the following parameters: radius 10 µm, eccentricity 0.5, randomness 0.1, mean intensity 0.5 and intensity standard deviation 0. For the cytoplasms, a nuclear-centric model was used with the following parameters: radius 50 µm, eccentricity 0.7, randomness 0.7, mean intensity 0.3 and intensity standard deviation 0. Default compositing with a container weight of zero was used for all images.

Intracellular fibrillar structures were generated for each cell using a MATLAB script. These structures resembled the myofibrils of hiPSC-CMs in our test set of epi-

fluorescence micrographs, stained with anti-cardiac Troponin T (red), anti-Myosin binding protein C3 (green) and DAPI (blue). In short, each fibril was constructed from a varying number of smaller subunits, each a two-dimensional Gaussian, organized at different orientations and varying intervals along a linear segment within the cell. First, several parameter values were obtained for each cell by random sampling from uniform distributions. The aspect ratio of the subunits, that is, the ratio of the subunit length to the subunit width, was sampled from the interval (1, 5). A fixed subunit width of 1.3 μm, corresponding to the full width at half maximum (FWHM) of the Gaussians along the longitudinal direction of the fibril, was used for all cells. Subunit length, that is, the FWHM of the Gaussians along the transverse direction of the fibril, thus ranged from 1.3 μm to 6.5 μm. Mean and standard deviation of the normally distributed distances between neighboring subunits along a fibril were sampled from the intervals (1.5, 2.5) μm and (0, 0.2) μm, respectively. Minimum spacing between fibrils was sampled from the interval (subunit length, subunit length + 10 μm). Parameters for the von Mises orientation distribution, that is, the mean direction and the concentration parameter were sampled from the intervals (0, π) and (0, 100), respectively.

Fibrils were located one by one, first sampling the midpoint of the fibril from a uniform distribution and the direction of the fibril from a von Mises distribution with the given mean direction and concentration parameters using the corresponding function from the CircStat toolbox (Berens 2009). Directions were converted to orientations by subtracting π from all direction values larger than or equal to π. Fibril length was then sampled from a normal distribution with an initial mean of 50 μm and an initial standard deviation of 10 μm. End points of the fibril were then calculated based on the known midpoint, length and orientation. Points along the entire length of the fibril were subsequently computed using linear interpolation with tenfold upsampling. The resulting region was then thickened by dilating using a square structuring element of ones with size 3x3 pixels. If any of the points of this region were located in a region of the image that is forbidden for fibrils, the fibril was rejected. Otherwise, the fibril was accepted. Initially, the forbidden region consisted of all pixels outside the cell or within the nucleus, dilated by a square structuring element of ones with size equal to twice the subunit length. The dilation was performed to avoid parts of the fibrils extending out from the cell or into the nucleus. When a fibril was accepted, the region consisting of the interpolated points of the fibril was first dilated using a square structuring element of ones with size equal to the given fibril spacing. This region was then added to the existing forbidden region, preventing the placement of new fibrils less than the given fibril spacing away from existing ones. If more than 10 000 consecutive iterations were rejected, the mean and the standard deviation of the fibril length were decreased by 10 % to fit more fibrils into the cell. The fibril placement process was continued until either a maximum number of 20 fibrils were accepted, the mean fibril length became shorter than the mean subunit-to-subunit distance times 15 or more than 100 000 consecutive iterations were rejected.

After the locations of the fibrils had been chosen, the image was formed by locating the fibril subunits. The subunits were located fibril by fibril, starting from one end of each fibril and proceeding towards the other end in steps whose sizes were sampled from a normal distribution with the given mean and standard deviation for subunit-to-subunit distance. The orientation of each Gaussian was sampled from a von Mises distribution having a mean orientation equal to the orientation of the fibril in question, and a concentration parameter equal to the concentration parameter of the fibrils, again using the corresponding function from the CircStat toolbox (Berens 2009). As in the case of the complete fibrils, π was subtracted from directions larger than or equal to π. The resulting image, consisting of Gaussians organized into fibrils, was used to construct the green channel of the synthetic image.

To add some interference to the green channel, resulting from the nucleus and the cytoplasm in the blue and red channels, respectively, an RGB image was first formed directly from the three color channels. The final green channel image was formed by converting the RGB image to grayscale format using MATLAB's *rgb2gray* function, that is, by eliminating the hue and saturation information while retaining only the luminance. As the relative intensities of the red, green and blue channels in the RGB image were 0.3, 1 and 0.5, respectively, the high intensity features of the resulting image represent the target fibrils, while weaker background signals originate from the cytoplasm and the nucleus. The motivation for this mixing operation is that in real images, the objects of interest are often accompanied by other interfering features, which arise, for example, in the case of fluorescence microscopy from unspecific binding of antibodies to different parts of the cell. The final RGB image was then constructed from the green channel image with the added interference, and the original red and blue channel images created by SimuCell. Finally, the red, green and blue channels of the image were rescaled to have relative intensities of 0.8, 1 and 0.9, respectively, in accordance with observations made from our test set images. The synthetic images were output in TIF format. An example image is shown in Figure 14. In total, 1000 images were generated. The PSD estimate of the green channel of the example image, expressed in polar coordinates, is shown in Figure 15. The estimate was obtained using the WOSA procedure and processed similarly to the PSD estimates of the real epi-fluorescence micrograph in Figure 1 and the artificial image in Figure 12. Spectral features similar to those seen in Figure 1 and discussed in Section 2.2.1 are visible also in the PSD estimate of the artificial image, but the lack of noise is evident especially in the high-frequency part of the artificial spectrum.

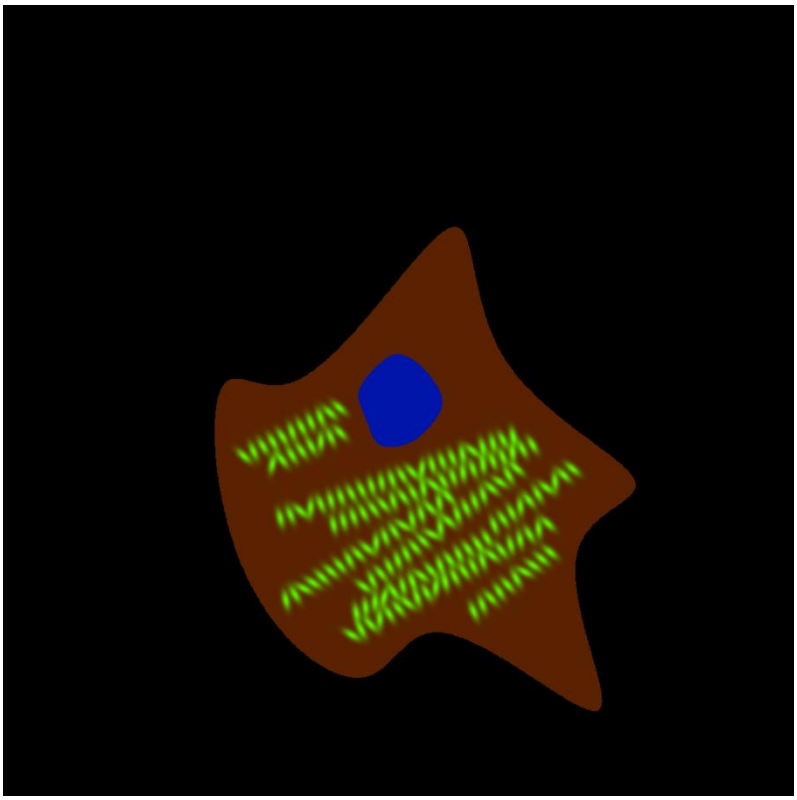

Figure 14. *An example of an artificial fluorescence microscopy image.*

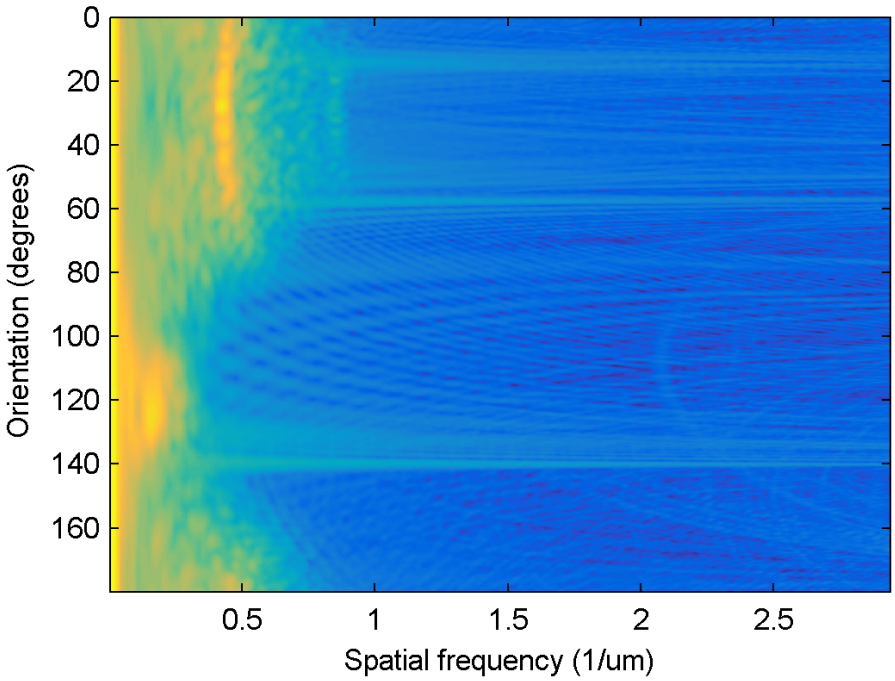

Figure 15. *The PSD estimate of the green channel in Figure 14 obtained using the WOSA procedure, log_{10}-transformed and expressed in polar coordinates. Bilinear interpolation was used for the polar transform. Only the upper half-circle of the original PSD estimate was retained for the polar representation.*

4.1.3 Degradation of artificial images

Synthetic images were blurred by filtering with MATLAB's *imfilter* function and a Gaussian filter having a prescribed standard deviation (Gonzalez & Woods 2001; Boudaoud *et al.* 2014). The standard deviation value was adjusted to obtain images with varying degrees of blurring. The dimensions of the square window were set equal to six times the standard deviation of the Gaussian kernel. Image boundaries were handled by assuming values outside the image to equal the value of the nearest pixel on the image border. Blurring was applied separately to each color channel in the case of RGB images.

Additive zero mean white Gaussian noise and Poisson noise were added using MATLAB's *imnoise* function (Paul *et al.* 2010). The variance of the Gaussian noise was adjusted to obtain images with different amounts of noise. The amount of Poisson noise was adjusted by first scaling the image intensity values down by division with the desired noise level and then scaling the values up with the same noise level value after the noise generation step. Noise was separately generated for each color channel in the case of RGB images. The PSD estimate of the green channel of Figure 14 after the addition of moderate zero mean Gaussian noise (normalized variance of 1 %) is shown in Figure 17. The PSD estimate was obtained and processed similarly to the noise-free image. In qualitative terms, the appearance of the artificial spectral estimate now closely resembles that of the real example in Figure 1.

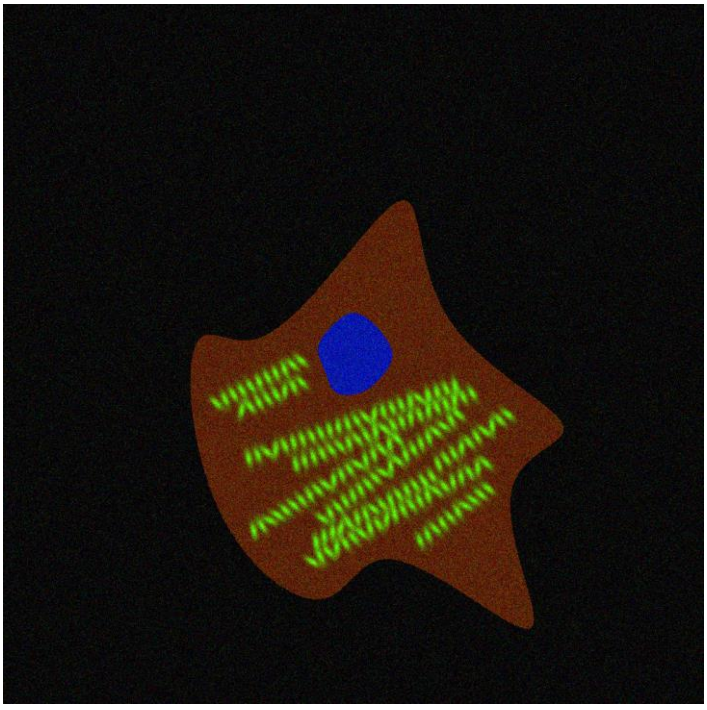

Figure 16. *An artificial fluorescence microscopy image after the addition of zero mean white Gaussian noise with normalized variance of 1 %.*

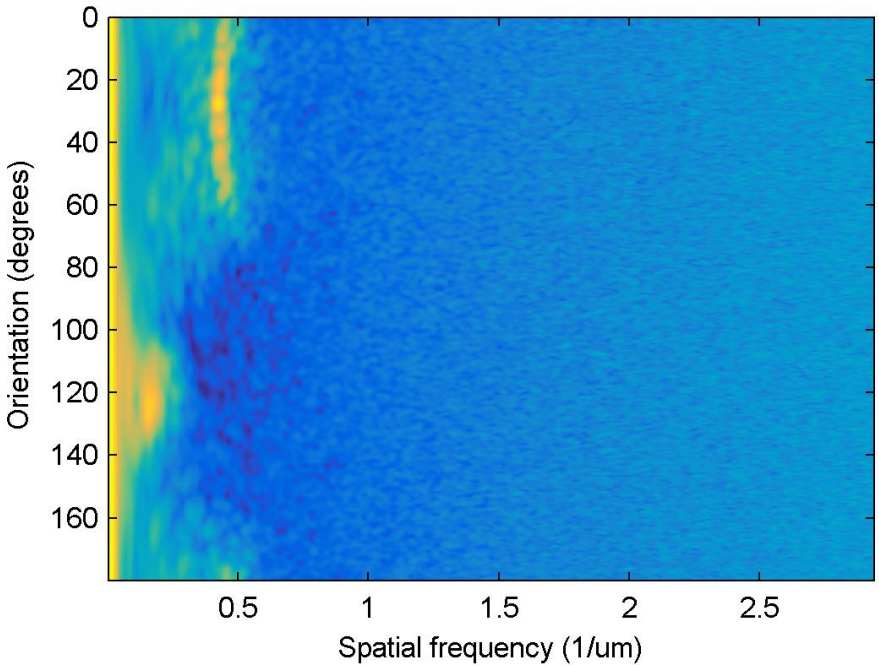

Figure 17. *The PSD estimate of the green channel in Figure 16. The PSD estimate was obtained and processed similarly to the estimate of Figure 15.*

4.2 Cell experiments and collection of real images

All of the laboratory experiments were performed at the University of Tampere by researchers of the Heart Group. The collection of biopsies for generating patient specific hiPSC lines was approved by the ethical committee of Pirkanmaa Hospital District (Aalto-Setälä R08070) and written informed consent was obtained from all the donors. hiPSC-lines had been established by Sendai viral (CytoTune® iPS reprogramming kit, Thermo Fisher Scientific, Waltham, MA, USA) or retroviral transfection of OCT3/4, SOX2, KLF4 and c-MYC (Ohnuki *et al.* 2009). Characterization of hiPSCs included confirmation of mutations by quantitative polymerase chain reaction (qPCR), karyotype analysis, confirmation of gene expression by polymerase chain reaction (PCR) and protein expression by immunocytochemistry. The pluripotency of hiPSC lines was confirmed by the formation of embryoid bodies and teratoma formation in mice. The cells were cultured as described before (Ojala *et al.* 2012).

Differentiation of hiPSCs into cardiomyocytes was carried out by co-culturing hiPSCs with murine visceral endoderm-like (END-2) cells (prof. Mummery, Humbrecht Institute, Utrecht, The Netherlands) (Mummery *et al.* 1991). hiPSCs formed spontaneously beating clusters after 15 days of co-culturing. Beating clusters of cardiomyocytes were cut and isolated with a scalpel and dissociated as described by Mummery and colleagues (Mummery *et al.* 2003). Differentiation of hiPSCs to peripheral sensory neurons was carried out according to the protocols of Goldstein *et al.* (Goldstein *et al.* 2010)

(hiPSC-PSN images 10-15) and Chambers *et al.* (Chambers *et al.* 2012) (hiPSC-PSN images 1-9).

The hiPSC-CMs were imaged using fluorescence microscopy. The cells were stained by double fluorescence protocol: culture medium was aspirated and the cells were washed with 1X phosphate buffered saline (PBS) (0.01M, pH7.4, Dulbecco's PBS, Lonza) two times at room temperature for 5 minutes. Cells were fixed with 4 % paraformaldehyde (PFA, Sigma-Aldrich) for 20 minutes at room temperature. Cells were washed again two times with 1X PBS for 5 minutes. Blocking and permeabilizing was done by applying PBS based solution (10 % normal donkey serum (NDS, Merck Millipore), 0.1 % TritonX-100 (Sigma Aldrich), 1 % bovine serum albumin (BSA, Sigma Aldrich)) for 45 minutes at room temperature. Cells were washed with PBS based solution (1 % NDS, 0.1 % TritonX-100, 1 % BSA) before adding primary antibodies. Primary antibodies of cardiac troponin T (goat IgG, 1:2000, Abcam), alpha actinin (mouse IgG, 1:1500, Sigma Aldrich) and myosin binding protein C (mouse IgG, 1:400, Santa Cruz) were used. Mixtures of primary antibodies were prepared in PBS based solution (1 % NDS, 0.1 % TritonX-100, 1 % BSA). Cells were incubated with primary antibody mixtures in +4°C overnight. Next day, cells were washed three times with 1 % BSA in PBS for 5 minutes before applying secondary antibodies. Alexa Fluor 488 (anti mouse, 1:800, Thermo Fisher Scientific) and Alexa Fluor 568 (anti goat, 1:800, Thermo Fisher Scientific) were used. Secondary antibody mixtures were prepared in 1 % BSA in PBS. Cells were incubated with secondary antibody mixtures for 1 hour in room temperature. Cells were washed three times with 1X PBS for 5 minutes and two times with phosphate buffer working solution (0.01M Na2HPO4/NaH2PO4, pH 7.4) for 5 minutes. Vectashield (Vector Laboratories Inc, Burlingame, CA, USA), containing 4',6-diamidino-2-phenylindole (DAPI) was used to stain nuclei and mount the samples on standard microscope glass. Imaging of immunostained samples was carried out with Zeiss (Carl Zeiss AG, Oberkochen, Germany) AxioScope A1 upright fluorescent microscope and Zeiss AxioCam MRc5 camera. Using 2x2 binning, the effective pixel size of the image sensor was 6.8µm x 6.8µm. The microscope was operated with Zeiss' Zen 2012 software via PC. 20x and 40x air objectives were used for imaging.

The hiPSC-PSNs were images using phase contrast microscopy. Live hiPSC-PSN cells were imaged at the age of around 20 days (depending on differentiation method) with Nikon (Tokyo, Japan) Eclipse TS100 inverted phase contrast microscope with Imperx IGV-B1620M camera. A 10x air objective was used without binning (1x1), with the pixel size of the image sensor being 7.4µm x 7.4µm.

5. RESULTS AND DISCUSSION

In this chapter, results concerning the performance of the software are presented. First, performance was evaluated using artificial phase contrast microscopy and epi-fluorescence microscopy images. These results are presented and discussed in Sections 5.1 and 5.2, respectively. Next, the performance of the software was evaluated by comparing the computational results to manual measurements performed by human experts for real phase contrast and epi-fluorescence micrographs. These experiments are presented in Sections 5.3 and 5.4, respectively. Results from a benchmarking experiment, where the software was compared against an existing method, FibrilTool, are the topic of Section 5.5. The chapter concludes with a presentation of the results of a sensitivity analysis, where the effect of tuning the values of user-adjustable parameters was studied, in Section 5.6.

5.1 Performance evaluation using artificial phase contrast micrographs

In previous studies, the performance of orientation analysis algorithms has typically been assessed using simple computer-generated test images featuring randomly placed Gaussians, ellipsoids or lines (Chaudhuri *et al.* 1987; Karlon *et al.* 1998; Karlon *et al.* 1999; Marquez 2006; Sander & Barocas 2009; Lo *et al.* 2012; Schriefl *et al.* 2012; Boudaoud *et al.* 2014). In line with this traditional approach, 1000 synthetic images featuring randomly placed two-dimensional Gaussian functions as targets, resembling clusters of cells in a phase contrast microscopy image were generated. In the remainder of the text, this set of images is referred to as the 'cell cluster dataset'. Details of the procedure used for generating these artificial images are explained in Section 4.1.1. The test images were analyzed with and without different types of image quality degradations and the estimated values of summary statistics were compared with true values. The values estimated for the mixed component were used in these comparisons, as the sizes of the targets within a single image were not particularly limited to a certain band of wavelengths and the images do not contain any other structures besides the targets of interest. To simulate real-life processes leading to image quality degradation (Paul *et al.* 2010), signal-dependent Poisson distributed noise and additive white Gaussian noise were added into the synthetic images. In addition to noising, blurring was applied by filtering the images with a Gaussian filter to simulate the effects of imperfect focus, leading to loss of sharp details in the images (Boudaoud *et al.* 2014). Each of the degradation operations was performed for all 1000 images and the images were then analyzed to quantify the effects of degraded image quality. The procedures for degrading image quality are

described in Section 4.1.3. The analysis settings used for these experiments are listed in Appendix 1.

First, the original images were analyzed without any degradation in image quality. The results of this experiment are summarized in Table 2. For each parameter, the means and standard deviations of the errors between true and estimated values were first calculated. To allow easier comparisons between different parameters having different numerical ranges, corresponding statistics were also computed for normalized errors, which were obtained by scaling the errors relative to the full range of the true values of each parameter. Moreover, Pearson's linear correlation coefficient between the true and estimated values was computed for each parameter. In the case of the mean orientation, only the absolute errors were analyzed, because normalizing the orientation values defined on a circular scale would have been meaningless. The results show that the mean orientation error is less than two degrees. This level of accuracy is highly likely to be sufficient for most applications since errors on the order of a few degrees are probably overshadowed by biological or technical variation in most experiments. It should be noted that even highly isotropic test images were not excluded from this analysis. In the most extreme cases, the structures are so isotropic that even the true mean orientation is not really indicative of any preferred orientation which can result in large, more or less random errors for such images. The obtained mean orientation error values can therefore be seen as rather pessimistic, representing a worst-case scenario. The other orientation and wavelength parameters, except for circular skewness, exhibit relatively high mean errors, indicating that the numerical values of these parameters should be interpreted with caution. However, the linear correlation coefficients are acceptable for all parameters, ranging from ~0.42 to ~0.95, which indicates that the errors are mostly attributable to systematic bias. This means that the estimated values can still be useful in experiments where a point of comparison is available. For example, values obtained from samples treated with a chemical of interest could be compared to a non-treated control sample. Such a comparative experimental setup including control samples is typical for many biological assays. Thus, the high systematic bias in most of the estimated parameters is not a major issue, as long as the user acknowledges the fact that the numerical values should not be interpreted without a point of reference. Moreover, even in the absence of any systematic bias, interpreting many of the parameters (for example circular variance) directly based only on their numerical values would probably be possible only for a statistics expert with previous experience of working with such measures. In other words, as the answer to the question of whether a particular value of circular variance, for example, is 'high' or 'low' is usually relative and depends on the case at hand, strong linear dependence between true and estimated values is much more important than the amount of systematic bias.

Table 2. *Performance evaluation using artificial phase contrast micrographs. Normalized errors are given relative to the full range of true values of each parameter.*

Parameter	Error, mean ± std	Normalized error, mean ± std	Pearson's r
Mean orientation	1.5534° ± 3.3967°	-	-
Circular variance	0.2384 ± 0.2162	0.2384 ± 0.2162	0.4156
Angular std	0.4042 ± 0.2632	0.2858 ± 0.1861	0.4230
Circular skewness	-0.0001 ± 0.0312	-0.0002 ± 0.0578	0.4686
Circular kurtosis	-0.3930 ± 0.2209	-0.3272 ± 0.1840	0.4980
Wavelength, mean	21.9955 µm ± 5.5363 µm	0.5281 ± 0.1329	0.9532
Wavelength, median	24.4053 µm ± 6.1943 µm	0.5703 ± 0.1447	0.9262
Wavelength, mode	31.7249 µm ± 10.1070 µm	0.7023 ± 0.2237	0.4770
Wavelength, std	2.1006 µm ± 3.6847 µm	0.1104 ± 0.1936	0.4722

Blurry images were obtained by filtering with Gaussian filters having standard deviations ranging from 0.5 to 5 pixels. Errors between true and estimated parameter values obtained at different levels of blurring are plotted in Figure 18A. An example of a moderately blurred image (filter standard deviation 2.5 pixels) is shown in Figure 18D. The effect of blurring on the results was negligible for all of the parameters with both the mean and the standard deviation of the errors staying almost constant. Images with different levels of additive Gaussian noise were obtained by adding noise with zero mean and normalized variance ranging from 0.25 % to 2.5 %. These values correspond to approximate standard deviations of 13 and 40, respectively, on the absolute 8-bit intensity scale (0-255). Errors between true and estimated parameter values obtained in the presence of different amounts of Gaussian noise are plotted in Figure 18B while an example of an image with moderate Gaussian noise (1 % normalized variance) is shown in Figure 18E. The estimates of parameters describing the orientation distribution, especially the mean orientation and circular skewness, are able to tolerate even extreme amounts of Gaussian noise with only a moderate increase in error. The wavelength statistics are more sensitive towards Gaussian noise. In the case of the mean, median and mode of the wavelength distribution, the mean errors actually decrease with increased noise but at the same time, the standard deviation of the error increases, indicating increased random variation. Finally, images with different levels of Poisson noise were obtained as in (Paul *et al.* 2010) by using scaling factors ranging from 2 to 20 during the noise generation process. Errors between true and estimated parameter values obtained in the presence of different levels of Poisson noise are plotted in Figure 18C while an example of an image with moderate Poisson noise (level 10) is shown in Figure 18F. Overall, the effect of Poisson noise on the results is very similar to that of Gaussian noise. It should be noted, that the more extreme amounts of noise and blurring tested in these experiments would rarely be encountered in practice.

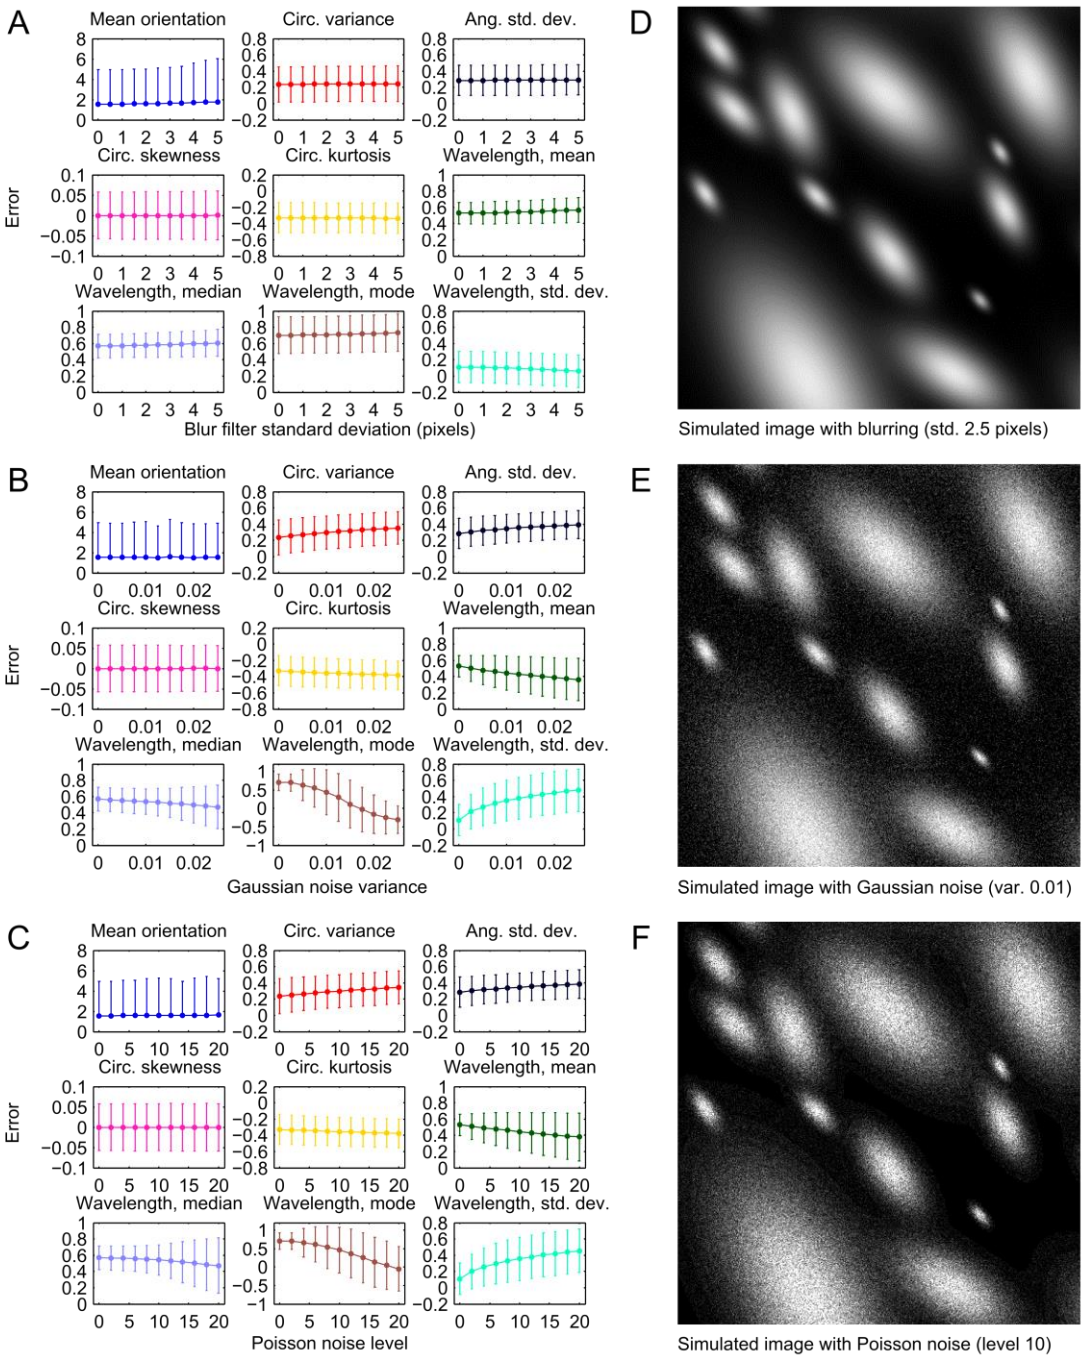

Figure 18. *Results of the performance evaluation using artificial phase contrast images with degraded quality. Errors (mean ± std) between true and estimated values are shown for each estimated parameter for the artificial cell cluster dataset images (N = 1000). For the mean orientation, errors are absolute and given in degrees. For the other parameters, errors normalized relative to the full range of true values are shown. Errors are shown for images blurred with Gaussian filters having standard deviations ranging from 0 to 5 pixels (A), images with additive white Gaussian noise having normalized variance ranging from 0 % to 2.5 % (B) and images with Poisson distributed noise at levels 0 to 20 (C). Examples of an image with blurring (filter std 2.5 pixels) (D), Gaussian noise (normalized variance 1 %) (E) and Poisson distributed noise (level 10) (F) are shown.*

5.2 Performance evaluation using artificial epi-fluorescence micrographs

In addition to the basic evaluation presented in the previous section, the performance of the software was also evaluated with more sophisticated test images containing 2D Gaussian subunits organized into intracellular fibrillar structures. SimuCell (Rajaram *et al.* 2012) was used to generate synthetic images of cells and nuclei, and an algorithm implemented in MATLAB was utilized to construct fibrillar structures within the cells. The image generation process is described in Section 4.1.2. As a result, a dataset of 1000 images with characteristics imitating those of epi-fluorescence microscopy images captured from cells with immunostained intracellular fibrils was generated. In the remainder of the text, this set of images is referred to as the 'fibril dataset'. In the case of this dataset, the targets of interest are constrained to a particular range of wavelengths and there are interfering larger features in the image, namely the nucleus and the cell borders. This made it possible to assess the capability of the detail component to capture the characteristics of the actual target fibrils rather than simply using the mixed component, which would also include contributions from the interfering structures. The image quality degradation experiments performed for the cell cluster dataset, described in the previous section, were repeated for the fibril dataset. Whenever the detail component could not be detected (*i.e.*, the SNR of the signal was too low to allow successful peak fitting during the detail component detection process), the image in question was excluded from the analysis. The total number of excluded images varied between zero and four (corresponding to 0 - 0.4 % of all images) depending on the type and severity of the degradations. The largest number of four exclusions was actually observed for the non-degraded images, which is probably a consequence of the less natural spectra of these images (see Section 4.1). The analysis settings used are listed in Appendix 1.

Table 3. *Performance evaluation using artificial epi-fluorescence micrographs. Normalized errors are given relative to the full range of true values of each parameter.*

Parameter	Error, mean ± std	Normalized error, mean ± std	Pearson's r
Mean orientation	1.8128° ± 7.1196°	-	-
Circular variance	-0.0119 ± 0.1427	-0.0119 ± 0.1427	0.5305
Angular std	0.0149 ± 0.1982	0.0106 ± 0.1401	0.5414
Circular skewness	0.0008 ± 0.0311	0.0015 ± 0.0616	0.3432
Circular kurtosis	-0.0509 ± 0.1808	-0.0453 ± 0.1609	0.6500
Wavelength, mean	0.3886 µm ± 0.5756 µm	0.3809 ± 0.5642	0.1681
Wavelength, median	0.2109 µm ± 0.5283 µm	0.2020 ± 0.5061	0.1978
Wavelength, mode	0.0061 µm ± 0.2002 µm	0.0054 ± 0.1778	0.7817
Wavelength, std	0.4912 µm ± 0.4384 µm	2.2042 ± 1.9671	0.0936

Similarly to the cell cluster dataset, the images of the fibril dataset were first analyzed without degrading the image quality. The results of this experiment are summarized in Table 3 similarly to the cell cluster dataset. The standard deviation of the errors in the estimated mean orientations was higher than for the cell cluster dataset but the mean error was still less than two degrees. These are again pessimistic values due to the inclusion of even the most isotropic cells in the analysis. The other orientation distribution parameters exhibited low errors as well, with all of the normalized mean errors below 5 %. Pearson's linear correlation coefficients between the true and estimated circular variance, angular standard deviation and circular kurtosis were rather high, ranging from ~0.53 to ~0.65, while circular skewness exhibited a moderate correlation of ~0.34. The decrease of systematic bias relative to the cell cluster dataset is probably explained by the ability of the detail component to capture only the relevant part of the power spectrum, which appears to be representative of the target fibrils. At the same time, the contributions of irrelevant parts of the spectrum are decreased. The estimated mode of the wavelength distribution had a mean error of only 6.1 nanometers (with the target size being on the micrometer scale) or 0.54 % of the full range of true values. Given the pixel size of 170 nm, this level of error indicates subpixel accuracy. A high linear correlation coefficient of ~0.78 also indicated that the estimated wavelength mode reflects the true values reliably. The mean, median and standard deviation parameters of the wavelength distribution did not show similarly improved results, however, with weak or nonexistent correlations ranging from ~0.09 to ~0.19. The mode of the wavelength distribution, however, reliably indicates the typical size of the targets of interest. This was expected, as the mode has been used previously, for example, as a measure of the sarcomere length of cardiomyocytes (Pasqualini *et al.* 2015).

The effects of blurring, Gaussian noise and Poisson noise on the errors in the case of the fibril dataset are visualized in Figure 19A-C with example images shown in Figure 19D-F. As opposed to the cell cluster dataset, serious blurring increased the errors of most parameters. This is natural, as the fibril subunits are much smaller and closer to each other than the Gaussian targets of the cell cluster dataset. Sharper images are therefore required to properly capture these detailed structures. However, the errors only start to increase markedly after the blur kernel standard deviation is increased beyond 2-3 pixels while less severe blurring is still tolerated well. Moreover, increased blurring smoothly increases the errors rather than causing sudden, complete breakdown of the detail component detection procedure. Tolerance against even extreme amounts of Gaussian and Poisson noise, on the other hand, is very high as none of the estimated parameters exhibit any clear sensitivity to either type of noise. The adverse effects of severe noise observed in the case of the cell cluster dataset are probably avoided due to the ability of the detail component to again extract only the relevant part of the power spectrum, decreasing the relative contribution of noise.

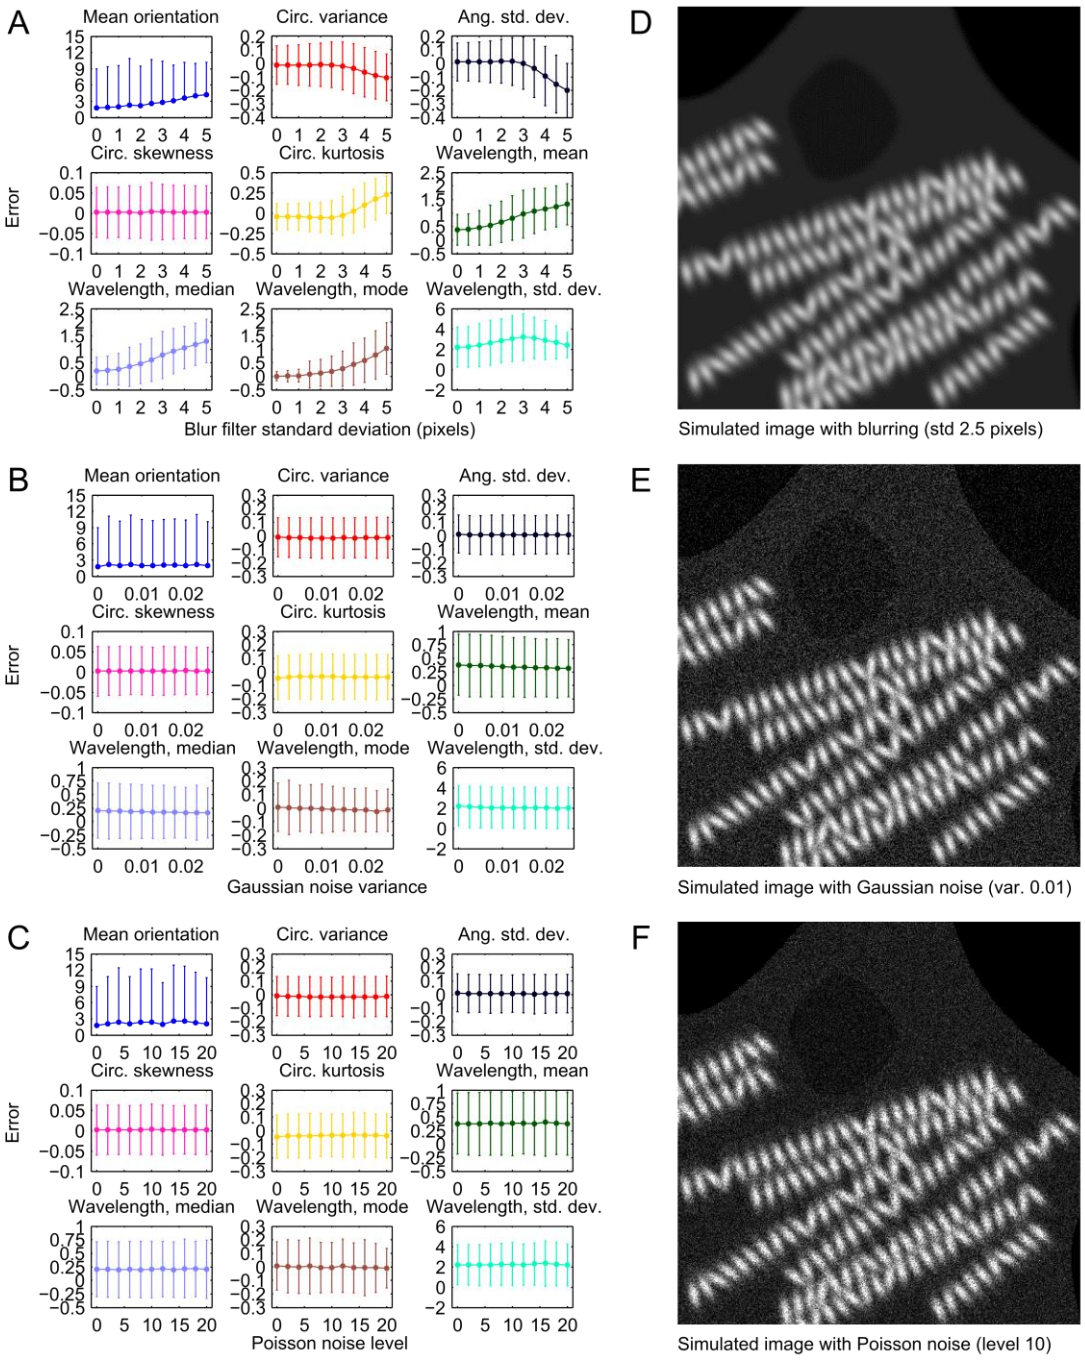

Figure 19. *Results of the performance evaluation using artificial epi-fluorescence microscopy images with degraded quality. Errors (mean ± std) between true and estimated values are shown for each estimated parameter for the artificial fibril dataset images (N = 1000). For the mean orientation, errors are absolute and given in degrees. For the other parameters, errors normalized relative to the full range of true values are shown. Errors are shown for images blurred with Gaussian filters having standard deviations ranging from 0 to 5 pixels (A), images with additive white Gaussian noise having normalized variance ranging from 0 % to 2.5 % (B) and images with Poisson distributed noise at levels 0 to 20 (C). Examples of an image with blurring (filter std 2.5 pixels) (D), Gaussian noise (normalized variance 1 %) (E) and Poisson distributed noise (level 10) (F) are shown.*

5.3 Performance evaluation using real phase contrast micrographs

Three sets of micrographs were used to evaluate the performance of CytoSpectre in the case of real images. All of the test sets were also analyzed subjectively by a panel of human experts to allow comparisons between computational and manual results. The test sets included images obtained of two different cell types using two different microscopy techniques and, in the case of epi-fluorescence microscopy, two different immunofluorescent stainings. Test Set 1 contained images obtained using phase contrast microscopy, depicting hiPSC-PSNs. The results for this set are presented in this section. Test Sets 2 and 3 included images obtained of hiPSC-CMs using epi-fluorescence microscopy and two different staining protocols. The results obtained for Test Sets 2 and 3 are presented in Section 5.4. The experimental methods used to differentiate, culture and image the cells are described in Section 4.2.

Test Set 1 included 15 images at 10X magnification. An example image is shown in Figure 20A. Test Set 1 was analyzed by 13 stem cell researchers from the University of Tampere, based only on visual examination of the images. The experts were asked to give an estimate of the mean orientation of the cells present in the images. The experts also had the option of leaving the mean orientation unspecified in the case of images exhibiting a high degree of isotropy (*i.e.*, lacking any preferred orientation). Some experts also specified multiple main orientations for a single image, in which case circular averaging (see Section 2.3.5) was used to obtain a single mean orientation value for each image. In addition to estimating the mean directions, the experts were asked to rank the images of each set based on the anisotropy of the cell population. This experiment was limited to the mean orientation and anisotropy only, since estimating the wavelength parameters would have been highly challenging based on visual examination and the value of such results would have been questionable. The images were analyzed using CytoSpectre and the computational results were compared with the expert evaluation. The analysis settings used for these experiments are given in Appendix 2.

Since the targets in these images are not clearly limited to a specific range of wavelengths and no significant interfering structures such as nuclei are visible, the mixed component was analyzed using the default settings for cutoff wavelengths. The most isotropic images were excluded from this analysis, as mean orientation is a meaningless quantity for highly isotropic cases. If ¼ or more of the experts were unable to specify the mean orientation for an image, the image was excluded. Ten images were retained based on this criterion. The mean orientations estimated by CytoSpectre were then compared with the manual estimates of each expert. The distributions of absolute differences between mean orientations estimated by CytoSpectre and the human experts are shown in Figure 20D. The red horizontal line represents the median difference, the edges of the blue box correspond to the 25^{th} and 75^{th} percentiles and the whiskers extend to

1.5 times the interquartile range. Red crosses signify values beyond these limits, which are considered outliers. The difference averaged over all images and experts was 15.4° ± 17.4° (mean ± std). This can be seen as an acceptable result in view of psychological studies which indicate that humans may overestimate or underestimate angles by up to 10° (Jastrow 1892; Fisher 1969). There was also considerable variation from expert to expert. The smallest mean difference of 11.9° ± 14.9° was observed with expert 9 and the highest mean difference of 19.8° ± 23.9° was observed with expert 5. This variation underlines the potential benefits of automated analysis and casts doubts on the reliability of manual measurements relying on a single human expert, as is often the case.

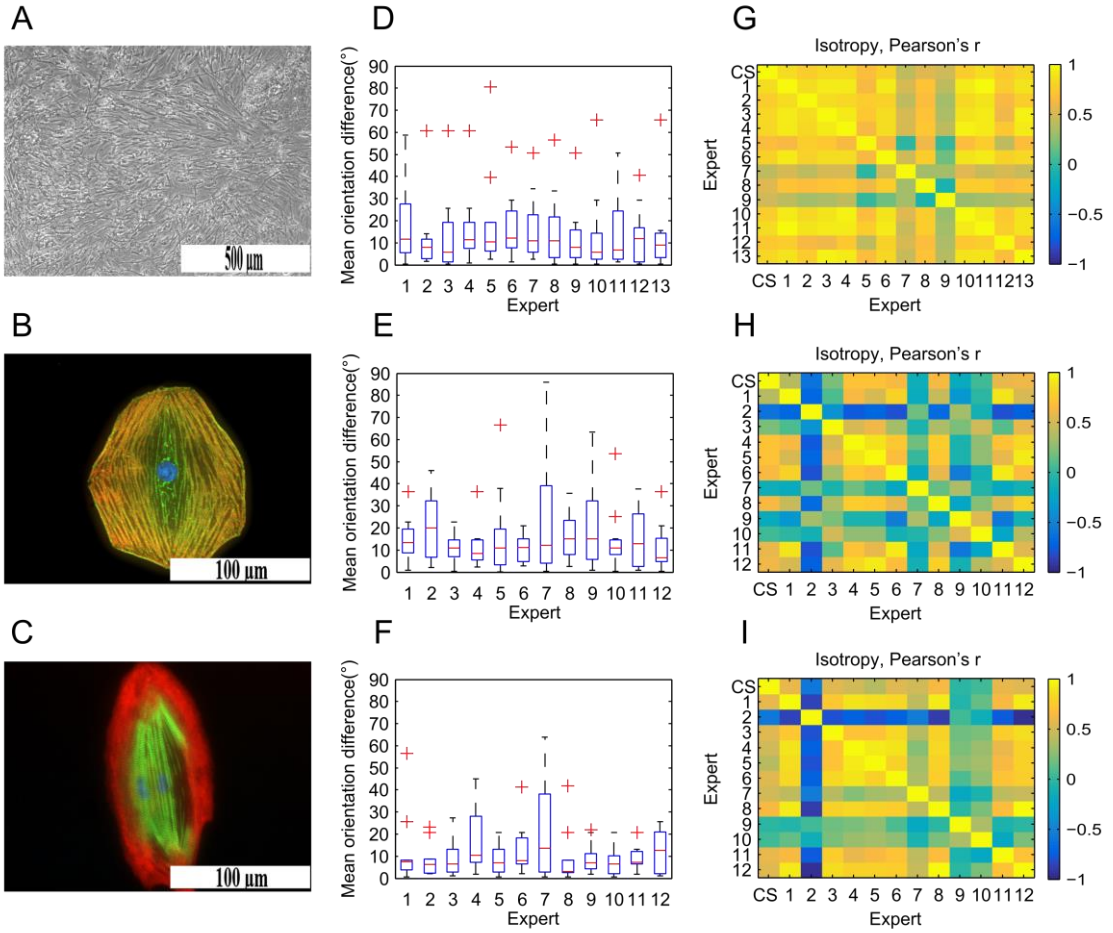

Figure 20. *Results of the performance evaluation using real images. Examples of phase contrast micrographs of hiPSC-PSNs (A) and epi-fluorescence micrographs of hiPSC-CMs with α-actinin (B) or Myosin binding protein C3 (C) staining in the green channel are shown. Distributions of absolute differences in degrees between mean orientations estimated by CytoSpectre and each human expert are shown for Test Set 1 (N = 10) (D), Test Set 2 (N = 11) (E) and Test Set 3 (N = 10) (F). Pearson's linear correlation coefficients between the circular variance values estimated by CytoSpectre and anisotropy rankings specified by human experts are shown as correlation matrices for Test Set 1 (N = 15) (G), Test Set 2 (N = 15) (H) and Test Set 3 (N = 15) (I).*

The correlation between the circular variance values estimated by CytoSpectre and the ranking of images performed by the human experts, based on the perceived degree of

anisotropy of the clusters of neurons, was also analyzed. Pearson's linear correlation coefficient between the values estimated by CytoSpectre and by each human expert as well as between each possible pair of humans was computed. All 15 images were used for this analysis. The results are visualized as a heatmap of correlation coefficients in Figure 20G. The degree of agreement over the isotropy was high both between CytoSpectre and human experts and also between pairs of humans. Pearson's linear correlation coefficient between the values estimated by CytoSpectre and the human experts was 0.72 ± 0.15 (mean ± std), while the correlation coefficients between pairs of human experts were slightly lower with a corresponding value of 0.67 ± 0.24.

5.4 Performance evaluation using real epi-fluorescence micrographs

Two sets of real images of hiPSC-CMs with different immunofluorescent stainings were used to evaluate the performance of CytoSpectre with epi-fluorescence microscopy images. The cell culture, differentiation and imaging protocols are described in Section 4.2. The first set (Test Set 2) included 15 images at 20x and 40x magnifications, stained with DAPI and antibody labels for cardiac Troponin T and α-actinin. An example image from this set is shown in Figure 20B. The second epi-fluorescence micrograph set (Test Set 3) included 15 images at 40x magnification, stained with the nuclear label DAPI and labels for cardiac Troponin T and Myosin binding protein C3. In these images, the striated sarcomere patterns were clearly visible. An example image is shown in Figure 20C. Test Sets 2 and 3 were analyzed by 12 human experts similarly to the images of Test Set 1. That is, the experts were asked to estimate the mean orientation of the myofibrils in each image and rank the images based on the perceived anisotropy of the fibrils.

The myofibrils visible in the images of the Test Set 2 were not limited to any clearly specific band of wavelengths in the power spectrum. The mixed component was therefore analyzed using a relatively broad wavelength range from 1 µm to 5 µm. This selection allowed us to exclude noise and irrelevant structures present at high spatial frequencies (*i.e.*, short wavelengths) as well as structures much larger than the myofibrils, such as the nucleus and the shape of the entire cell. The mean orientations estimated by CytoSpectre were then compared with the manual estimates of each expert. Based on the same criterion used for Test Set 1 to exclude highly isotropic images, 11 images were retained for the mean orientation comparison. The distributions of absolute differences between manual and computational mean orientations are visualized for each expert in Figure 20E. The difference averaged over all experts and all images was 15.6° ± 14.6° (mean ± std). Similarly to Test Set 1, considerable expert-to-expert variation was observed. For example, the difference between CytoSpectre and expert 3 was only 10.6° ± 5.9° while comparison with expert 7 produced a difference of 22.8° ± 24.8°.

The results of the correlation analysis are shown in Figure 20H. The linear correlation coefficient between the values estimated by CytoSpectre and the rankings performed by human experts was 0.30 ± 0.43 (mean ± std) while the linear correlation coefficient between human experts was only 0.18 ± 0.50. While the mean correlation coefficient of 0.30 is rather low, there is considerable variation from expert to expert, similarly to the mean orientation estimates. Many of the human analysts, especially experts 1, 4, 5, 6, 8, 11 and 12 agree with the computational results much more strongly with correlation coefficients of approximately 0.40, 0.71, 0.70, 0.63, 0.65, 0.59 and 0.60, respectively. The remaining experts tend to disagree not only with CytoSpectre but also with each other, highlighting the diversity of ways of perceiving the images and performing scoring. These results seem to indicate that it is more likely for a randomly selected human expert to agree with the computational results than with another randomly selected human expert. While the software was able to estimate circular variance values which corresponded well with the majority of manual rankings, the observed expert-to-expert variation again confirmed the need for objective analysis methods.

Next, the second set of fluorescence microscopy images was analyzed to evaluate the compatibility of the software with different immunofluorescent stainings. The staining used for the images of Test Set 3 highlighted the characteristic striated patterns of the myofibrils (see Figure 20C). The striations reflect the repeating structure of sarcomeres, which have a size of approximately 2 µm and are delimited by protein structures known as Z-disks (Pasqualini *et al.* 2015). In the power spectrum, sarcomeres appear as a prominent peak constrained within a narrow range of spatial frequencies close to 0.5 μm^{-1} (or equivalently within a narrow range of wavelengths close to 2 µm). Based on this information about the target structures, the detail component was analyzed using an expected wavelength range of 1.5 µm to 2.5 µm. Computational results were compared with manual results similarly to the Test Set 2 images. Results of the mean orientation comparison are shown in Figure 20F. Ten images with sufficient anisotropy were retained for this analysis based on the same criterion as with Test Sets 1 and 2. For these images, the difference in computationally and manually estimated mean orientations was 11.3° ± 12.0° (mean ± std) when averaged over all images and experts. The variation from expert to expert was again considerable, with the lowest mean difference of 8.1° ± 6.6° observed with expert 5 and the largest mean difference of 20.1° ± 20.0° observed again in the case of expert 7. Results of the correlation analysis for computationally and manually estimated measures of isotropy, performed for all 15 images, are shown in Figure 20I. These results are very similar to the values observed for Test Set 2, with Pearson's linear correlation coefficient between circular variance values estimated by CytoSpectre and rankings performed by human experts having the value of 0.36 ± 0.36 (mean ± std). The corresponding value between human experts was 0.30 ± 0.54. Again, the relatively low mean linear correlation coefficients are largely explained by the remarkable variation among the human experts. The general pattern of the corre-

lation heatmap is similar to the images of Test Set 2, that is, the same experts tended to agree and disagree with the computational results as with the first dataset.

Overall, the results of these experiments indicate that CytoSpectre is not limited to phase contrast microscopy but is also suitable for epi-fluorescence images. These results also show that the detail component extraction procedure functions properly also in the case of real images and produces results that most human experts agree with relatively well. Moreover, these results support the idea that CytoSpectre is not limited to a particular fluorescent label but can be used in combination with different stainings. Moreover, there are no theoretical or practical reasons to doubt the compatibility of the software with images captured using other 2D microscopy techniques, although only fluorescence and phase contrast images have been systematically evaluated thus far. This experiment also confirmed that the results obtained on a cellular level for the images of peripheral sensory neurons can be generalized to images with different cell types and structures on a different, subcellular scale.

5.5 Comparison with FibrilTool

CytoSpectre was benchmarked against FibrilTool (Boudaoud *et al.* 2014), a recently published orientation analysis plug-in for the popular ImageJ (Schneider *et al.* 2012), introduced in Section 2.3.4. While FibrilTool cannot be used to estimate any wavelength statistics, it was still possible to compare the orientation analysis capabilities of the two methods. FibrilTool allows the estimation of mean orientation and anisotropy of structures within user-specified regions of interest. For each image, a simple ImageJ macro was utilized to select the complete image as a ROI, because CytoSpectre analyzes the whole images as well. The default line width setting of one was used for all images, as suggested in the FibrilTool paper (Boudaoud *et al.* 2014). FibrilTool was applied in this manner for the first 100 images from the artificial cell cluster and fibril datasets first without any degradation and then with moderate amounts of blurring (Gaussian filter standard deviation 2.5 pixels), Gaussian noise (normalized variance 1 %) or Poisson noise (level 10). The total number of images to analyze was thus 400 per dataset, that is, 800 in total for the two datasets. Only these subsets of images were used because FibrilTool requires some manual operation and analyzing 800 images in total, instead of 8000 images, was still feasible. The mean orientations estimated by FibrilTool were adjusted to follow the CytoSpectre angle convention (0°-180°), allowing a direct comparison. In the case of the fibril dataset, FibrilTool was observed to often detect the orientation parallel to the fibril subunits as the main orientation. This orientation is perpendicular to the actual orientation of the fibrils and the mean orientation error could be lowered by applying a correction of 90 degrees. This correction was thus performed for all main orientations estimated by FibrilTool for this dataset. For the cell cluster dataset, this adjustment was not necessary due to the absence of such problematic structures in these images. In contrast to CytoSpectre, FibrilTool does not estimate

the circular variance for each ROI but instead an anisotropy index, which is a measure of anisotropy rather than isotropy. However, as both measures are defined in the range of 0 to 1, an isotropy index for each ROI was simply obtained as the complement of the corresponding anisotropy index (*i.e.*, unity minus the anisotropy index). The resulting values are then directly proportional to circular variance, even though the numerical values of the two measures are not necessarily identical.

Next, the results obtained by CytoSpectre in the manner described in the previous sections were compared with the results produced by FibrilTool. The mixed component parameters were again used in the case of the cell cluster dataset and the detail component parameters in the case of the fibril dataset. For each dataset and type of image degradation, the mean and standard deviations of the absolute main orientation errors in degrees were calculated for both CytoSpectre and FibrilTool. Similarly for each dataset and degradation type, Pearson's linear correlation coefficient between the true circular variance and the estimated measure of isotropy was calculated. The estimated measure of isotropy was circular variance in the case of CytoSpectre and isotropy index in the case of FibrilTool. The resulting values are shown in Table 4 for the cell cluster dataset and in Table 5 for the fibril dataset. In the case of non-degraded or blurred images of the cell cluster dataset, FibrilTool produced main orientation estimates with slightly lower mean error than CytoSpectre, although the difference was less than one degree. Although the relative difference in the errors is large, the practical significance of such small errors is non-existent. However, as observed also in the image quality degradation experiment for both artificial image datasets, the main orientation estimates of CytoSpectre were almost unaffected by blurring or noise. The corresponding estimates of FibrilTool were more sensitive towards image quality degradation, although in the case of the cell cluster dataset, blurring was tolerated well. In the case of added noise and for the fibril dataset in general, the main orientation estimates of CytoSpectre compared favorably against those of FibrilTool. For the isotropy estimates, CytoSpectre obtained higher correlation coefficients than FibrilTool in all of the test cases for both datasets. However, the difference was less pronounced for the non-degraded and blurred cell cluster images.

Table 4. *Performance comparison between CytoSpectre (CS) and FibrilTool (FT) using artificial phase contrast images. The result of the better-performing method is made bold in each case.*

Image quality	CS orientation, error mean ± std (°)	FT orientation, error mean ± std (°)	CS isotropy, Pearson's r	FT isotropy, Pearson's r
Non-degraded	1.7884 ± 3.6128	**1.3904 ± 3.5042**	**0.39844**	0.31512
Blurred	1.8748 ± 3.9533	**1.3432 ± 3.0325**	**0.39645**	0.31258
Gaussian noise	**1.7101 ± 3.4064**	6.4346 ± 10.067	**0.38642**	0.15947
Poisson noise	**1.8325 ± 3.623**	3.3052 ± 3.0916	**0.38501**	0.22644

Table 5. *Performance comparison between CytoSpectre (CS) and FibrilTool (FT) using artificial fluorescence microscopy images. The result of the better-performing method is made bold in each case.*

Image quality	CS orientation, error mean ± std (°)	FT orientation, error mean ± std (°)	CS isotropy, Pearson's r	FT isotropy, Pearson's r
Non-degraded	**2.6261 ± 10.482**	29.673 ± 39.693	**0.54506**	0.19542
Blurred	**3.2969 ± 10.063**	56.685 ± 38.136	**0.15807**	0.10137
Gaussian noise	**2.2876 ± 10.11**	15.785 ± 25.552	**0.45714**	0.2711
Poisson noise	**3.7733 ± 13.852**	19.829 ± 28.444	**0.49203**	0.22492

In summary, the accuracy of CytoSpectre was found to be comparable to that of FibrilTool for the cell cluster dataset without any added noise. However, even if the accuracy of the two methods is similar in this case, FibrilTool requires approximately 20 seconds of manual operation per ROI (Boudaoud *et al.* 2014), which could mean minutes of labor for a single image in the case of multiple ROIs. In contrast, CytoSpectre only requires a few seconds of computation time per image, freeing the user for other tasks during an analysis run. In the presence of Gaussian or Poisson noise, CytoSpectre performed better than FibrilTool in terms of mean orientation error and linear correlation between the true circular variance and the estimated measure of isotropy. In the case of the fibril dataset, CytoSpectre outperformed FibrilTool in all cases. One of the main difficulties for FibrilTool in the images present in the fibril dataset is probably the shape of the fibril subunits, as they have variation in intensity also along the longitudinal orientation of the fibrils. The contributions from intensity gradients along the longitudinal and transverse orientations cannot be easily separated from each other in the spatial domain. Another issue present in many real images that is incorporated into the synthetic images of this dataset is the presence of interfering structures. The nucleus and the exterior of the cell's cytoplasm also contribute to the orientation distribution estimated by FibrilTool and they too cannot be easily separated from the actual targets of interest in the spatial domain. This separation is possible in the frequency domain, allowing CytoSpectre to obtain better results for such images. On the other hand, the cell cluster dataset does not feature such interfering structures and this advantage of spectral analysis does not bring any added benefit for CytoSpectre in that case. It is therefore not surprising, that the results produced by the two methods are very similar for that dataset, although CytoSpectre still appears to be less sensitive to noise. Finally, it should be noted that this comparison is not perfectly fair, since FibrilTool and CytoSpectre are based on quite different philosophies. FibrilTool allows (and requires) users to manually select ROIs and careful selection of these regions could probably improve the results when compared to the approach of selecting the entire image as a ROI. CytoSpectre, on the other hand, is meant to be used in automated fashion for large quantities of images with varying quality. This high-throughput approach is reflected in the experimental setting

of this comparative analysis, as manual selection of ROIs would have been unfeasible for such a large number of images and would have introduced a source of human error.

5.6 Sensitivity analysis of adjustable parameters

In order to test how sensitive CytoSpectre is to the exact values of user-adjustable parameters, the fibril dataset was re-analyzed using different combinations of adjustable parameters. The value of a single parameter was varied at a time within reasonable limits while keeping all other settings at their default values. The value of the maximum iterations parameter N_{max} was not varied, as the number of iterations performed during detail component extraction is mainly controlled by the convergence threshold parameter, as long as the maximum iterations parameter has been set at a high enough value. The parameters and their tested values are given in Table 6 with the default values made bold. The values estimated for the detail component were then compared with the true values. Relative change in RMS error was used to quantify the impact of different parameter values, allowing comparison between the different estimated parameters. The relative change was expressed as a percentage of the RMS error obtained using default settings. The results are plotted for each parameter in Figure 21-Figure 26.

Table 6. *Adjustable parameters and tested values, with default values made bold.*

Parameter	Description	Values tested
Resolution parameter *W*	Controls the tradeoff between spectral resolution and noise.	1 (min. noise), 2 (low noise), **3 (balanced)**, 4 (high resolution), 5 (max. resolution)
Segment length *L*	Controls the length of the spectral segment used during estimation of spectral background and detail component main orientation.	0.02, 0.04, 0.06, 0.08, **0.10**, 0.12, 0.14, 0.16, 0.18, 0.20
Significance level *A*	Significance level for statistical testing during estimation of detail component main orientation.	0.01, 0.02, 0.03, 0.04, **0.05**, 0.06, 0.07, 0.08, 0.09, 0.10
Angular tolerance ε	Controls the number of angular spectral bins used during estimation of detail component wavelength range.	0, 5, **10**, 15, 20, 25, 30, 35, 40, 45
Convergence threshold τ	Controls the stopping threshold for the iterative detail component extraction procedure.	**1**, 2, 3, 4, 5, 6, 7, 8, 9, 10
Prior size *E*	Controls the initial angular width of the Gaussian mask used during the iterative detail component extraction procedure.	0.1, 0.2, 0.3, 0.4, **0.5**, 0.6, 0.7, 0.8, 0.9, 1.0

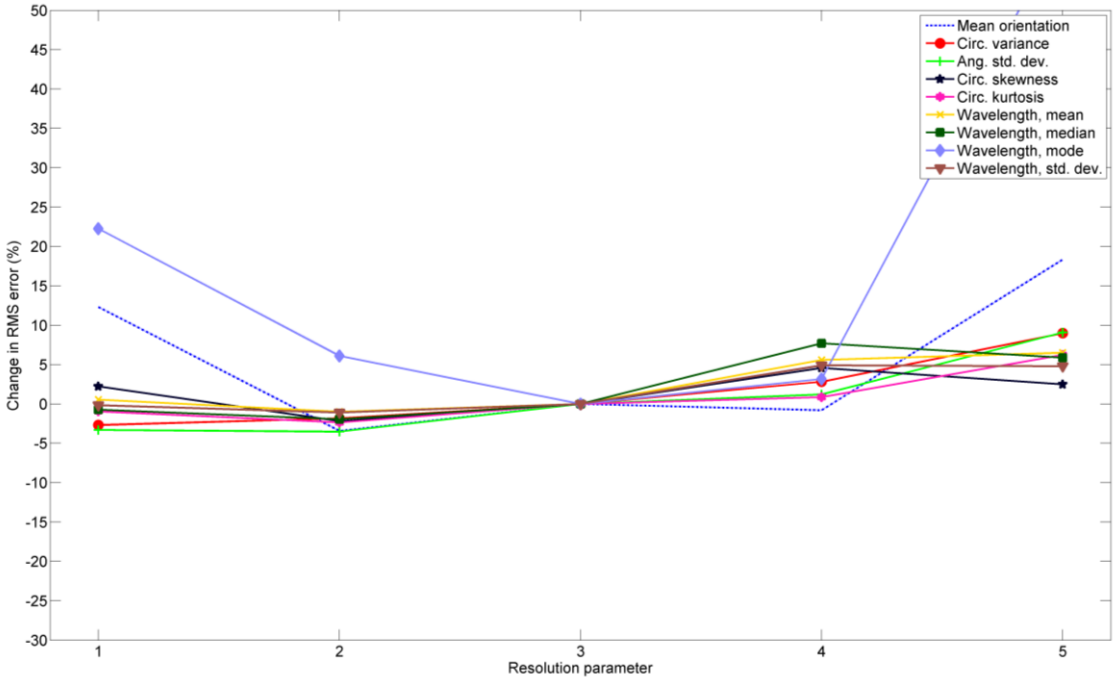

Figure 21. *Effect of the resolution parameter on RMS errors of estimated parameters relative to default setting.*

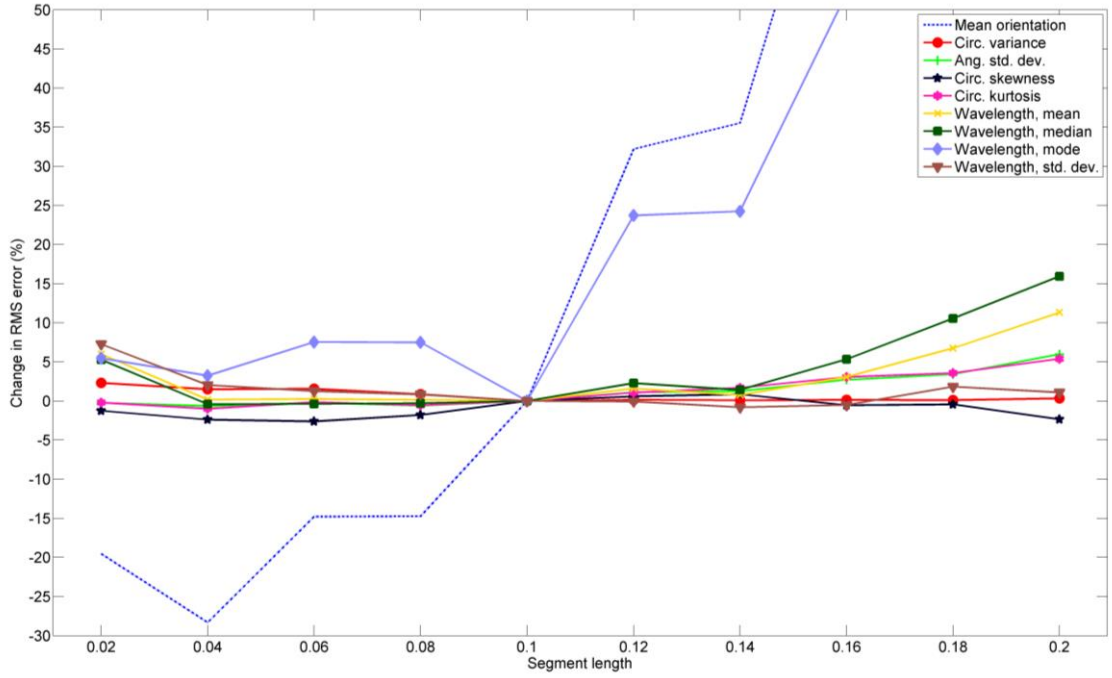

Figure 22. *Effect of the segment length parameter on RMS errors of estimated parameters relative to default setting.*

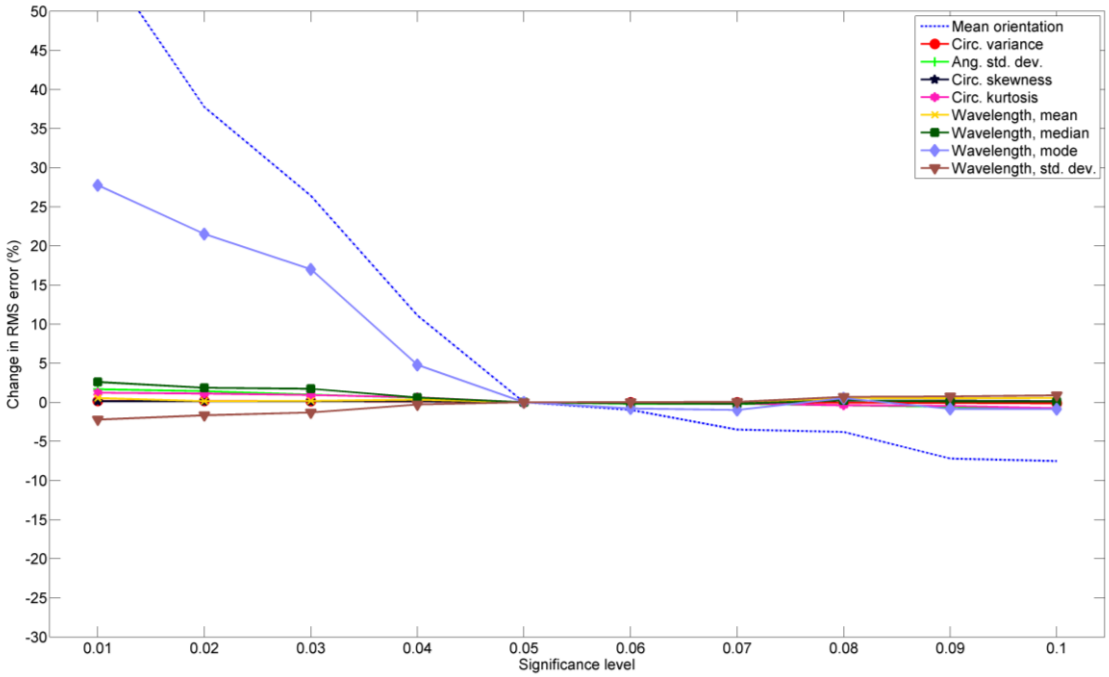

Figure 23. *Effect of the significance level parameter on RMS errors of estimated parameters relative to default setting.*

Based on the sensitivity analysis, most of the adjustable parameters may be safely assumed to have relatively little impact on the results produced by CytoSpectre, as long as the adjustable values are kept within certain limits. The apparently large variations in RMS error observed for the mean orientation in the case of several parameters are mainly explained by the low absolute errors of this parameter. It is doubtful that an increase in the mean orientation error from, for example, two degrees to three degrees would be catastrophic in typical applications, even though the relative change would be 50 %. In most cases, errors of this magnitude can be expected to be overshadowed by biological or technical variation. For most of the adjustable parameters and their tested values, the increases in RMS error were less than 5 % and exceeded 10 % only near the limits of the tested parameter value ranges. The prior size parameter E is the only exception, as increasing its value beyond 0.6 rapidly increased the RMS errors of most of the estimated parameters. However, there was still a relatively stable range of values for this parameter between 0.2 and 0.6, where the increase in RMS error stays below 10 % for all estimated parameters. Based on these results, it appears that tuning of the adjustable parameters by the user is rarely necessary and the default values can be expected to represent an acceptable compromise for most images.

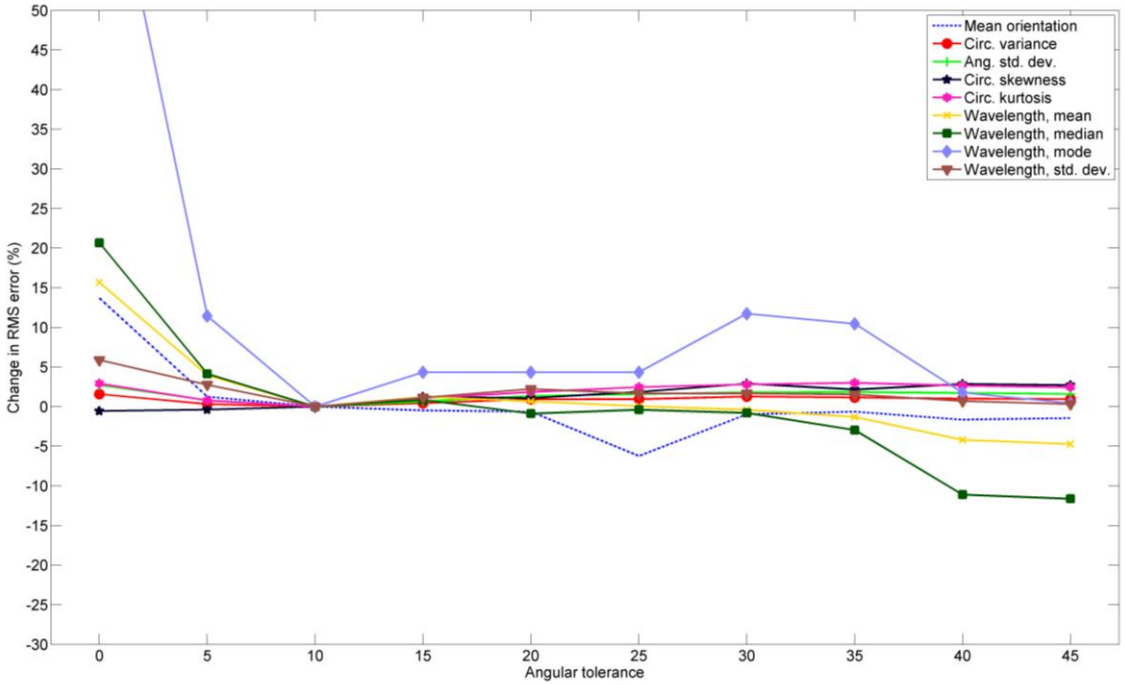

Figure 24. *Effect of the angular tolerance parameter on RMS errors of estimated parameters relative to default setting.*

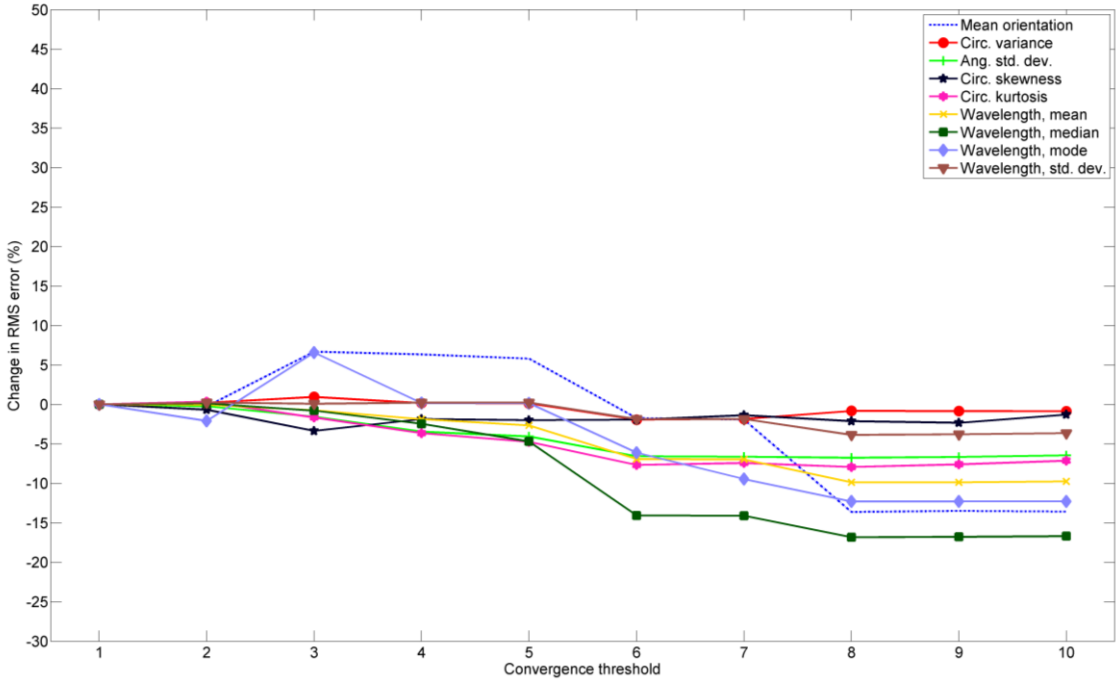

Figure 25. *Effect of the convergence threshold parameter on RMS errors of estimated parameters relative to default setting.*

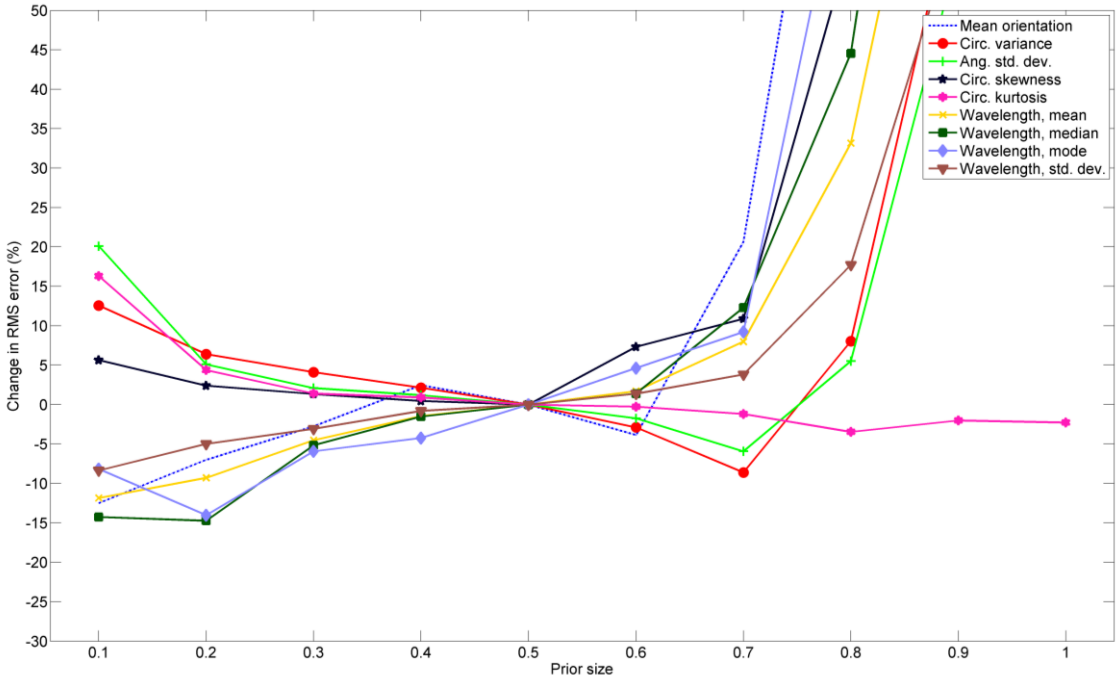

Figure 26. *Effect of the prior size parameter on RMS errors of estimated parameters relative to default setting.*

6. CONCLUSIONS

In this study, a software tool allowing spectral analysis of micrographs, CytoSpectre, was developed. The software was implemented as a standalone MATLAB application, which can be operated via a GUI on basic hardware without prior experience of programming or image processing. Spectral analysis allows users to extract information about the orientation and size distributions of targets within the images. In addition to analyzing complete images, the analysis can be targeted to features of a particular size range in order to obtain information only on these targets of interest while excluding other structures present in the image. This feature is implemented by constraining the spectral analysis to a particular spatial frequency range. Using CytoSpectre, analyses can be performed automatically for large quantities of images and the results can be exported to spreadsheets and text files or plotted as images for further study.

The performance of the software was analyzed by extensive simulations using artificial images with varying types and levels of image quality degradations. In these experiments, high tolerance against realistic amounts of blurring and Gaussian or Poisson distributed noise was observed. The software was benchmarked against FibrilTool (Boudaoud *et al.* 2014), a state-of-the-art orientation analysis tool, and it was found out that for simple targets without any degradation in image quality, CytoSpectre offers accuracy comparable to that of FibrilTool with the important added benefit of rapid, fully automated operation. For images with degraded quality or more realistic targets, CytoSpectre was superior to FibrilTool also in terms of accuracy. The performance of CytoSpectre in the case of real images was evaluated using epi-fluorescence microscopy images of stem cell derived cardiomyocytes and phase contrast micrographs of stem cell derived peripheral sensory neurons. Comparison with manual measurements performed for these images by a panel of stem cell biologists showed that CytoSpectre is not limited to a single microscopy technique or cell type and is able to produce results which are in agreement with the majority of human experts. Furthermore, significant expert-to-expert variation was observed in this experiment, underlining the need for automated computational tools such as the one developed in this study. Finally, the effects of tuning the values of several user-adjustable parameters were quantified, and the software was found to perform consistently across a wide range of parameter values. This result implies that fine-tuning of these parameters by the user is not necessary in most cases, which further improves the user-friendliness of the software.

Hopefully, CytoSpectre will be useful in various applications where the orientations and/or size distributions of biological structures are of interest. The main benefits of-

fered by the software in comparison to previously published methods are the high level of automation, ease of use, versatility in terms of applications and good tolerance against degraded image quality. Since the software is in principle compatible with most 2D images, it could also be useful in biology related or even non-biological applications benefiting from spectral analysis such as (bio)materials research (Sander & Barocas 2009; Ayres *et al.* 2008) or geophysics (Perron *et al.* 2008).

REFERENCES

Ayres, C.E., Jha, B.S., Meredith, H., Bowman, J.R., Bowlin, G.L., Henderson, S.C. & Simpson, D.G. (2008). Measuring fiber alignment in electrospun scaffolds: A user's guide to the 2D fast Fourier transform approach, Journal of Biomaterials Science, Polymer Edition, Vol. 19(5), pp. 603-621.

Bartlett, M.S. (1948). Smoothing periodograms from time series with continuous spectra, Nature, Vol. 161(4096), pp. 686-687.

Berens, P. (2009). CircStat: a MATLAB toolbox for circular statistics, Journal of Statistical Software, Vol. 31(10), pp. 1-21.

Boudaoud, A., Burian, A., Borowska-Wykręt, D., Uyttewaal, M., Wrzalik, R., Kwiatkowska, D. & Hamant, O. (2014). FibrilTool, an ImageJ plug-in to quantify fibrillar structures in raw microscopy images, Nature Protocols, Vol. 9(2), pp. 457-463.

Bray, M.A.P., Adams, W.J., Geisse, N.A., Feinberg, A.W., Sheehy, S.P. & Parker, K.K. (2010). Nuclear morphology and deformation in engineered cardiac myocytes and tissues, Biomaterials, Vol. 31(19), pp. 5143-5150.

Chambers, S.M., Qi, Y., Mica, Y., Lee, G., Zhang, X.-J., Niu, L., Bilsland, J., Cao, L., Stevens, E., Whiting, P., Shi, S.-H. & Studer, L. (2012). Combined small-molecule inhibition accelerates developmental timing and converts human pluripotent stem cells into nociceptors, Nature biotechnology, Vol. 30(7), pp. 715-720.

Chaudhuri, S., Nguyen, H., Rangayyan, R.M., Walsh, S. & Frank, C.B. (1987). Fourier domain directional filtering method for analysis of collagen alignment in ligaments, IEEE Transactions on Biomedical Engineering, Vol. BME-34(7), pp. 509-518.

Daniell, P. (1946). On the theoretical specification and sampling properties of autocorrelated time-series (discussion), JR Stat Soc, Vol. 8, pp. 88.

Davidson, M.W. & Abramowitz, M. (2002). Optical Microscopy, in: Encyclopedia of Imaging Science and Technology, John Wiley & Sons, Inc.

Eliceiri, K.W., Berthold, M.R., Goldberg, I.G., Ibáñez, L., Manjunath, B.S., Martone, M.E., Murphy, R.F., Peng, H., Plant, A.L., Roysam, B., Stuurmann, N., Swedlow, J.R., Tomancak, P. & Carpenter, A.E. (2012). Biological imaging software tools, Nature Methods, Vol. 9(7), pp. 697-710.

Evennett, P.J. & Hammond, C. (2005). MICROSCOPY | Overview, in: Encyclopedia of Analytical Science, 2nd ed., Elsevier, Oxford, UK, pp. 32-41.

Feinberg, A.W., Alford, P.W., Jin, H., Ripplinger, C.M., Werdich, A.A., Sheehy, S.P., Grosberg, A. & Parker, K.K. (2012). Controlling the contractile strength of engineered cardiac muscle by hierarchal tissue architecture, Biomaterials, Vol. 33(23), pp. 5732-5741.

Ferraioli, L., Congedo, G., Hueller, M., Vitale, S., Hewitson, M., Nofrarias, M. & Armano, M. (2011). Quantitative analysis of LISA pathfinder test-mass noise, Physical Review D - Particles, Fields, Gravitation and Cosmology, Vol. 84(12).

Field, D.J. (1987). Relations between the statistics of natural images and the response properties of cortical cells, Journal of the Optical Society of America.A, Optics and image science, Vol. 4(12), pp. 2379-2394.

Field, D.J. & Brady, N. (1997). Visual sensitivity, blur and the sources of variability in the amplitude spectra of natural scenes, Vision research, Vol. 37(23), pp. 3367-3383.

Fisher, G.H. (1969). An experimental study of angular subtension. The Quarterly journal of experimental psychology, Vol. 21(4), pp. 356-366.

Goldstein, R.S., Pomp, O., Brokhman, I. & Ziegler, L. (2010). Generation of neural crest cells and peripheral sensory neurons from human embryonic stem cells. Methods in molecular biology (Clifton, N.J.), Vol. 584, pp. 283-300.

Gonzalez, R.C. & Woods, R.E. (2001). Digital Image Processing, 2nd ed., Addison-Wesley Longman Publishing Co., Inc, Boston, MA, USA.

Hsiao, W.H. & Millane, R.P. (2005). Effects of occlusion, edges, and scaling on the power spectra of natural images, Journal of the Optical Society of America A: Optics and Image Science, and Vision, Vol. 22(9), pp. 1789-1797.

Jahne, B. (2004). Practical Handbook on Image Processing for Scientific and Technical Applications, 2nd ed., CRC Press, Inc, Boca Raton, FL, USA.

Jang, H.-Y., Kim, H.-R., Kang, M.-S., Kim, M.-H. & Zhang, B.-T. (2015). The demand for quantitative techniques in biomedical image informatics, Biomedical Engineering Letters, Vol. 4(4), pp. 319-327.

Jastrow, J. (1892). Studies from the University of Wisconsin: On the Judgment of Angles and Positions of Lines, The American Journal of Psychology, Vol. 5(2), pp. 214-248.

Johnson, P.E. & Long, D.G. (1999). The probability density of spectral estimates based on modified periodogram averages, IEEE Transactions on Signal Processing, Vol. 47(5), pp. 1255-1261.

Jokinen, H., Ollila, J. & Aumala, O. (2000). On windowing effects in estimating averaged periodograms of noisy signals, Measurement, Vol. 28(3), pp. 197-207.

Kamentsky, L., Jones, T.R., Fraser, A., Bray, M., Logan, D.J., Madden, K.L., Ljosa, V., Rueden, C., Eliceiri, K.W. & Carpenter, A.E. (2011). Improved structure, function and compatibility for cellprofiler: Modular high-throughput image analysis software, Bioinformatics, Vol. 27(8), pp. 1179-1180.

Karlon, W.J., Covell, J.W., McCulloch, A.D., Hunter, J.J. & Omens, J.H. (1998). Automated measurement of myofiber disarray in transgenic mice with ventricular expression of ras, Anatomical Record, Vol. 252(4), pp. 612-625.

Karlon, W.J., Hsu, P.-P., Song, L.I., Chien, S., McCulloch, A.D. & Omens, J.H. (1999). Measurement of Orientation and Distribution of Cellular Alignment and Cytoskeletal Organization, Annals of Biomedical Engineering, Vol. 27(6), pp. 712-720.

Khan, M., Xu, Y., Hua, S., Johnson, J., Belevych, A., Janssen, P.M.L., Gyorke, S., Guan, J. & Angelos, M.G. (2015). Evaluation of changes in morphology and function of human induced pluripotent stem cell derived cardiomyocytes (hiPSC-CMs) cultured on an aligned-nanofiber cardiac patch, PLoS ONE, Vol. 10(5).

Kim, K.-M., Kim, S.-Y., Minxha, J. & Palmore, G.T.R. (2011). A novel method for analyzing images of live nerve cells, Journal of neuroscience methods, Vol. 201(1), pp. 98-105.

Lehmussola, A., Ruusuvuori, P., Selinummi, J., Huttunen, H. & Yli-Harja, O. (2007). Computational Framework for Simulating Fluorescence Microscope Images With Cell Populations, Medical Imaging, IEEE Transactions on, Vol. 26(7), pp. 1010-1016.

Lo, W., Chen, W.-L., Hsueh, C.-M., Ghazaryan, A.A., Chen, S.-J., Hui-Kang Ma, D., Dong, C.-Y. & Tan, H.-Y. (2012). Fast fourier transform-based analysis of second-harmonic generation image in keratoconic cornea, Investigative Ophthalmology and Visual Science, Vol. 53(7), pp. 3501-3507.

Mardia, K.V. & Jupp, P.E. (2009). Directional Statistics, John Wiley & Sons, Inc.

Marple, S.L. (1987). Digital spectral analysis: with applications, Prentice-Hall, Englewood Cliffs, NJ, USA.

Marquez, J.P. (2006). Fourier analysis and automated measurement of cell and fiber angular orientation distributions, International Journal of Solids and Structures, Vol. 43(21), pp. 6413-6423.

Mummery, C., Ward-van Oostwaard, D., Doevendans, P., Spijker, R., Van den Brink, S., Hassink, R., Van der Heyden, M., Opthof, T., Pera, M., Brutel de la Riviere, A., Passier, R. & Tertoolen, L. (2003). Differentiation of human embryonic stem cells to cardiomyocytes: Role of coculture with visceral endoderm-like cells, Circulation, Vol. 107(21), pp. 2733-2740.

Mummery, C.L., Van Achterberg, T.A.E., Van den Eijnden-Van Raaij, A.J.M., Van Haaster, L., Willemse, A., De Laat, S.W. & Piersma, A.H. (1991). Visceral-endoderm-like cell lines induce differentiation of murine P19 embryonal carcinoma cells, Differentiation, Vol. 46(1), pp. 51-60.

Murphy, D.B. (2002). Fundamentals of Light Microscopy and Electronic Imaging, John Wiley & Sons, Inc., USA.

Ohnuki, M., Takahashi, K. & Yamanaka, S. (2009). Generation and characterization of human induced pluripotent stem cells. Current protocols in stem cell biology, Chapter 4.

Ojala, M., Rajala, K., Pekkanen-Mattila, M., Miettinen, M., Huhtala, H. & Aalto-Setälä, K. (2012). Culture Conditions Affect Cardiac Differentiation Potential of Human Pluripotent Stem Cells, PLoS ONE, Vol. 7(10), pp. e48659.

Palmer, B.M. & Bizios, R. (1997). Quantitative Characterization of Vascular Endothelial Cell Morphology and Orientation Using Fourier Transform Analysis, Journal of Biomechanical Engineering, Vol. 119(2), pp. 159-165.

Pang, Y., Ucuzian, A.A., Matsumura, A., Brey, E.M., Gassman, A.A., Husak, V.A. & Greisler, H.P. (2009). The temporal and spatial dynamics of microscale collagen scaffold remodeling by smooth muscle cells, Biomaterials, Vol. 30(11), pp. 2023-2031.

Pasqualini, F.S., Sheehy, S.P., Agarwal, A., Aratyn-Schaus, Y. & Parker, K.K. (2015). Structural phenotyping of stem cell-derived cardiomyocytes, Stem Cell Reports, Vol. 4(3), pp. 340-347.

Paul, P., Duessmann, H., Bernas, T., Huber, H. & Kalamatianos, D. (2010). Automatic noise quantification for confocal fluorescence microscopy images, Computerized Medical Imaging and Graphics, Vol. 34(6), pp. 426-434.

Perron, J.T., Kirchner, J.W. & Dietrich, W.E. (2008). Spectral signatures of characteristic spatial scales and nonfractal structure in landscapes, Journal of Geophysical Research: Earth Surface, Vol. 113(4).

Petroll, W.M., Cavanagh, H.D., Barry, P., Andrews, P. & Jester, J.V. (1993). Quantitative analysis of stress fiber orientation during corneal wound contraction, Journal of cell science, Vol. 104(2), pp. 353-363.

Rajaram, S., Pavie, B., Hac, N.E.F., Altschuler, S.J. & Wu, L.F. (2012). SimuCell: A flexible framework for creating synthetic microscopy images, Nature Methods, Vol. 9(7), pp. 634-635.

Rezakhaniha, R., Agianniotis, A., Schrauwen, J.T.C., Griffa, A., Sage, D., Bouten, C.V.C., Van De Vosse, F.N., Unser, M. & Stergiopulos, N. (2012). Experimental investigation of collagen waviness and orientation in the arterial adventitia using confocal laser scanning microscopy, Biomechanics and Modeling in Mechanobiology, Vol. 11(3-4), pp. 461-473.

Ruderman, D.L. (1997). Origins of scaling in natural images, Vision research, Vol. 37(23), pp. 3385-3398.

Sander, E.A. & Barocas, V.H. (2009). Comparison of 2D fiber network orientation measurement methods, Journal of Biomedical Materials Research - Part A, Vol. 88(2), pp. 322-331.

Schneider, C.A., Rasband, W.S. & Eliceiri, K.W. (2012). NIH Image to ImageJ: 25 years of image analysis, Nature Methods, Vol. 9(7), pp. 671-675.

Schriefl, A.J., Reinisch, A.J., Sankaran, S., Pierce, D.M. & Holzapfel, G.A. (2012). Quantitative assessment of collagen fibre orientations from two-dimensional images of soft biological tissues, Journal of the Royal Society Interface, Vol. 9(76), pp. 3081-3093.

Sinnecker, D., Laugwitz, K.-L. & Moretti, A. (2014). Induced pluripotent stem cell-derived cardiomyocytes for drug development and toxicity testing, Pharmacology and Therapeutics, Vol. 143(2), pp. 246-252.

Soffer, B.H. & Lynch, D.K. (1999). Some paradoxes, errors, and resolutions concerning the spectral optimization of human vision, American Journal of Physics, Vol. 67(11), pp. 946-953.

Stoica, P. & Moses, R.L. (2005). Spectral Analysis of Signals, Pearson Prentice Hall, Upper Saddle River, NJ, USA.

The MathWorks, Inc., Fast Fourier Transform, webpage. Available (accessed 20.07.2015): http://se.mathworks.com/help/matlab/ref/fft.html.

Tolhurst, D.J., Tadmor, Y. & Chao, T. (1992). Amplitude spectra of natural images, Ophthalmic and Physiological Optics, Vol. 12(2), pp. 229-232.

Umeno, A. & Ueno, S. (2003). Quantitative analysis of adherent cell orientation influenced by strong magnetic fields, IEEE Transactions on Nanobioscience, Vol. 2(1), pp. 26-28.

Welch, P.D. (1967). The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms, Audio and Electroacoustics, IEEE Transactions on, Vol. 15(2), pp. 70-73.

Wu, G., Mainprize, J.G. & Yaffe, M.J. (2012). Spectral analysis of mammographic images using a multitaper method, Medical physics, Vol. 39(2), pp. 801-810.

Zanella, F., Lyon, R.C. & Sheikh, F. (2014). Modeling heart disease in a dish: From somatic cells to disease-relevant cardiomyocytes, Trends in cardiovascular medicine, Vol. 24(1), pp. 32-44.

APPENDIX 1: ANALYSIS SETTINGS USED IN PERFORMANCE EVALUATION EXPERIMENTS WITH ARTIFICIAL IMAGES

Cell cluster dataset

Magnification	10
Camera pixel size	6.8 µm
Segmentation	Disabled
Target channel	Combined
Rotate mixed component by 90°	Enabled
Rotate detail component by 90°	Enabled
Spectral resolution/noise	Balanced
Wavelength cutoff, low	4 µm
Wavelength cutoff, high	Automatic
Exclude wavelengths	Disabled
Detail component detection	Automatic
Detail component search settings	Default

Fibril dataset

Magnification	40
Camera pixel size	6.8 µm
Segmentation	Disabled
Target channel	Green
Rotate mixed component by 90°	Enabled
Rotate detail component by 90°	Disabled
Spectral resolution/noise	Balanced
Wavelength cutoff, low	Automatic
Wavelength cutoff, high	Automatic
Exclude wavelengths	Disabled
Detail component detection	Semi-automatic: 1 µm (low), 3 µm (high)
Detail component search settings	Default

APPENDIX 2: ANALYSIS SETTINGS USED IN PERFORMANCE EVALUATION EXPERIMENTS WITH REAL IMAGES

Test Set 1

Magnification	10
Camera pixel size	7.4 µm
Segmentation	Disabled
Target channel	Combined
Rotate mixed component by 90°	Enabled
Rotate detail component by 90°	Enabled
Spectral resolution/noise	Balanced
Wavelength cutoff, low	Automatic
Wavelength cutoff, high	Automatic
Exclude wavelengths	Disabled
Detail component detection	Automatic
Detail component search settings	Default

Test Set 2

Magnification	20/40 (depending on image)
Camera pixel size	6.8 µm
Segmentation	Disabled
Target channel	Green
Rotate mixed component by 90°	Enabled
Rotate detail component by 90°	Enabled
Spectral resolution/noise	Balanced
Wavelength cutoff, low	1 µm
Wavelength cutoff, high	5 µm
Exclude wavelengths	Disabled
Detail component detection	Automatic
Detail component search settings	Default

Test Set 3

Magnification	40
Camera pixel size	6.8 µm
Segmentation	Disabled
Target channel	Green
Rotate mixed component by 90°	Enabled
Rotate detail component by 90°	Disabled
Spectral resolution/noise	Balanced
Wavelength cutoff, low	Automatic
Wavelength cutoff, high	Automatic
Exclude wavelengths	Disabled
Detail component detection	Semi-automatic: 1.5 µm (low), 2.5 µm (high)
Detail component search settings	Default